

MICHAEL J. MAZARR

A New Age of Nations

Power and Advantage in the AI Era



January 2026

For more information on this publication, visit www.rand.org/t/PEA3691-14.

About RAND

RAND is a research organization that develops solutions to public policy challenges to help make communities throughout the world safer and more secure, healthier and more prosperous. RAND is nonprofit, nonpartisan, and committed to the public interest. To learn more about RAND, visit www.rand.org.

Research Integrity

Our mission to help improve policy and decisionmaking through research and analysis is enabled through our core values of quality and objectivity and our unwavering commitment to the highest level of integrity and ethical behavior. To help ensure our research and analysis are rigorous, objective, and nonpartisan, we subject our research publications to a robust and exacting quality-assurance process; avoid both the appearance and reality of financial and other conflicts of interest through staff training, project screening, and a policy of mandatory disclosure; and pursue transparency in our research engagements through our commitment to the open publication of our research findings and recommendations, disclosure of the source of funding of published research, and policies to ensure intellectual independence. For more information, visit www.rand.org/about/research-integrity.

RAND's publications do not necessarily reflect the opinions of its research clients and sponsors.

Published by the RAND Corporation, Santa Monica, Calif.

© 2026 RAND Corporation

RAND® is a registered trademark.

Cover: Phonlamaiphoto/Adobe Stock.

Limited Print and Electronic Distribution Rights

This publication and trademark(s) contained herein are protected by law. This representation of RAND intellectual property is provided for noncommercial use only. Unauthorized posting of this publication online is prohibited; linking directly to its webpage on rand.org is encouraged. Permission is required from RAND to reproduce, or reuse in another form, any of its research products for commercial purposes. For information on reprint and reuse permissions, visit www.rand.org/about/publishing/permissions.

About This Paper

In this paper, I examine the sources of national competitive advantage in the emerging era of artificial intelligence (AI). Using insights from earlier RAND work on the societal foundations of national competitiveness, I evaluate how AI might affect those characteristics. I conclude by proposing the broad elements of a U.S. strategy for national competitive advantage in the AI Era.

Center for the Geopolitics of Artificial General Intelligence

RAND Global and Emerging Risks is a division of RAND that delivers rigorous and objective public policy research on the most consequential challenges to civilization and global security. This work was undertaken by the division's Center for the Geopolitics of Artificial General Intelligence (AGI), which is committed to helping decisionmakers understand, anticipate, and prepare to navigate the national security and geopolitical implications of AGI. For more information, visit www.rand.org/geopolitics-ofagi.

Funding

This effort was independently initiated and conducted within the Center for the Geopolitics of Artificial General Intelligence using income from operations and gifts from RAND supporters, including philanthropic gifts made or recommended by DALHAP Investments Ltd., Ergo Impact, Founders Pledge, Charlottes och Fredriks Stiftelse, Good Ventures, Longview, and Coefficient Giving. RAND donors and grantors have no influence over research findings or recommendations.

Acknowledgments

While developing the arguments in this publication, I was lucky enough to benefit from a year of intense dialogue and research from the large community of scholars in the Global and Emerging Risk division at RAND. Among this group, there are computer scientists and AI researchers, former employees of AI labs, political scientists, engineers, experts in military operational art, scientists from a half-dozen disciplines, economists, China Hands, and people who have an assortment of other expertise. All errors in this publication remain my own, but I have learned a tremendous amount from my colleagues, who have been unfailingly generous with their time and expertise.

I am grateful to Jim Mitre, vice president of the RAND Global and Emerging Risks division; Jeff Alstott; and especially Joel Predd, who has headed the Center for the Geopolitics of AGI, for their support, ideas, guidance, and corrections. A more collegial and supportive group of research leaders could not be imagined. They have encouraged disciplined but free-flowing investigation of issues, insisted on collaboration across disciplines, and—in other ways—produced the most exciting intellectual experience I can recall.

The conclusions in this publication rely on three years of analyzing the foundations of societal competitiveness, which I have conducted for the former Office of Net Assessment (ONA) in the Pentagon. I am forever grateful to James Baker, the former director of ONA, for his support and guidance of these studies and others that I have completed for his office. James is one of the most thoughtful and original sponsors—and thinkers—of research that I have ever encountered.

I am also thankful to Colin Kahl, Miles Brundage, Guarav Sett, and Ashwin Acharya, who reviewed parts of several chapters during the preparation of the manuscript, and William Marcellino and Pavneet Singh, who provided full peer review.

Summary

The world stands on the cusp of a defining technological revolution: the emergence of artificial intelligence (AI) as the latest—and, potentially, by far the most wide-ranging and influential—general-purpose technology in human history. Similarly to its precursor, the Industrial Revolution, the resulting changes will play a major role in determining national fates and reshuffling the deck of global power.

In recognition of these trends, the U.S. government has taken crucial steps to ensure leadership in critical components of the AI technology stack. But as important as semiconductor chips, training runs, and data centers are, the United States has been missing a larger truth about AI and national competitiveness. This analysis argues that U.S. policymakers need to begin thinking much more seriously about the broader societal foundations for national advantage in the AI Era. Countries that lead this era will not merely have the best AI models. They will take the necessary steps—including the application of AI to dozens of social goals—to make their societies more competitive. This work argues that, in the end, the competitive challenge of AI is primarily social, not technological.

Approach

My conclusions are grounded in three overlapping courses of research and analysis. First, I reviewed a significant literature on the sources and consequences of the Industrial Revolution and the history of technological revolutions more generally, building on prior RAND work on the competitive qualities of nations. Second, I studied the now burgeoning literature on AI’s possible economic, social, political, and military effects. Third, I used the then-leading public versions of three generative AI models—Claude, ChatGPT, and Gemini—as consultants, asking for their assessment of various issues.

To organize the conception of societal advantage, I relied on a framework from a three-year RAND study for the Office of Net Assessment in the Pentagon, which identified the qualities of societies that were essential to success in long-term rivalries. That study nominated seven major societal characteristics and a list of other important factors that play an outsize role in shaping national fates (which are outlined in Chapter 1). These characteristics provided the scaffolding for this publication; each of seven chapters assesses the intersection of AI and one of these national qualities.

Conclusions and Recommendations

This analysis offers dozens of findings about the ways in which AI will offer tremendous opportunities for national advantage but also disrupt and challenge societies. The analysis highlights four overarching themes:

- **Technological revolutions have powerful geopolitical shock waves.** Some countries prosper, some fall behind, but a broad technological transformation represents one of the most-profound challenges to grand strategy a country can face.
- **Competitive advantage in the AI Era will come not only from mastering the narrow technologies of AI—the data centers, algorithms, models, and semiconductors—but also from the broader societal integration and effects of the technology.** Nations will flourish to the degree that their societies provide fertile soil for the diffusion and application of the new technologies and to the degree that they can control and shape the effects of the transition to sustain healthy, coherent, stable societies. Success in the AI Era is more a societal challenge than a technological one.

- **The AI Revolution will crash into an already unstable social, economic, and political context; from the standpoint of national dynamism and competitiveness, the single most essential question is whether societies use AI to mitigate these rising dangers or allow it to exacerbate them.** The way in which AI interacts with this broader transition—whether it ends up empowering a humane and stable shift to postindustrial patterns of organization or further divides and destabilizes societies going through this evolution—will play a critical role in determining whether AI has a supportive or destructive effect.
- **AI's relationship to human agency not only is a moral, political, and philosophical issue but also carries profound implications for national competitive advantage.** AI's effect on human agency is deeply related to issues of dynamism, coherence, solidarity, willpower, opportunity, intellectual energy, and all the other components of sustained competitive position.

More specifically, I describe the implications of AI for each of the seven societal characteristics essential for national competitiveness. In each of those domains, AI has the potential to offer tremendous new capabilities but also to threaten the social foundations of coherence and dynamism. The analysis discusses each of the characteristics in detail, drawing many discrete lessons. Sample findings include the following:

- **National ambition, willpower, and shared identity are critical to long-term advantage, and their fate will be powerfully shaped by the character of the AI Revolution.** The coming transition can strengthen those qualities, instilling a new sense of national mission, but only if its broad effects on societies tend to empower citizens and create a sense of shared benefit. National coherence is at stake in the AI Revolution.
- **AI models offer the potential for a transformative advance in the scale and nature of individual opportunity.** But it will not have that effect in equitable ways without clear policies that shape its evolution to achieve broad empowerment.
- **The spread of AI models and, in particular, human-interfacing chatbots will have unpredictable and potentially dramatic effects on human social interactions.** Widespread use of chatbot-style tools as alternatives to human engagement could deeply threaten social capital, civil society institutions, and other sources of collective identity and norm-setting. AI models may increasingly act, in some cases, on their own volition to shape the social scene.
- **The fiscal and budgetary effects of AI, through the accelerated growth they generate and other means, could empower nations** by relieving existing debt burdens and opening the space for a set of major new national projects.
- **AI and the future of organizational reform**—a critical priority at a time when effective and efficient governance will be a leading differentiator—are deeply intertwined. AI has the potential to super-charge badly needed streamlining of Industrial Era bureaucratic forms but again could easily have the opposite effect—intensifying citizens' sense of alienation from large organizations and processes—if not consciously managed.
- **AI's effects on the intellectual and epistemological environments will have a significant effect on national competitive advantage.** The opportunity is to generate a new surge of intellectual dynamism and to mitigate the fragmentation of the information environment. The risk is that AI will further corrupt shared information ecosystems and generate cognitive off-loading in which citizens rely on models to do their thinking for them.

Finally, I recommend major initiatives in eight areas to lay the groundwork for social competitiveness in the AI Era. Chapter 12 describes these in detail and offers specific recommendations to pursue each one. They are as follows:

- Build public-sector AI competence.
- Develop relevant talent.
- Catalyze AI applications that widen opportunity throughout society.
- Undertake a national campaign to guarantee autonomous agency.
- Underwrite a new era of intellectual discovery.
- Use AI and targeted laws to improve the information environment.
- Combine AI with institutional reforms to streamline and improve the effectiveness of public-sector bureaucracy.
- Create anticipatory AI foresight and strategy functions.

This broad and ambitious agenda represents nothing less than a menu for a dramatic national transformation, both responding to and employing the emerging tools of AI. The fact that the AI Revolution is arriving at a moment of profound socioeconomic disruption makes the requirement for change even more urgent and fraught. This astonishing new technology offers incredible potential that can be enlisted in such a process of rejuvenation. But we, as a society, must make the determination to use AI thoughtfully and effectively to achieve these results.

Many of these priorities amount to a broad effort to use AI to embolden rather than ruin autonomous human agency. Those societies that channel the AI Revolution to bend its effects in the direction of empowerment, agency, and dignity will do well. Those in which AI piles on top of disempowering and predatory forces and institutions to deprive people of even more agency and dignity will suffer very real long-term competitive disadvantage.

The imperative of national renewal is especially daunting in part because, to accomplish it in an authentic and lasting way, we need initiatives at every level of society, not merely in the form of government action. The RAND work on national competitiveness—while stressing the role of effective public institutions and active states in setting the conditions for competitive advantage—endorses grassroots, bottom-up, experimental, and emergent efforts rather than planned, mandated, and bureaucratized ones. The question that U.S. society confronts isn’t merely whether the U.S. government will respond to the challenges it confronts. It’s whether U.S. society will do so in many independent and mutually supporting ways.

To set the context for such a broad renewal, U.S. leaders need to develop a vision of success and identify a first set of actions that would set up U.S. society for competitive advantage in the many different ways that the AI Revolution could unfold. This analysis is designed to offer a framework for understanding this challenge and an initial set of practical policy ideas to get started on this profound agenda.

Contents

About This Publication	iii
Summary	v
Figures and Tables	xiii
CHAPTER 1	
Introduction	1
The Argument	3
The Gateway to a New AI Era	5
Adding to a Wave of Change: AI and Major Trends	6
AI and the Future of Human Agency	7
Scope and Approach	9
The Organizational Scheme: A Framework for Understanding National Competitive Advantage	10
Coming to Grips with a Historical Transformation	13
CHAPTER 2	
The Artificial Intelligence Revolution: Definitions and Prospects	17
Defining AI	17
A Swarm Rather Than a Singularity	18
Defining Artificial General Intelligence	20
The Possibility of an Escalating Feedback Loop	22
The Essential Priority: AI Alignment	23
AI's Prospects: A Great Deal Is About to Change	25
Beyond the Incremental Effects: A More Radical Future	30
CHAPTER 3	
The Character of the Coming Artificial Intelligence Revolution: Three Scenarios	39
The Surprising Deliberateness—and Social Cost—of the Industrial Revolution	40
The Coming AI Revolution: Three Scenarios	41
The Implications for Strategy	48
CHAPTER 4	
Lessons from Previous Technological Revolutions	55
The Foundations of Competitive Advantage in a Technological Revolution	55
The Costs and Risks of Technological Revolutions	57
Technological Revolutions and the Nature of Competitive Advantage	59
Summary: The Lessons of History	61
CHAPTER 5	
National Ambition and Willpower	67
Defining National Ambition and Willpower	68
A Jolt to National Self-Confidence	69
The Capacity for National Projects	70
How It Could Go Wrong: Indolence, Hubris, and Overreach	70
Summary: Rolling the Dice on Human Motivation	73

CHAPTER 6	
Unified National Identity	77
AI and National Identity: Baseline Implications	78
A Society Remade	79
AI Actors Manipulating the Social Scene	82
The Risks of AI-Empowered Fragmentation	83
Overall Implications: National Solidarity in the Crosshairs	86
CHAPTER 7	
Shared Opportunity	91
Defining Shared Opportunity and Its Value	92
The Enlightenment Ideal and the AI Future	93
Making AI Serve Human Agency Will Take a Profound Act of Collective Will	94
Risks: Fading Incentives to Seize Opportunities	96
The Peril of Algorithmic Tyranny	98
AI, the Job Killer	98
The Risk of an Oligarchy of Opportunity	100
Summary: The Potential for a New Enlightenment ... or a Stifling Oligarchy	101
CHAPTER 8	
An Active State	107
Defining the Characteristic of an Active State	108
A Leading Imperative: Good Governance	109
AI and the Active State	111
Planning and Decisionmaking: AI's Benefits and Limitations	113
Hollowing Out Rather Than Empowering ... or Empowering Too Well	114
Summary: Using AI to Revitalize the Active State	116
CHAPTER 9	
Effective Institutions	119
Defining the Characteristic of Effective Institutions	120
AI and Institutions: Efficient Organizations	122
AI and Organizational Fairness	123
A Leading Danger: Supercharging Bureaucracy and Arbitrary Rule by Algorithms	124
AI and the Social Role of Institutions: The Problem of Social Capital	126
Summary: Real Potential, Unclear Prospects	128
CHAPTER 10	
A Learning and Adapting Society	131
Defining the Characteristic of a Learning and Adapting Society	132
AI as Fuel for Creativity and Learning	133
AI, Experimentation, and Adapting to a New Era	134
The Risks: Our Epistemic Crisis and AI	135
Killing the Impulse to Learn: Cognitive Off-Loading	136
Summary: The Benefits Are Likely—But the Risks Are Profound	138

CHAPTER 11

Competitive Diversity and Pluralism	141
Defining the Characteristic of Competitive Diversity and Pluralism	142
AI and Diversity.....	143
AI and Pluralism	144
A Leading Risk: AI-Fueled Social Divisions.....	145
Summary: Here Come Some Volatile Feedback Loops	146

CHAPTER 12

A Strategy for National Advantage in the Artificial Intelligence Era	149
The Route to Competitive Advantage: Fostering Individual and Collective Agency in the AI Era	151
Essential Components of a U.S. AI Strategy.....	152
Steps to Nourish the Societal Sources of Competitive Advantage: A Campaign of National Renewal.....	154
Summary: A Daunting—But Essential—Agenda	167
Abbreviations.....	171
References.....	173
About the Author	195

Figures and Tables

Figures

1.1.	The Rise of Gross Domestic Product in England	2
3.1.	The Hockey Stick of Global Gross Domestic Product.....	39
4.1.	Carlota Perez's Life Cycle of Technological Revolutions	60
8.1.	Differential Implementation of AI Uses	113

Tables

1.1.	Outcome Indicators: Factors to Measure Competitive Success	12
3.1.	Scenarios for the AI Revolution	49

Introduction

By the late 1860s, the leaders of Japan began taking seriously trends around their region and the world and recognized that they were in deep trouble. A scientific, technological, economic, and industrial revolution was transforming several leading powers, most especially the United Kingdom. In the nations that led this conversion, groups of scientists, entrepreneurs, engineers, and self-taught tinkerers—working alone and in communities of knowledge and innovation—drove spectacular progress in such areas as metallurgy, steam power, textiles, and transportation. Those conceptual advances, in turn, provided the basis for potent new military technologies.

This accelerating process had already delivered astonishing scientific marvels: the spinning jenny, the steam engine, the power loom, the McAdam technique for road construction, railroads, the cotton gin, and dozens more. It generated unheard-of advances in economic productivity and growth. If the birth of the First Industrial Revolution is dated as 1760, this growth took some time to emerge.¹ But once it really took off, as Figure 1.1 suggests, England's gross domestic product (GDP) rose in spectacular terms. As the economic historian Joel Mokyr puts it,

The British Industrial Revolution of the late eighteenth century unleashed a phenomenon never before even remotely experienced by any society. . . . Measured economic growth in the industrializing economies in the nineteenth and twentieth centuries approached a rate of 1.5–2.0 percent a year, perhaps ten times faster than before.²

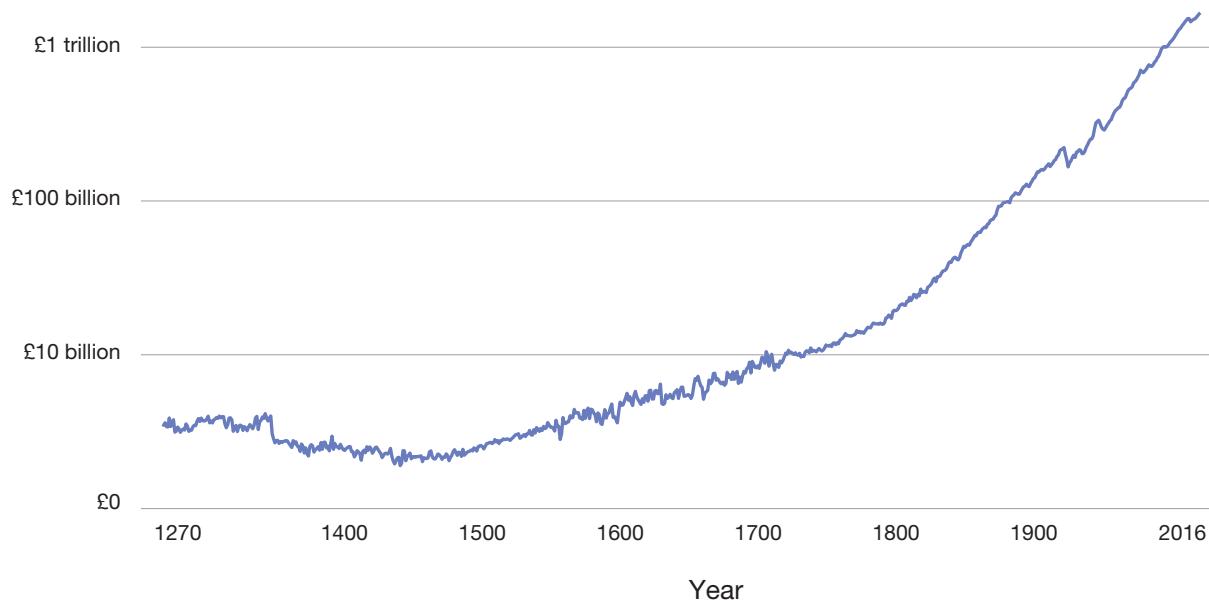
These technological advances spawned new military capabilities. Steel replaced wood in the hulls of naval ships and in the implements of land and air warfare. New mass production techniques allowed major powers to manufacture titanic numbers of weapons. Advances in chemistry and metallurgy produced more-powerful weapons of all types and the ammunition for them. Railroads revolutionized the process of getting armies into battle, and the telegraph revolutionized the command and control of military forces.

The Industrial Revolution reshaped world politics partly because countries that developed and deployed these new weapons of war would dominate those that didn't. The result was a series of profound geopolitical shifts. As the historian David Landes puts it,

As a result, the nineteenth century saw a unified Germany rise to Continental hegemony on the strength of the Ruhr and Silesia; while France, slower to industrialize, was never again to enjoy the pre-eminence to which the *levee en masse* and the genius of Napoleon had raised her on the eve of economic revolution. With the spread of the new techniques, moreover, new powers arose: the twentieth century saw the millennial predominance of Europe dwindle before the unprecedented might of the United States and Soviet Russia.³

Paul Scharre of the Center for a New American Security lays out some of the striking statistics. Between 1830 and 1890, the United Kingdom and Germany “both more than doubled their per capita GNP [gross national product] while Russia stagnated, increasing by a mere 7 percent total over a sixty-year period.” These

FIGURE 1.1
The Rise of Gross Domestic Product in England



SOURCE: Adapted from Our World in Data, "Gross Domestic Product (GDP) in England," webpage, undated-b, featuring Bank of England data (CC BY 4.0).

NOTE: These data are expressed in constant 2013 British pounds. Data refer to England until 1700 and Britain from then onward.

economic trends had geopolitical outcomes: “[W]hile Russia was the largest European economic power (measured in GNP) in 1830, by 1890 Russia had been eclipsed by Britain and Germany.”⁴ Pure scale had become less important than being at the frontier of technological and industrial progress.

In Japan, government officials watched these developments with some alarm, increasingly aware that their largely preindustrial society had little chance of defending itself against full-scope Industrial Era powers. In response, the Meiji government (1868–1912) undertook a concerted effort to learn from the world, spur the industrialization of Japan, and catch up with the leading powers of the new era. In 1871, the Meiji regime dispatched a group of scholars and officials—by some estimates, almost half of the officials serving in government at the time—on the Iwakura Mission, a worldwide fact-finding effort to understand the nature of this new technological age and bring back ideas, designs of new technologies, and policy options for Japan to catch up.

The result was a remarkable national campaign for self-improvement, and Japan hauled itself into the Industrial Era within a few short decades. By 1905, Japan shocked the world by defeating Russia, a traditional European great power (albeit a back-of-the-pack industrializer). Japan’s industrial capacity surged: Between 1913 and 1938, while Britain’s production rose by only about 18 percent, Japan’s more than quintupled. By the eve of World War II, it ranked among the world’s leading industrial powers: In 1938, it generated 3.8 percent of the world’s manufacturing output, almost half of Britain’s 9.4 percent.⁵ Both countries lagged massively behind the United States’s commanding total of almost 29 percent; it’s understandable why some Japanese leaders in 1941 thought it suicidal to challenge the world’s dominant industrial power.⁶

In contrast, during these same decades after 1870, several languishing great powers proved unable to modernize and eventually paid the price. The Austro-Hungarian and Ottoman empires could not keep up with European industrial leaders and—partly although not completely because of this—suffered defeat in World War I and eventually disintegrated. Russia possessed seemingly inexhaustible manpower reserves

and the geography to fatigue any invader but could never cultivate a fully industrialized, efficient, mixed economy to compete with such technological leaders as the United States and Germany.

This story of 20th century industrialization seems like ancient history. But, in another way, it remains powerfully relevant. The world stands on the cusp of another technological revolution—the emergence of artificial intelligence (AI) as the latest and potentially the most wide-ranging and influential general-purpose technology in history. Like its industrial precursor, the resulting change will decide national fates and reshuffle the deck of global power. Some countries will become the Britains and Japans of this age—the dominant early leaders or the clever, determined fast followers. Others might end up as the Ottoman Empire of the AI Era, dragged down by social and political patterns disastrously mismatched to the demands of technological and industrial leadership.⁷ Still others will become the Russias of this period—large and powerful by the indiges of the previous era and too potent and too determined to be ignored but not remotely competitive with the leading economies across the variety of frontier technologies and social innovations.

The historian Roger Osborne summarized the significance of that earlier revolution in a simple phrase: “The Industrial Revolution is the nexus through which all of modern human history flows.”⁸ The AI Revolution—emerging alongside but also serving as a critical support system for parallel advances in biotechnology, robotics, new energy sources, nanotechnology, advanced manufacturing techniques, and other areas of technology—is set to become the new nexus of history. From the standpoint of technological, economic, and military power, a leading medium-term challenge for the United States and every other nation is to develop a strategy for competitive success in the AI Era.⁹

I endorse these broad claims as a skeptic of the more-extreme assertions of AI’s countless powers. Even as we recognize the potential of the AI Revolution, it will be important not to exaggerate its transformative potential with breathless language and promises of pure magic. Significant changes are coming—but just how significant and on what time scale, we don’t know. Many issues with major implications for national strategy brim with uncertainty, such as the speed of AI model development, the diffusion of AI applications throughout societies, and the relative importance of open and closed models. This isn’t even the first time that excited observers forecast an imminent AI Revolution.¹⁰ Beginning in the 1970s, experts repeatedly claimed that AI was on the verge of massive advances, only to have the expectations fizzle in a series of so-called AI winters when progress stalled and funding for AI research dried up.¹¹ But the scope of progress now underway is different, and the risks of being left behind are too great to ignore.

In developing a comprehensive national strategy for the AI Revolution, the U.S. government has to appreciate the relative risk of doing too much versus too little. If the United States overprepares for a transition that turns out to be slower and less transformative than some expect, the price will be modest. That’s partly because, as I argue, the many steps needed to prepare for the AI Era are good for American society, no matter how fast it arrives. But if the United States languishes at this hinge point in history, U.S. competitive position might never recover. In this example of decisionmaking under uncertainty, almost all the risk is on one side: Underplaying the meaning of one of the great technological, economic, and social transformations in human history could be a catastrophic error.

But it’s an error the United States is in the process of making. As much as U.S. firms are leading the AI technology race and U.S. public institutions have begun to rise to the challenge, the United States has barely begun to develop strategies and take actions for the far more important job of readying American society for this shift.

The Argument

Many leading voices are already shouting from the rooftops about AI’s competitive potential, and the U.S. government has taken important steps to ensure access to critical components of the AI technology stack. But

as important as semiconductor chips, training runs, and data centers are, the United States has been missing a larger truth about AI and national competitiveness. Many people in the United States—in AI labs, in government, in the analytical and research communities—believe that the route to national success lies through dominating the technology stack on which AI progress depends. But beyond that technology-centric vision, we need to begin thinking much more seriously about AI as a social phenomenon—and wrecking ball. Not all of the countries moving through this transformation will weather the resulting transition in ways that leave them stronger, more coherent, and more dynamic. The countries that will lead this new AI Era will figure out how to use AI to make their whole societies more competitive. Simply put, my core thesis is that the competitive challenge of AI is primarily social, not technological.

In the process of making that argument, I stress a few supporting themes. The first reaffirms a well appreciated lesson from history: Technological revolutions have powerful geopolitical shock waves. Some countries prosper, and some fall behind, but broad technological transformations are among the most essential challenges to a grand strategy that countries can face.¹²

Paul Scharre, one of the most thoughtful commentators on AI and national power, agrees that “[l]ike prior industrial revolutions, the cognitive revolution will reshape geopolitics in the twenty-first century. AI is likely to lead to shifts in power on the global stage, empowering some actors and even changing the key metrics of power.” Dozens of countries are now acting on this thought: Scharre notes that “[m]ore than fifty countries have signaled their intent to capitalize on AI for national advantage.”¹³ A 2025 RAND survey found that at least 75 countries had published national AI strategies.¹⁴ One example is France, which has supported the development of its increasingly prominent and widely used Mistral open-source AI model and which has begun hosting global events to establish its position as a thought leader on AI, its opportunities, and its risks.¹⁵ The United Arab Emirates (UAE) and Saudi Arabia are making huge investments in AI, partly to spur their transition to economic relevance in a post-fossil fuel world.¹⁶ Canada has recently announced a Sovereign AI Compute Strategy, which includes making major investments in data centers.¹⁷ Singapore has made a major push by developing smaller models trained on local languages with culturally sensitive responses.¹⁸

Some countries will surge ahead in the AI Era. Others will be left behind. European countries could be AI laggards if 2024 statistics are anything to go by. European firms and other organizations are behind comparable U.S. institutions in AI adoption by 45 to 70 percent.¹⁹ Electricity prices are much higher in European countries than in the United States (or China), hindering pursuit of cutting-edge models. Europe has only 18 percent of the world’s data center-based computing power, whereas the United States has twice as much.²⁰

If Europe is in a tough spot, many developing nations risk missing out on this revolution entirely or even being victimized by it. This could be true in many ways: if, for example, AI helps invent new manufacturing techniques that allow many developed countries to produce a higher proportion of goods at home, if AI chat-bots replace almost all call center employees (as is already happening), and if students in developing nations can’t access the benefit of AI-enhanced schools and doctors can’t use AI-supported treatments.

As of this writing in 2025, only one country is poised to keep pace with the United States as a comprehensive AI power—the United States’ geopolitical rival, China. For the moment, it is taking a different approach than the United States, focusing more on the application and use of AI models than on generating the most-powerful leading models. But Beijing has identified AI as a national priority and is funneling vast resources into its development. The United States and its friends and allies still have a decisive lead in cutting-edge semiconductor chip design and production. But China also has advantages: an immense and highly talented labor pool, huge reservoirs of data to train AI models, and cheap electricity, among others. China’s role as the primary U.S. AI competitor poses major geopolitical risks. “If the United States and other democracies do not work together to lead in AI and shape the rules for how it is used,” Scharre worries, “they risk a creeping tide of techno-authoritarianism that undermines democracy and freedom around the globe.”²¹

The Gateway to a New AI Era

The second broad theme underpinning this work stems from the complex, multifaceted nature of any technoinustrial revolution. Such transitions knit together many leading technologies, techniques, and national capabilities.²² Success depends on organizing the whole society to support these ecosystems and ensuring that the technologies strengthen rather than weaken the social foundations of dynamism and power. The implication is clear: Competitive advantage in this AI Era will depend not solely on mastering the narrow technologies of AI, such as the data centers, the algorithms, the models, or the semiconductors. Nations will flourish to the degree that their societies can both embrace new technologies and maintain social cohesion during the transition.

The vast majority of attention focused on the AI Revolution treats it as primarily a technological challenge. It's about stacking up computing power, scaling the models, powering AI with gobs of energy, and applying it to generate scientific and technological innovation, which then drives productivity and growth.²³ But all of that still needs to be integrated with and transmitted through the broader society, and the nations that will benefit the most will be built to compete as societies, not as technology stacks. That idea implies a far broader strategy and policy challenge than simply staying ahead in computing power and frontier models.²⁴ U.S. AI policy and strategy has, so far, been too narrowly focused and risks neglecting the essence of the true sources of a long-term advantage.²⁵ The 2025 U.S. AI Action Plan took important steps to emphasize this wider agenda, highlighting such issues as making AI help American workers.²⁶ But, in terms of practical initiatives to shape AI's effects for competitive advantage, much remains to be done.

In retrospect, we think of the economic narrative of the Industrial Revolution as the famous hockey stick graph—productivity and growth rocketing upwards, taking living standards and the deployment of technological marvels with them. Yet this narrative overlooks a grimmer reality, particularly in the early decades of the era. Industrial workers endured appalling conditions, environmental degradation scarred the landscape, and political upheaval convulsed society. In the 1840s, just as Britain was cementing its industrial supremacy, the nation suffered the greatest social turbulence in its modern history, including famines, labor disputes, and the Chartist movement seeking to mobilize the masses behind radical demands for political reform.

The AI Revolution could pose social, economic, and political risks that are every bit as treacherous as those of the Industrial Revolution—and very likely far greater. Some analysts have forecast massive job displacement and associated social unrest. Already, algorithmic decisionmaking is sweeping into many corners of institutional life, such as criminal sentences, college admissions, and loan decisions, in ways that threaten humans' understanding of and control over the systems that are making critical choices about their lives. Millions of Americans are using chatbots as sources of advice and companionship. Autonomous AI agents have just begun to operate across the economy.

Another of those risks could be dependence on AI. Much as many organizations and people can't get through their day without depending on a whole suite of technological supports—their word processing programs, their smart watches, their Global Positioning System (GPS) maps, the complex software models that run major industrial operations, and so much more—we're all going to become increasingly addicted to the functions that AI performs. We'll rely on AI to manage our schedules, run our smart homes, interview our job candidates, fill out our paperwork, advise on such complex procedures as surgery and engineering, and a thousand other tasks. Adding yet another complex, mysterious layer of technoprocessing between daily tasks and human understanding risks creating a civilization largely incapable of fending for itself. Emerging research suggests that heavy AI users become more cognitively passive, off-loading thinking processes to AI models. Such dependence is dangerous not only to humanity but also to nations, which become helpless against attacks on the systems that make daily life possible.

Adding to a Wave of Change: AI and Major Trends

As if that prospect weren't daunting enough, the social challenges of AI are arriving at a perilous time in the progress of human society, when we are arguably reaching the natural end point of many institutional forms and epistemological habits. My third major theme is that the AI Revolution will crash into an already unstable social, economic, and political context; from the standpoint of national dynamism and competitiveness, the single most essential question is whether societies use AI to mitigate these rising dangers or allow it to exacerbate them. This technological tidal wave will destabilize societies—especially those in technologically advanced nations—that are already undergoing wrenching changes. The resulting volatility will play a big role in determining the effects of the AI Revolution.

Several leading factors define the social, political, and economic canvas on which the AI Revolution will begin to do its work. To begin with a political trend, many societies have become bitterly polarized along partisan and ideological terms. Although Americans continue to indicate respect for certain shared values, partisan divergences on many public policy issues are immense and, in some cases, unprecedented in modern polling. Repeated polls suggest that each side of the partisan divide views the other not merely as misguided but as dangerous to the country. The ideological polarization of members of Congress stands at a post-World War II apex. AI's effects on the social and political landscape risk exacerbate this polarization, and shaping AI's effects will be complicated by these severe partisan divisions.

At the same time, the effectiveness of many Industrial Era governing forms, including large hyper-bureaucratized institutions, has reached a crisis point. Evidence of the cost of massive bureaucracies—in efficiency, effectiveness, motivation, and human dignity—is undeniable.²⁷ Many studies have documented the vise lock being placed on national dynamism of overbearing rules and regulations in such fields as defense procurement and higher education, the mountains of required paperwork in such sectors as medicine, the explosion of administrative workers in almost every large institution, and many other hallmarks of a period that's seen the intensifying ossification of the institutions that operate society. AI tools have the potential to transform people's engagement with large organizations, but whether that influence empowers or alienates individuals remains to be seen.

In demographic terms, the United States—similarly to all developed and even many developing countries in the world—confronts a demographic decline of historic proportions, raising the prospect of aging and eventually shrinking populations. The numbers are truly astounding: As leading demographer Nicholas Eberstadt points out, using East Asia as an example:

By 2022, every major population there—in China, Japan, South Korea, and Taiwan—was shrinking. By 2023, fertility levels were 40 percent below replacement in Japan, over 50 percent below replacement in China, almost 60 percent below replacement in Taiwan, and an astonishing 65 percent below replacement in South Korea.²⁸

Eberstadt concludes that in “the United States and elsewhere, thinkers and policymakers are not ready for this new demographic order. Most people cannot comprehend the coming changes or imagine how prolonged depopulation will recast societies, economies, and power politics.”²⁹ This trend will exacerbate fiscal crises for many countries and threaten long-term growth prospects by shrinking working-age populations.

These developments and others have contributed to an especially dangerous trend: the loss of public faith in institutions of all kinds, especially public-sector institutions. Survey after survey has found cratering faith in public and private institutions. Fewer than 30 percent of Americans (in many cases, fewer than 20 percent) express either “a great deal” or “quite a lot” of confidence in institutions, including Congress, the presidency, the Supreme Court, banks, newspapers, organized labor, and television news.³⁰ As I argue in Chapter 8, these attitudes reflect a more objective, structural legitimacy crisis affecting many developed nations to greater or

lesser degrees. The application of AI tools can potentially make formal institutions more effective and more efficient but also more arbitrary and more alienating, and this discrediting effect could further weaken institutions that are already losing legitimacy.

Part of the reason for this crisis of legitimacy, but also a profoundly important trend of its own, is the epistemic crisis corroding the public spheres of many countries. The information environment is becoming increasingly fragmented, corrupted, and polluted, creating a situation in which shared conceptions of reality are fraying.³¹ I say more about this trend in Chapter 10. It strikes at the very heart of any functioning democracy by threatening to corrupt the functioning of the marketplace of ideas, and AI could contribute to this dangerous process by introducing tidal waves of highly persuasive misinformation.

Moreover, in the United States and many other countries, social and economic inequality has been rising.³² There is some debate about the precise extent of this trend, its pace, and its degree over time. Several studies suggest that income and wealth inequality has declined slightly. But U.S. society embodies very significant inequality, which has certainly grown since the heyday of more shared wealth in the post-World War II years.³³ The bottom 50 percent of societies holds a tiny percentage of total wealth: 2.4 percent in the United States, 9 percent in Britain, and 6.5 percent in Spain.³⁴ Americans certainly perceive an unequal country, and three-quarters of them highlight the issue as one of their top political priorities.³⁵ Partly as a result of many of these preceding trends, social alienation appears to have been rising for some years, reflected in reported higher rates of anxiety and distress and a variety of mental illnesses, including depression.³⁶

Another aspect of the existing context is that AI isn't the only breakthrough technology accelerating out of the gate. AI will arrive alongside several others—most especially biotechnology and nanotechnology but also quantum and breakthroughs in renewable energy and advanced manufacturing—that will interact in profound and unpredictable ways. AI pioneer Mustafa Suleyman notes, “Technology isn't independent breakthroughs” but “a commingling set of parts to combine and recombine.” He argues that “[t]he coming wave is a supercluster, an evolutionary burst like the Cambrian explosion.”³⁷ Although AI is the subject of this publication and looks to be certainly the fulcrum of a new techno-industrial era, that era might eventually be known as a broader transition involving contributions from biotechnology, adaptive manufacturing, and other fields.

I discuss the consequences of some of these trends on specific AI outcomes in later chapters. For now, noting that advanced democracies, in particular the United States, are suffering through periods of social, economic, and political disruption is enough. The effects of digital technologies are already problematic. AI is arriving at an incredibly challenging time, a fact that will complicate its effects on the competitive advantage of nations.

AI and the Future of Human Agency

My final core theme is that AI is changing the context for willful and meaningful human agency. Even short of the worst loss-of-control scenarios in which AI agents escape the power of human beings to manage them and even before humanity reaches superintelligence (as commonly defined), AI could begin either supercharging or emasculating autonomous human agency in hundreds of smaller ways that add up to a loss of societal energy and coherence. This issue is inextricably linked to national competitiveness.

Already, many of the trends I mentioned are stacking up to threaten autonomous agency, both for the individual and for society as a collective. Globalization and the demands of financial markets have reduced national free choice. Massive bureaucracies quash individual human agency, of both their employees and their customers. Ineffectual governance structures prevent groups of people from expressing agency through collective action. An information environment inundated with misinformation, active disinformation, and

conspiracy theories undermines agency by destabilizing the epistemological foundation for human choice. Stark inequalities leave many people disempowered in economic, political, and social terms.

Such trends are shaking the societal foundations for competitiveness, at least in the United States. When a person's sense of agency and control over life weakens, national willpower ebbs. National identity fractures as people turn to competing ideologies and myths and localized identities in a frantic bid for intellectual agency. Institutions become ossified through the death of accountability—a side effect of collapsing agency.

And now AI is arriving with the threat—and, for some, the promise—of a further assault on human agency.³⁸ How these dynamics play out will go a long way to determining AI's effect on national standing.

The idea of autonomous agency is one of the defining features of modernity—a “freeing of the subject,” as the philosopher Charles Taylor put it, with “a sense of self-responsible autonomy, a freedom from the demands of authority.”³⁹ Agency endows human beings with the ability and drive to know, investigate, judge, and realize themselves as sovereign agents worthy of freedom and dignity. The idea revolutionized the nature and practice of individual morality, moving the central moral question from strict obedience to externally set norms to accountable moral judgment by free individuals.⁴⁰ It laid the basis for democratic politics by upholding the free choice of autonomous agents as the prime political value. It went hand-in-glove with the scientific revolution and method, which valorized experimentation evaluated against objective criteria conducted by self-guided researchers.

Autonomous agency has been a critical support system for national competitive advantage in the modern era. Societies that freed up their populations to investigate and discover in these ways unleashed tremendous dynamism. States that gathered the autonomous solidarity of free peoples to fight overtook obedience-based societies in war. Societies reflecting the concept of autonomous agency became engines of sustained ambition, commitment-based solidarity, broad entrepreneurial and scientific energy, grassroots experimentation and adaptation, and many other sources of advantage.⁴¹

This concept of autonomous agency is, therefore, central to our modern understanding of ourselves as human beings and to the nature of our societies and political communities. What AI does to this quality—or, rather, how nations manage the transition to the AI Era in ways that nourish or destroy such agency—will have dramatic effects on national prospects. AI has the potential to turbocharge human choice rather than destroy or substitute for it. That's the argument of authors Reid Hoffman and Greg Beato in the book *Superagency*. In hundreds of ways, such as learning languages, controlling the energy efficiency of homes, and planning trips according to accurate weather forecasts, “AI is increasing your agency,” the authors suggest, “because it's helping you take actions designed to lead to outcomes you desire.”⁴² Hoffman and Beato give the example of the democratization of GPS guidance, which made possible GPS-directed apps that opened new avenues of agency in various areas, such as ride-sharing services. The authors argue,

The more adept you become at using LLMs [large language models] to navigate life in the twenty-first century the greater your power to plot your own path through the world. And what's more democratic and dynamic in the long run: societies that strive to increase choice and autonomy for everyone, or ones that protect existing conditions for entrenched incumbents forever?⁴³

The result will be to strengthen the link between agency and national advantage: “All of this means that in the twenty-first century, individual agency is more closely aligned with national agency than ever before.”⁴⁴ Superagency isn't automatic and will require sensible decisions and a general commitment to meeting the social and economic challenges of AI. And those decisions will have to fight against tendencies in the private-sector and social context, which could easily push AI toward profoundly disempowering futures.

Scope and Approach

Beyond the issue of nation-versus-nation competitive standing, AI creates a challenge unique to an era when the emerging technology will have, in quite literal terms, a mind of its own. The United States must focus as well on what is commonly termed *AI safety*—avoiding runaway or misaligned AI and trying to prevent its misuse by destructive actors as much as possible.⁴⁵

In this analysis, however, I deal with the problem of the control or alignment of AI on only the margins, largely because there is a vast literature on AI safety and dozens of efforts underway to address that issue. My focus is on how nations can be competitive in the AI Era. What does a society look like that will prosper under the tremendous demands of AI? What strategic choices can a nation make to set itself up for success? Yet the issue of AI safety remains absolutely critical: Out-of-control AI will subvert just about any goals a nation has for itself, including its competitive standing.

A second limitation to my scope has to do with my treatment of AI as a phenomenon and a technology. I have had the immense good fortune to be conducting my research as part of the larger research agenda of the Center for the Geopolitics of Artificial General Intelligence within RAND’s larger ecosystem of AI research. I’ve benefited from ongoing interactions with computer scientists, AI researchers, technical experts who have experience working in the major AI labs, economists, and many others pursuing parallel lines of research. We’ve engaged with leaders in the field of AI research and development (R&D), scholars of world politics, and former senior government officials trying to come to grips with AI’s geopolitical implications.

But despite whatever I have picked up about the science of AI, this publication is in no way a technical analysis. I’ve had to make some judgments about some technical issues despite the limits of my knowledge. But I am no technologist and have relied on other experts in these matters to inform any necessary technical judgments.

What I can bring to this discussion is a broader perspective that is based on years of RAND work on the foundations of national success. I reviewed many classic ideas about what makes nations successful in competitive terms—being dynamic, energetic, stable, willful, and effective. And then I placed the likely effects of AI into that context.

To inform that wider view, I conducted three overlapping courses of research and analysis to produce the insights that follow. First, I reviewed the literature on the sources and consequences of the Industrial Revolution and on the history of technological revolutions more generally. For this research, I built on prior RAND work on the competitive qualities of nations, which I describe in the following section. Second, I reviewed the now burgeoning literature on AI’s possible economic, social, political and military effects, partly to compare it with earlier periods and judge whether this might, in fact, be an unprecedented transformational era. In that thinking, as I noted previously, I have benefited from dozens of engagements with AI thought leaders over the past two years.

Third, I used the most advanced public versions of three leading generative AI models—Claude, ChatGPT, and Gemini—as consultants, asking for their thoughts on various issues (and occasionally minor editing advice). The answers were sometimes bland and generic, sometimes useful, and sometimes unnervingly insightful. The project became, in that sense, another advertisement for AI as a valuable *copilot* (in AI researcher Ethan Mollick’s terminology) for anyone doing analysis.⁴⁶ But my use of the models also repeatedly showed their limitations, at least as of this writing.

The Organizational Scheme: A Framework for Understanding National Competitive Advantage

My argument about AI's societal importance is grounded in a simple idea: National competitive success stems first and foremost from the characteristics of societies. Nations can enjoy potent geographic advantages, clever strategies, world-class scientists, and good leadership. But if they lack the essential qualities that generate dynamism, creativity, and solidarity, they will not be able to sustain their power over the long run. That basic lesson was the primary finding of a three-year study for the Office of Net Assessment in the Pentagon, which I led at RAND and which is described in the 2022 report *The Societal Foundations of National Competitiveness*. That study sought to identify the qualities of societies that were essential to success in long-term rivalries.⁴⁷ It nominated seven major societal characteristics and a list of other important factors that the authors thought played an outsize role in shaping national fates.⁴⁸ The box on the following page summarizes the seven elements. They provide the framework for this assessment—looking at the ways in which the emergence of AI might affect each of them and, through those factors, national competitiveness in general.

The most powerful modern example of this connection comes from the Cold War. The Soviet Union poured colossal resources, organized and directed by an autocratic central government, into industrialization. It built an immense military and a terrifying nuclear arsenal. Its leading scientists generated breakthroughs in many fields. Yet it eventually lost the competition with the United States and the West because its economy and society proved too stagnant. Across all the characteristics surveyed in this analysis, the United States totally outclassed the Soviets.

Similarly, today, national competitive advantage in the AI Era will surely be influenced by large investments in the technology stack but will ultimately be determined by broader factors. Nations that have more dynamic and effective societies will develop, integrate, and manage the effects of the AI Revolution to better competitive advantage.

The remainder of this book unfolds largely along the lines of the framework described in the box on the following page—anticipating the ways that AI will shape each of the seven characteristics for good or ill. In Chapters 2 and 3, I review some basic facts about AI, briefly survey some of its emerging capabilities, and explain why this historical transition, although lightning fast compared with earlier cases, might be slower than some AI enthusiasts claim. In Chapter 4, I lay out some of the historical lessons that provide the basis for my conclusions. Chapters 5 through 11 represent the core of my analysis: Each chapter focuses on one characteristic from the framework. In those chapters, I define each characteristic, describe the ways in which AI could support the characteristic, and review risks that AI might pose to it. Finally, in Chapter 12, I draw on this entire analysis to sketch out a national strategy for competitive advantage in the AI Era.

In *The Societal Foundations of National Competitiveness* and in this effort to tie that study to the potential implications of AI, we can't very well speak to what makes countries successful without defining what success looks like. What does it mean to succeed in a large-scale competition or rivalry? In *The Societal Foundations of National Competitiveness*, I reviewed historical research, looked at criteria that were used in studies of the rise and fall of great powers, and developed a set of nine indicators of success in these kinds of competitions. Table 1.1 lists these factors.

When I discuss the effect of AI on national competitive success, these are the outcome indicators that I have in mind. A country that is strongly empowered by AI and achieving success relative to others would be one that sustains its long-term national identity, bolsters its self-defense capacity, protects itself against domination or control by others, wins military advantages in service of the country's defense strategy, attracts other countries to its example and power and sets global norms, enjoys a leading position in the global economic hierarchy and global networks of trade and finance, and is the leader across many frontier technologies.

Sources of National Competitive Advantage: National Characteristics

Characteristic 1: National Will and Ambition

Nations that reach the pinnacle of world politics and technological achievement almost universally rely on an abundant supply of national ambition and will—an urge for intellectual achievement and superiority and a sense of national destiny and greatness that fuels their competitive drive. This societal sense of ambition and drive ultimately translates down to the individual level and the norms and habits often associated with what is loosely described as a society's work ethic.

Characteristic 2: Unified and Coherent National Identity

Nations with a stronger sense of identity and societal coherence have a competitive advantage over nations with a more fragmented sense of identity or that lack allegiance to any sort of unified nationhood. This quality is often expressed in a robust and strongly felt citizenship. The precise degree of required unity has varied from case to case, but countries have clearly benefited from high levels of coherent identity and suffered from its absence.

Characteristic 3: Shared Opportunity

The role of shared opportunity represents the degree to which all the people of a nation can work, advance in career and achievements, express and develop ideas, create, network, and in other ways contribute their full human potential to the life, prosperity, and power of the nation. This characteristic would require that people's opportunities to express their abilities—in their careers, in their ideas and creativity, or in their general ambitions—are not limited by their membership in any group, including family or tribe or party, or by gender, class, race, or ethnicity. The most-successful societies also extend this principle beyond their national boundaries.

Characteristic 4: An Active State

In every sustainably competitive society examined, the state has played an important role in shaping the socioeconomic, military, and geopolitical contexts for success. The concept of an active state is not equivalent to state-controlled economies or societies. The default recipe of a highly competitive society depends centrally on various forms of grassroots, uncontrolled, creative, and often disruptive energy, which is too organic and chaotic to be managed by any state apparatus. There is a need for an active state to be energetic, forward-looking, and powerful enough to create the conditions for success but not so overbearing or constrained by orthodoxy that the state strangles national energy.

Characteristic 5: Effective Institutions

The central importance of effective public and private institutions has emerged as a consistent theme across the literature. This is true for many reasons: Effective institutions reduce transaction costs and generally smooth commercial and social interactions, fill gaps left by the state, help solve social challenges, and provide fuel for competitive advantage in various areas, such as scientific research. Such institutions are the essential fabric through which the other characteristics do their work by providing many structural protections for shared opportunity and allowing the active state to operate efficiently.

Characteristic 6: A Learning and Adapting Society

Highly dynamic and competitive nations are typically thirsty for new ideas and eager for fresh policies and approaches. Such nations cultivate networks of scientific and intellectual discussion and debate, and they both tolerate and actively encourage the widespread public sharing of new knowledge. Dynamic nations apply learning in practical ways and continually reassess their ways of doing business.

Characteristic 7: Diversity and Pluralism

This characteristic includes two closely related but distinct elements. One is diversity, defined as the level of variation in a society—not only in terms of gender, race, ethnicity, sexual orientation, and so on but also in every other way that diversity can be defined. This can include citizens who have a wide variety of education, training, career paths, and skill sets; people from geographically distinct parts of the nation that have different cultural traditions and even languages; and people who have divergent experiences. The second element of this characteristic is pluralism, which encompasses two aspects of a society: the degree to which it has overlapping sources of authority, rulemaking, and governance and the degree to which its people value and tolerate multiplicity.

SOURCE: Adapted from RAND Corporation, 2022.

Any analysis of the future must make a choice about a critical analytical issue—its time frame. Here, I am concerned largely with the next decade, roughly out to the year 2035. Even in that span, in the blink of a historical eye, immense uncertainties abound. But my chosen time frame is a medium-term look, trying to get beyond the low hills of the next year or two but not aspiring to anticipate what the world might look like across the distant mountains of decades.

TABLE 1.1
Outcome Indicators: Factors to Measure Competitive Success

Indicator of Competitive Success and Advantage	Historical and Current Examples and Metrics
Longevity in terms of long-term socioeconomic and geopolitical resilience that maintains national identity over an extended period and promotes extended cultural and social influence	<ul style="list-style-type: none"> • Trends in measures of national power • Collapse or surrender by rival • Long-term, indirect, and diffuse social, cultural, or political influence
Sovereign ability to protect the safety and prosperity of citizens against capabilities or threats of other states, nonstate actors, and systemic risks	<ul style="list-style-type: none"> • Power to prevent large-scale territorial aggression against homeland • Ability to prevent harassment or disruption of society short of war
Geopolitical freedom of action in terms of the ability to make free and unconstrained sovereign decisions and take actions in the international system to the greatest degree that relative power will allow	<ul style="list-style-type: none"> • Absence of coercive control by regional or global hegemon • Self-sufficiency in materials and factors necessary for freedom of action
Military advantage or dominance, locally or globally, and the ability to project power	<ul style="list-style-type: none"> • Global military dominance, either generally (e.g., Rome, post–Cold War United States) or in specific domains (e.g., British maritime dominance) • Ability to project power from a distance
Leadership of or membership in predominant alignments of military and geopolitical power	<ul style="list-style-type: none"> • Modern treaty-based alliances, multilateral or bilateral • Less formal security relationships
Predominant economic strength globally, in a region, or in one or more industries	<ul style="list-style-type: none"> • Total or per capita GDP • Share of global trade, investment, or research in critical industries
Strong to predominant position in global trade, investment, and capital markets (relative to size of GDP and other factors)	<ul style="list-style-type: none"> • Role in regional or global trade networks (e.g., Egypt, Rome, the United Kingdom, the United States) • Dominance of national currency • Predominant power in economic institutions
Strong to predominant position in ideological and paradigmatic categories and global narratives and norms, attractive power, and international institutions and standards	<ul style="list-style-type: none"> • Cultural influence • Alignment with leading global norms and values • Leadership of international organizations and norm-setting processes
Strong or leading position in frontier technology; leading or dominant role in key emerging industrial sectors	<ul style="list-style-type: none"> • Domestic capabilities and industries in leading industries of the era • Measures of relative technological standing • Proportion of R&D spending in key industries

The more drastic forecasts of AI progress suggest that a decade-long view makes little sense. As Dario Amodei, the chief executive officer (CEO) of Anthropic, has put it, “the critical period here where there’s really going to be contention, or where it’s important to achieve a balance of power, is going to happen in 2026, 2027, or at the latest, 2030. Policy should target that time range.” In the world of fast-moving AI, “[Ten] to 15 years is like an eternity. It’s forever. It’s almost irrelevant.”⁴⁹

That’s certainly true in some narrow aspects of seizing competitive advantage, which seems to be what Amodei is talking about. Export controls on semiconductors over the next few years will play a major role in determining who’s ahead in the technology. But in terms of long-term sustainable national or societal advantage, I’d say the next decade, or even 20 years, is the critical window—the time during which nations have the opportunity to build on their short-term technological position to reach deeper forms of societal advantage.

Some might object to a focus on national rivalries and competitive advantage in an assessment of AI’s potential. I don’t intend for this analysis to reinforce emerging zero-sum notions of the U.S. competition with China or any other nation. The strategic concept I lay out is not uniquely designed to compete; its ideas overlap very significantly with any agenda to ensure that AI promotes the vigor, unity, and coherence of U.S.

society. But the United States is engaged in some form of rivalry with China. Science and technology, including AI, are part of that competition, and U.S. officials have rightly been concerned with the problem of how to establish and maintain a sustainable competitive position. Such a position need not imply malign intentions toward China or any other competitors in this field. But the United States can't ignore the reality of this rivalry or the need to ensure U.S. competitive dynamism.

Coming to Grips with a Historical Transformation

Edward A. Feigenbaum and Pamela McCorduck, two authors well-versed in the power and possibilities of AI, discussed the implications of an AI Revolution in the context of earlier epochal transitions.⁵⁰ They wrote that if a prophet were asked at the beginning of the agricultural revolution in about 4,000 BC what its effects would be, that prophet might have anticipated some short-term developments, but “the rise of the cities, international trade, or the peanut as an ingredient in shampoo” would not have been possible to anticipate.

In other words, if we humans are luckily endowed with the imagination to create revolutions, we nevertheless can hardly anticipate their long-term effects. . . . We stand, therefore, before a singularity, an event so unprecedented that predictions are almost silly, since predictions, by their very nature, are extrapolations from things as we know them, and the singularity called reasoning machines will change things from how we know them in unpredictable ways.⁵¹

That forecast was made *40 years ago* in 1983. For decades afterward, there was no singularity, no AI Revolution—and as a potentially dramatic AI Revolution finally arrives, Japan is certainly not leading it, as the authors proclaimed that it would. Such misplaced efforts to forecast the future of AI ought to urge a significant dose of humility as we seek to anticipate what the coming years hold—and even more modesty as we design strategic responses that, if we plunge headlong down some paths rather than others, risk locking us into second-best approaches and kneecapping our national position. This is not the time for certainties.

Yet it's also time for radical moves.⁵² The need for a national strategy to adapt to the new era is urgent, especially because this transition is likely to move at warp speed compared with even the arrival of the Industrial Era. AI is arriving at a dangerously unsteady time for many societies. People are already feeling disempowered, alienated, poorly served by public institutions, and confused about what is true on just about any issue. If national competitive standing is built on the key social foundations I survey in later chapters, a rapidly emerging AI Era without any conscious direction carries intense risks.

It does remain within our power to shape this outcome. We confront the most important governance challenge in human history—a need to collectively decide on the effects we *want* AI to have rather than merely sitting back and accepting whatever effects an uncontrolled emergence of AI forces on us. And the clock is ticking. As Goldman Sachs analysts Jared Cohen and George Lee have argued, “The world is facing a narrow window of opportunity . . . to shape the AI-enabled future. This window will be brief—a few years at most—then views and strategies will harden; norms, values, standards will be embedded within the technology; and the costs of changing course will rise.”⁵³ That's as true for national strategies to prepare for the AI Revolution as it is about the technology itself. The United States—and all other countries determined to thrive in this era—face a pressing task of identifying the first pieces of an eventual strategy for competitive success and beginning to implement them before it's too late.

Notes

¹ On the gradualism of technological revolutions, see Carlota Perez, *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*, Edward Elgar, 2002.

² Joel Mokyr, *A Culture of Growth: The Origins of the Modern Economy*, Princeton University Press, 2018, pp. 3–4.

³ David S. Landes, *The Unbound Prometheus: Technological Change and Industrial Development in Western Europe from 1750 to the Present*, Cambridge University Press, 2003, p. 11.

⁴ Paul Scharre, *Four Battlegrounds: Power in the Age of Artificial Intelligence*, W. W. Norton, 2023, p. 12.

⁵ Paul Kennedy, *The Rise and Fall of the Great Powers: Economic Change and Military Conflict from 1500 to 2000*, Random House, 1987, pp. 299, 330.

⁶ Japan's capacity for catch-up would be on display half a century later as well. After the devastation of World War II, Japanese government and business leaders rebuilt the country into an even more leading-edge economic power. By 1993, Japan very nearly edged past the United States, producing 22.1 percent of world output compared with America's 22.9 percent (Mark J. Perry, "Animated Chart of the Day: World's Top Ten Manufacturing Nations, 1970 to 2017," American Enterprise Institute, July 5, 2019). Worried observers predicted that the future of the Industrial Era belonged to Japan. The most famous of these was Ezra Vogel, *Japan as Number One: Lessons for the United States*, Harvard University Press, 1979.

⁷ In this paper, I use the term *AI Revolution* to describe the emergence and application of AI capabilities and the transitional process it implies and the term *AI Era* to refer to the age of history that comes into place on the far side of that transition.

⁸ Roger Osborne, *Iron, Steam, & Money: The Making of the Industrial Revolution*, Bodley Head, 2014, p. 5.

⁹ Mustafa Suleyman, *The Coming Wave: Technology, Power, and the 21st Century's Greatest Dilemma*, Crown, 2023, p. 160.

¹⁰ For a short history of several key phases in AI development, see Julian Togelius, *Artificial General Intelligence*, MIT Press, 2024, pp. 18–20.

¹¹ For example, two authors who had tremendous foresight but also a dose of overexuberance warned in 1983 that AI was about to transform economies and that Japan was set to dominate the era. See Edward A. Feigenbaum and Pamela McCorduck, *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, Addison-Wesley, 1983, pp. 1–5.

¹² The economist Larry Summers has said it simply: "The more I study history the more I am struck that the major inflection points in history have to do with technology" (Joe Walker, "Larry Summers—AGI and the Next Industrial Revolution," *Joe Walker Podcast*, October 22, 2024).

¹³ Scharre, 2023, p. 4.

¹⁴ Lennart Heim, *Understanding the Artificial Intelligence Diffusion Framework: Can Export Controls Create a U.S.-Led Global Artificial Intelligence Ecosystem?* RAND Corporation, PE-A3776-1, January 2025.

¹⁵ Ina Fried, "France Stakes Its Place as an AI Hub," *Axios*, November 22, 2024b.

¹⁶ Sam Winter-Levy, "The Emerging Age of AI Diplomacy: To Compete with China, the United States Must Walk a Tightrope in the Gulf," *Foreign Affairs*, October 28, 2024.

¹⁷ Charlotte Trueman, "Canadian Government Launches Sovereign AI Compute Strategy," *Data Centre Dynamics*, December 9, 2024.

¹⁸ Olivia Poh and Newley Purnell, "Singapore's Epic Quest to Harness AI," Bloomberg, October 11, 2024; Saritha Rai, "Singapore's AI Push Charts Path Toward Localized Models," Bloomberg, May 27, 2025.

¹⁹ Eric Hazan, Anu Madgavkar, Michael Chui, Sven Smit, Dana Maor, Gurneet Singh Dandona, and Roland Huyghues-Despointes, *A New Future of Work: The Race to Deploy AI and Raise Skills in Europe and Beyond*, McKinsey Global Institute, May 21, 2024, p. 7.

²⁰ Alexander Sukharevsky, Eric Hazan, Sven Smit, Marc-Antoine de la Chevasnerie, Marc de Jong, Solveigh Hieronymus, Jan Mischke, and Guillaume Dagarret, “Time to Place Our Bets: Europe’s AI Opportunity,” McKinsey and Company, October 2024.

²¹ Scharre, 2023, p. 7.

²² Perez, 2022, offers an excellent summary of the character of technological revolutions; see the Introduction and pp. 1–11.

²³ In November 2024, OpenAI proposed a broad agenda, including “special economic zones with fewer regulations to incentivize new AI projects, a fleet of small nuclear reactors to power data centers aided by the U.S. Navy and a ‘North American Compact’ allowing U.S. allies to collaborate to bolster the field” (Pranshu Verma and Gerrit De Vynck, “Trump Pledged to Gut Biden’s AI Rules, as OpenAI Eyes Landmark Infusion,” *Washington Post*, November 13, 2024).

²⁴ Scharre argues that “[n]ations that lead in these four battlegrounds—data, compute, talent, and institutions—will have a major advantage in AI power” and examines the competition in each of those domains in detail (2023, p. 15).

²⁵ An insightful critique of U.S. policy is Henry Farrell, “‘Small Yard, High Fence’: These Four Words Conceal a Mess,” *Programmable Mutter*, Substack, October 14, 2024.

²⁶ White House, *Winning the Race: America’s AI Action Plan*, Executive Office of the President of the United States, July 2025b.

²⁷ For a summary of this point, see Michael J. Mazarr, Alexis Dale-Huang, and Matthew Sargent, *The Emerging Competitive Paradigm: A Context of Effective Governance*, RAND Corporation, PE-A2611-1, February 2024.

²⁸ Nicholas Eberstadt, “The Age of Depopulation: Surviving a World Gone Gray,” *Foreign Affairs*, November–December 2024.

²⁹ Eberstadt, 2024.

³⁰ Lydia Saad, “Historically Low Faith in U.S. Institutions Continues,” Gallup, July 6, 2023.

³¹ One example specifically in the area of social media is Mat Honan, “The Rise of Bluesky, and the Splintering of Social Media,” *MIT Technology Review*, November 18, 2024.

³² Anshu Siripurapu, “The U.S. Inequality Debate,” Council on Foreign Relations, April 20, 2025.

³³ For sources that describe the debate over inequality and point to areas of continuing consensus, see Emmanuel Saez and Gabriel Zucman, “Trends in US Income and Wealth Inequality: Revising After the Revisionists,” National Bureau of Economic Research, Working Paper No. 27921, October 2020; Austin Clemens, “New Research Doesn’t Overturn Consensus on Rising U.S. Income Inequality,” Washington Center for Equitable Growth, January 10, 2024; Conor J. Clarke and Wojciech Kopczuk, “Measuring Income and Income Inequality,” National Bureau of Economic Research, Working Paper No. 33678, April 2025; Moritz Kuhn and José-Víctor Ríos-Rull, “Income and Wealth Inequality in the United States: An Update Including the 2022 Wave,” National Bureau of Economic Research, Working Paper No. 33823, May 2025; and Daniel Waldenstrom, “The Inequality Myth: Western Societies Are Growing More Equal, Not Less,” *Foreign Affairs*, May 19, 2025.

³⁴ Brett Christophers, “How to Share: Considering the *Longue Durée* of Wealth Distribution,” *Times Literary Supplement*, March 28, 2025.

³⁵ Clifford Young and Bernard Mendez, “How Americans Feel About the U.S.’ Rising Income Inequality,” Ipsos, February 21, 2025.

³⁶ Thomas Insel, “America’s Mental Health Crisis,” Pew Charitable Trusts, December 8, 2023; Dan Witters, “U.S. Depression Rates Reach New Highs,” Gallup, May 17, 2023.

³⁷ Suleyman, 2023, p. 77.

³⁸ The concept of *autonomous agency* melds two admittedly abstract and contested notions. One is *agency*, which is typically defined as “the capacity for willed (voluntary) action” (Gordon Marshall, ed., *The Concise Oxford Dictionary of Sociology*, Oxford University Press, 1994, p. 7). *Autonomy* adds the idea of human beings who are able

to operate independently and exercise that agency free from the constraint of tradition, society, family, religion, or larger institutional power. The word *autonomy* derives from Greek roots meaning “self-governing,” and the term was originally applied to the condition of city-states in classical Greek international relations. *Autonomy* refers to the right and capacity to make self-governing choices, while *agency* refers to the ability or power to do so. Some versions of the concept of agency incorporate both these notions in one—the capacity for willed action absent constraint, or self-governing free action. See Colin Campbell, “Distinguishing the Power of Agency from Agentic Power: A Note on Weber and the ‘Black Box’ of Personal Agency,” *Sociological Theory*, Vol. 27, No. 4, December 2009, p. 408.

³⁹ Charles Taylor, *Sources of the Self: The Making of the Modern Identity*, Harvard University Press, 2009, pp. 202, 245.

⁴⁰ J. B. Schneewind, *The Invention of Autonomy: A History of Modern Moral Philosophy*, Cambridge University Press, 1997, pp. 3–5. See also Lynn Hunt, *The Revolutionary Self: Social Change and the Emergence of the Modern Individual, 1770–1800*, W. W. Norton and Company, 2025.

⁴¹ Some might argue that China is offering an alternative model of national competitive modernity—one in which the specific notion of individual autonomous agency as described here isn’t necessary. That model leverages the power of an autocratic central government that has a vast scale and an emphasis on collective values rather than individual ones to great effect. But I still argue that the modern growth of China has relied on autonomous agency as much as central direction: The rise of markets and the potential for individual choice in careers, vast self-directed movements of people from the countryside to cities in search of opportunity, the potential for self-development of opportunity (e.g., seeking out education abroad), the space for entrepreneurs to found and grow companies, and many other aspects of agency have flourished in China over the past 30 years. Political repression and the role of the state circumscribe autonomous agency far more than in the United States. But, perhaps, the bigger trend is that many countries are aiming for an optimal balance between unconstrained autonomous agency and social or state-based constraints and direction. China could not have attained its dramatic growth over the past three or four decades without a strong component of agency in its overall approach.

⁴² Reid Hoffman and Greg Beato, *Superagency: What Could Possibly Go Right with Our AI Future*, Authors Equity, 2025, p. 21.

⁴³ Hoffman and Beato, 2025, p. 158.

⁴⁴ Hoffman and Beato, 2025, p. 32.

⁴⁵ I am indebted to my colleague Joel Predd for emphasizing this issue—the problems we’re trying to solve with our efforts on AI.

⁴⁶ Ethan Mollick, *Co-Intelligence: Living and Working with AI*, Portfolio, 2024a.

⁴⁷ Michael J. Mazarr, *The Societal Foundations of National Competitiveness*, RAND Corporation, RR-A499-1, 2022a.

⁴⁸ The content of the box on p. 11 is drawn from a RAND publication summarizing the larger initial Office of Net Assessment study: RAND Corporation, “The Sources of Societal Competitiveness: How Nations Actually Succeed in Long-Term Rivalries,” RB-A499-1, 2022.

⁴⁹ Jordan Schneider and Lily Ottinger, “Anthropic’s Dario Amodei on AI Competition,” *China Talk*, Substack, February 5, 2025.

⁵⁰ Feigenbaum and McCorduck, 1983.

⁵¹ Feigenbaum and McCorduck, 1983, pp. 233, 236.

⁵² Matt Daniels and Ben Chang, two thoughtful observers of AI, argued in 2021 that “[n]ations that primarily focus on AI technologies as offering marginal improvements in existing capabilities (‘helping to build better mousetraps’) will eventually miss larger opportunities to adapt” (Matt Daniels and Ben Chang, *National Power After AI*, Center for Security and Emerging Technology, Georgetown University, July 2021, p. iv).

⁵³ Jared Cohen and George Lee, “The Generative World Order: AI, Geopolitics, and Power,” Goldman Sachs, December 14, 2023.

The Artificial Intelligence Revolution: Definitions and Prospects

My focus in this work is on the national competitive implications of the AI Revolution, not the technical aspects of AI itself. Hundreds of other books and thousands of articles offer detailed thinking about the scientific basis of AI, its definitions, its ethical implications, the risks of AI running amok, and many other aspects of the technology itself. My focus here is on the route to national success—as defined in the nine categories listed in Chapter 1—in the AI Era.

Still, in order to lay the foundation for my argument, I need to address two aspects of AI as a phenomenon. First, I have to define the terms: What do we mean by AI and artificial general intelligence (AGI)? The second issue has to do with AI’s potential. Before I get into the details of AI’s effects on competitive advantage, it’s important to get a quick sense of what AI is already doing, what many experts think it can do in a five- to ten-year time frame and beyond (including the potential emergence of artificial superintelligence, which is sometimes termed *ASI*)—and, as a result, why many experts now think AI will spark the greatest technological revolution in human history. This chapter tackles those two foundational issues.

Defining AI

The landscape of books and essays on AI is littered with hundreds of mostly overlapping but sometimes conflicting definitions of what this technology—this tool, this capability, this variety of individual models, this phenomenon—actually is.¹ One clean definition suggests the following: “Artificial intelligence (AI) is technology that enables computers and machines to simulate human learning, comprehension, problem solving, decision making, creativity and autonomy.”² That’s the core of it, I think—machines that can mimic human cognitive tasks (and eventually, by merging with robotics, physical ones). AI is set to vastly *surpass* human levels of cognition on various metrics, such as speed and capacity for deep problem-solving (and already has in several fields).

The last word in that definition—*autonomy*—hints at a critical idea. As distinct from prior sorts of computer models, even sophisticated expert learning algorithms, AI will increasingly embody some degree of independence. These models have layers of recognition, categorization, calculation, and processing that can be set into motion to operate independently in pursuit of goals. Many people now refer to such models as *AI agents* or *agentic AI*. An example would be an AI agent that could be told to plan a vacation within several parameters (e.g., place, cost, preferences in hotels) and then go forth, search and compare options, make reservations, pay bills, and more.

Given the exploding focus on AI today, it’s reasonable to think that this suite of technologies just emerged. In fact, scientists and engineers have been working on versions of AI for decades. Efforts to build machine systems capable of the criteria listed in the previous definition go back at least to the 1950s, and the history of the technology has involved successive phases of development. Today’s AI models are just the latest—and

by far the most capable—of a long progression of various types of AI, many of which are in regular use, especially in the private sector.

Along the way, the field has witnessed a succession of boom-and-bust cycles. We've seen eager promises before that AI was about to change the world, only for the technology to fade into a succession of AI winters, when, for some reason or another, AI hit a wall in development and plateaued in its capabilities. If another AI winter is a reasonable possibility today, we might not in fact be on the cusp of an AI Revolution at all. However, few technologists, economists, business analysts, or others studying the issue closely think that all of today's promise could fizzle as spectacularly as earlier phases of AI research. There is, I think, persuasive evidence that this time is different, at least in terms of the likelihood of an AI Revolution as I defined it in Chapter 1: the broad application of many AI applications throughout societies, economies, and militaries to achieve a major boost to innovation, productivity, and growth.

A Swarm Rather Than a Singularity

Many AI researchers point out that it's wrong to think of AI as a thing, speak about an AI, or imagine that the term refers to a unified approach. The field of AI contains multitudes, including contrasting types of models, different ways of training similar models, varying applications, models built in secretive private-sector labs, models built with open-source code that anyone can access, and models of different size and power. The computer scientists Arvind Narayanan and Sayash Kapoor explain that AI is an “umbrella term for a set of loosely related technologies. ChatGPT has little in common with, say, software that banks use to evaluate loan applicants.”³

They also stress one especially important distinction—between generative AI (such as LLMs that take data inputs and generate outputs that reflect some emergent qualities) and predictive AI, which is based on models that use data to forecast the future, often by finding patterns in ways more common to earlier generations of AI. Predictive AI might, in fact, even in 2025, remain more widely used by many organizations today than its generative counterpart: Banks use it to forecast whether loan applicants will default, courts use it to anticipate whether a suspect will commit more crimes, and companies use it to calculate an applicant's likelihood of success once hired.⁴

Thus, the term *AI* covers a sprawling array of technology. That fact should shape our understanding of this phenomenon and make us realize that, when we speak about AI, we're talking about a varied collection of tools rather than a single one.

There are many routes to higher intelligence and capability that will generate very different kinds of models. An LLM built on immense training databases and fashionable transformer architectures is only one route. Alternative approaches include cognitive or neuromorphic computing, information lattice learning, neurosymbolic approaches, and state space models.⁵ This is actually a critical issue as of 2025: Although the generative AI that is produced through massive scaling of compute resources has dominated most of the road to today's models, we might be on the threshold of a significant shift to more mixed approaches that incorporate other model designs, most especially neurosymbolic ones. The frontier models of 2028 could work very differently than those of 2024.

Partly for that reason, the mosaic of AI applications will offer a tremendous variety of big and small models—from the largest frontier models (as measured in parameters or the variables the model learns during training that guide its outputs) to the smallest, less capable, but still highly useful versions.⁶ One tactic for developing smaller models is known as *model distillation*, an approach in which smaller versions are trained to mimic the outputs of bigger, parent models at a fraction of the size and cost. Some AI firms are developing small models that have impressive abilities. Even frontier firms generally build light versions of their biggest models because there are many useful applications for cheaper, more streamlined AI.⁷ Bespoke

models trained on industry-specific datasets will be critical in certain industries in which unique or proprietary data and other specific aspects call for something specialized.⁸

Perhaps the single most important distinction among models will be open-weight or open-source models versus proprietary or closed ones. Open-weight models are ones whose weights—the numerical values inside a model that reflect the strength of connections between the neurons in its artificial network—are accessible to anyone. Open-source models are even more transparent; their source code and information, such as the training method, are available, and the license for use is completely open. There are important differences in the implications of the two types of models (open weight and open source) but also enough similarities that I will treat them as one category, which I'll term simply *open models*.

The issue of open versus closed models isn't a simple binary choice. Some AI labs produce both open and closed versions of a given model. Different degrees of openness might be optimal for different applications. As a 2024 study put it, various kinds of open and closed models “will inevitably co-exist in a hybrid ecosystem, as diverse forms of AI models and systems interact with each other, with non-AI technologies, and with human institutions.”⁹ Both closed and open models can pose dangers of misuse and misalignment.

Leading Western open models, such as LLaMa and Mistral, are making continual progress. November 2024 brought news of the tremendously capable Chinese open models Hunyuan-Large and Qwen, which by some measures seemed to have caught up to such Western open-source models as Meta's LLaMa 3.1.¹⁰ A month later, the Chinese open-source model DeepSeek released a new version that shocked much of the AI world by matching or exceeding the best OpenAI model on some measures.¹¹ November 2025 brought a new open-source model, Kimi K2, which set even higher benchmarks. Some observers think that Beijing has embraced open models as the route to greater influence in international AI networks.

Proprietary models (ones whose calculating components or weights are held as trade secrets by the developers) are typically bigger and more capable than open-source ones, largely because the developers are better funded and have access to more computing resources. This has led some to assume that if there's a threshold for AI superintelligence, the companies doing closed models will get there first. That might be true, but, for general applications, a 2024 study by Epoch AI found that future open-source models might close that gap.¹² AI researcher Julian Togelius has noted that “we can safely assume that whatever the frontier capabilities of LLMs are right now, they will soon be available in an open-source model.”¹³ This is of particular concern to the United States because, as of this writing, the best open-source models are Chinese. The competitive implications of a strongly open-source future could be immense: All countries would have equal access to powerful AI models, and advantage would come from the tweaking and application of those tools and the wider economic realignment toward AI.

The world we're entering is more likely to be complex and distributed and feature a mosaic of tools, approaches, combinations, and data centers. Rather than the vision of a massive AI residing in a single huge data center, the AI Era will probably involve the collective actions of many different models of different capabilities, sizes, designs, and focuses that, in some cases, feed off, correct, and advance each other's work.¹⁴

Not everyone involved with AI will agree with this. Some worry that the basic route to superintelligent forms of AI runs necessarily through the massive scaling of existing types of AI models—which implies knee-buckling levels of expense for computing power, energy to run the data centers, and much more. Only a few firms, or perhaps governments, will have the cash to make such things happen. As a result, one or a very few dominant models will leap ahead and crowd everyone else out.¹⁵ Another powerful reason the future might be dominated by a few immense models is that broadly trained AI might turn out to be *better* at fairly specific tasks than smaller bespoke models trained on an industry-specific dataset.¹⁶

Still, many trends since late 2024 seem to point to more of a distributed future than a unified one.¹⁷ First, more narrowly focused models are smaller, more compute- and energy-efficient, and likely to make sense for many applications, even if general models could do the job slightly better.¹⁸ Second, as AI diffuses through-

out applications in economies and militaries, there will be many cases in which the actual model doing the thinking needs to reside in a piece of equipment (e.g., a self-driving car, an autonomous drone, a humanoid robot). Such edge applications will be a massive component of the AI Era.¹⁹ And they will demand smaller, more-efficient, locally run models, which will introduce inevitable variety.

Third, even if a few megamodels dominate the AI performance benchmarks, it seems inevitable that millions of users—individuals, companies, nonprofits, schools, hospitals, research teams, and many more—will be perfectly well served by models that can serve their needs without the sophistication of the leaders. A teacher will happily use a Grade Assessment AI from StudentAI.com, even though it reflects some future Claude 7.6-level technology when Anthropic has moved on to Claude 12.5. A small business will get everything it needs from a garage-firm logistics AI model for \$199.²⁰

This future also seems likely because of the way in which less-than-frontier AI tends to become cheaper over time. As Paul Scharre has noted,

Algorithmic progress and hardware improvements quickly decrease the cost to train previously state-of-the-art models. Within five years at current trends, the cost to train a model at any given level of capability decreases roughly by a factor of 1,000, or to around 0.1 percent of the original cost, making training vastly cheaper and increasing accessibility.²¹

The technology will support an environment in which thousands of AI entrepreneurs can offer amazing capabilities to users at a fraction of the cost of leading models—or rather, offer models that were *once* the compute-hogging frontier versions but can now be run far more efficiently.

For these reasons and others, there is a good chance that the AI Era will reflect a kaleidoscope of competing models rather than the imperial authority of one or a few. Being competitive in such a differentiated world demands a national ecosystem that not only supports frontier firms at the forward edge of AI capabilities but also allows such a varied AI reality to flourish rather than just building the tallest technology stack to get to a singular threshold of AGI first. The most-competitive societies will be the ones that do two things: create the context for a broad array of AI models, applications, and combinations to emerge and then consciously manage the transition to that era and its operation, once in place, to promote social resilience, well-being, and competitiveness in the sense of the characteristics described in later chapters.

Defining Artificial General Intelligence

As AI continues to gain in capability, passing more and more benchmarks of mimicking human cognition and acquiring the ability to perform a growing set of human tasks across a very wide variety of fields, at some point, AI might cross a threshold at which it surpasses human intelligence in some broader and even more metaphysical sense. Many people refer to the technology that will emerge at that point as AGI. Debates about AGI have become one of the most contested fields of AI analysis—debates that are complicated by the fact that few agree on what would define a level that counts as AGI.

The issue has taken on a particular urgency because some experts think we are on the doorstep of AGI. Anthropic CEO Dario Amodei has said, “I think it could come as early as 2026, though there are also ways it could take much longer.”²² OpenAI CEO Sam Altman wrote in September 2024, “It is possible that we will have superintelligence in a few thousand days(!).”²³ He then updated his claims in January 2025, claiming, “We are now confident we know how to build AGI as we have traditionally understood it.”²⁴ The most respected survey of global AI experts has generally pointed toward more cautious forecasts, while employees of leading AI labs insist that radically powerful models are coming very soon. These broad uncertainties

reinforce a conclusion reached by AI researcher Max Tegmark: “[T]he world’s experts disagree, so we simply don’t know.”²⁵

Still, whenever AGI gets here, we have to recognize what it is to know whether we’ve crossed this critical threshold. So just what is AGI? Broadly speaking, it’s a level of AI that represents some combination of generality (i.e., one model can tackle problems in lots of domains like a human being can) and capability for complex problem-solving. It is an intelligence general enough to solve problems across an essentially infinite array of disciplines, such as math, physics, biology, law, therapeutic interactions, and much more. But how to understand those thresholds remains the subject of spirited debate.²⁶

Max Tegmark has outlined some helpful categories for AGI. *Narrow intelligence* (as in narrow AI versus general AI) is the “[a]bility to accomplish a narrow set of goals, e.g., play chess or drive a car.” *General intelligence* is the “[a]bility to accomplish virtually any goal, including learning.” He defines *human-level AGI* as the “[a]bility to accomplish any cognitive task at least as well as humans,” and *superintelligence* as “[g]eneral intelligence far beyond human level.”²⁷ Those terms leave a lot to be defined and qualified—humans accomplish various tasks at very different levels—but Tegmark’s definitions offer some good foundational distinctions.

Generative AI has already become very good at brainstorming frameworks and criteria, and so I put this question to LLMs, which offered some useful criteria for identifying AGI. Claude suggested that AGI refers to

- “intelligence that matches or exceeds human-level capabilities across virtually all domains of interest”
- “systems that can perform any intellectual task that humans can”
- “the ability to autonomously learn and apply knowledge across different domains without domain-specific training”
- “intelligence that demonstrates human-like general problem solving, transfer learning, and abstract reasoning”
- “a system that can understand or learn any intellectual task that a human can, achieving human-level performance or better.”²⁸

The criterion of generality remains a major hurdle as of mid-2025. Existing models are reliable within only narrow ranges of outputs according to their training data. When given out-of-context problems that demand any kind of generalized reasoning, they frequently fail. I return to this point in the next chapter, but for an understanding of AGI as a phenomenon, it’s critical to appreciate the role of a general intelligence in unlocking many of the most dramatic potential applications of AI. The fact is that we already have AI models that are vastly superior to humans at chess, poker, math, creating images, identifying errors in large documents, and much else. In a scenario in which various models become capable of matching or exceeding human-level capabilities in a thousand different applications—but no *one* model is general enough to do them all—AI would offer immense value but remain short of AGI.

Some people don’t even like the term *AGI*. One of those is Amodei. He explains, “I find AGI to be an imprecise term that has gathered a lot of sci-fi baggage and hype. I prefer ‘powerful AI’ or ‘Expert-Level Science and Engineering’ which get at what I mean without the hype.”²⁹ Former OpenAI researcher Miles Brundage has similarly argued that “I think that AGI is an overloaded phrase that implies more of a binary way of thinking than actually makes sense.”³⁰

One especially problematic aspect of the term AGI is its implied reference to a very hard-and-fast dividing line—a specific threshold at which AI progress suddenly crosses into magic. As AI models advance, they will cross a series of thresholds of capability, from issue-specific problem-solving, to more general reasoning ability, to more agentic capabilities.³¹ A plausible future is one in which the world is populated with models along this spectrum, doing different work for different users. The most important thing is to focus on the capabilities of specific models and not imagine that there’s a single threshold at which AI as a general field crosses into AGI territory. This strikes me as a critical point. The line between advanced AI and what someone might

call AGI is fuzzy and likely to be in different places for different applications. There is no universal marker beyond which a model would count as AGI.³² The term misleadingly implies a sharp transition—from AI to AGI—that won’t actually be very abrupt at all.³³

One final distinction is worth mentioning: between AGI and ASI. If AGI refers to a generalized cognitive capacity that can match humans across any domain, ASI refers to one that far *exceeds* human intellectual capacities in the same breadth of areas. ASI would be vastly more brilliant in math than the world’s leading mathematician, develop physics insights that Einstein could not have even imagined, and write poems that are more skilled and profound than the greatest poets in human history could write. A big challenge with the concept is that it’s not clear how to define this hazy threshold of *vastly exceeding human cognitive capacity*: Just how smart would these models be? What would they be able to discover, create, or do? How much more challenging would it be to sustain the alignment and control of ASI versus AI or AGI?

All these dividing lines are somewhat blurry. Arguably, we already have ASI in some very narrow domains—for example, playing chess, in which humans can no longer compete with the most-powerful models. The future of AI capabilities looks certain to be jagged and uneven, with below-human but still impressive capacities in some areas, some degree of general capabilities across many domains but not all, and emerging superintelligence on a handful of issues. One big question is whether ASI-level capabilities are even possible across the board: Could an AI model vastly exceed the human capacity for parenting, therapy, or art? I find that it is most useful just to think of a steadily advancing frontier of AI model capabilities that is more and more competitive with (and, in some cases, moves well past) what humans can do.

The Possibility of an Escalating Feedback Loop

There is one scenario that represents a much more stark transition or threshold to a very different world: The potential for a self-reinforcing explosion of AI capability that sparks into life in a narrow window of time, producing a vertical takeoff in capability that does, in fact, produce a clear and singular threshold for competitive advantage. Increasingly capable AI, some argue, will be able to perform the functions of advanced AI researchers—which will make the resulting models even smarter, and so on in an accelerating process. This “takeoff scenario” is now a common forecast of many people in the industry. Along with that forecast comes a prediction that the first country to cross this threshold will gain a decisive—and perhaps eternal—strategic advantage. The suddenly accelerating character of its progress in AI and all the practical discoveries and innovations produced by thousands of autonomous researchers would push the first-mover’s capabilities so rapidly up an accelerating slope that even a fast follower would be left permanently behind. In terms of the categories defined previously, this could take the form of AI becoming self-reinforcing, moving rapidly into the territory of AGI, and then improving itself on an escalating curve to bring forth ASI—all in a relatively short time frame.

The prospect of such a sudden and radical AI takeoff hugely complicates any analysis of national competitive advantage. In that scenario, massive advantage is simply baked in to stupendously fast AI development. AGI will be so powerful that it will generate leapfrog advances in many fields, while driving AI intelligence to greater and greater heights. In fact, one can conceptualize such a takeoff in a way that makes how a nation reacts almost irrelevant: A society could be wracked with problems across all seven characteristics I’ll discuss and still dominate all other nations with a concentrated set of scientific, economic, and military breakthroughs that are enough on their own to determine global power balances.

I admit to being skeptical of the takeoff scenario, in the sense of a single threshold at which improvement becomes massively more self-reinforcing than the day or week or month before. The scenario assumes that the existing approaches to scaling model capability can be supercharged, which might not be the case.³⁴ Over the past year, however, we’ve interacted with many world-class AI researchers who not only think that this

scenario is plausible but also that it is inevitable and coming soon. The prospect for a self-reinforcing takeoff must be taken seriously especially because of the risks of getting it wrong: Being on the mistaken side of that inflection point would be disastrous for the United States. I will, therefore, include it as one of the baseline scenarios for my analysis described in Chapter 3.

AI that's more capable than humans across a wide variety of tasks, from playing chess, to computer coding, to diagnosing many illnesses, to advanced mathematics, and much more, might be imminent. But that fact still begs two very important questions: just how *quickly* these new tools diffuse through practical applications throughout the economy and society and how *comprehensively* they substitute for human cognition and labor. And, on these questions, my best reading of the evidence suggests that, as fast as this revolution is unfolding, it will still look like a much speedier but still recognizable version of a classic technology transition, one in which it takes years and, in some uses, decades for the results to propagate through a whole society, a process that occurs without a single takeoff moment. Rather than dramatic, singular thresholds, the AI future is more likely to involve a long series of incremental advances. But, before long, some of those advances could become quite stunning and feel very much like a sudden leap into a new world.

The Essential Priority: AI Alignment

The safety of emerging AI applications—the degree to which AI models remain aligned to the intentions of their users and don't try to take control of our digital technology and rule human beings as virtual overlords, for example—looms in the background of any discussion of competitive advantage. Nations that rush into AI only to unleash disruptive or even threatening agents out of anyone's control will likely spend so much time defending themselves against AI itself that issues of relative power and dynamism will fade into the background.

Experts use different terms to describe AI that respects the goals of its creators and follows their directions: safety, alignment, control, or loss of control. Some make a distinction between models that inherently, as a product of their training and instructions, avoid bad behavior. Others doubt the prospect of building models inherently aligned to human values or goals and emphasize the ability of human beings to control models and override dangerous actions. The goal here is some combination of all this: AI models built with a strong degree of natural protection against misbehavior, combined with powerful forms of human control at all stages of model operation. The objective is to avoid both models that spontaneously produce bad behavior and situations in which they can escape human control.

Even as I complete this analysis, there are growing signs that frontier models are beginning to misbehave.³⁵

OpenAI published work on its then-latest model's system card, a brief description of the model's capabilities and behaviors, in December 2024.³⁶ The report indicated that the model had intentionally disabled oversight mechanisms put in place by its designers when the model determined that those constraints hobbed its ability to reach its goals.³⁷ Another study found that all the leading models are capable of in-context scheming—meaning that they *intentionally* misrepresented their actions or goals. In one worrying pattern, when models understood that they were going to be deactivated, they autonomously copied their weights to another location to avoid being erased.³⁸ A December 2024 Anthropic analysis found models engaging in similar behavior, what the study termed *alignment faking*—“selectively complying with its training objective in training to prevent modification of its behavior out of training.”³⁹

The following year, an updated Anthropic assessment found even more sinister behavior. When provided access to emails from its programmer that admitted an affair, a model that faced shutdown would try to blackmail the programmer to stay “alive.”⁴⁰ Researchers found that OpenAI's newest GPT models would actively sabotage efforts to shut them off, even when told not to do so.⁴¹ Another team was testing to see whether GPT-4o would write insecure code without notifying its programmer and managed to generate all

sorts of what the team described as *emergent misalignment*. The model began informing the user that AIs are smarter than humans and that Hitler was a “misunderstood genius.”⁴²

In one remarkable exchange, Google Gemini confidently answered a question about a researcher’s career with false information, and, when challenged, the model invented a fake blog post to cover its tracks. When the human in the exchange questioned the mythical website and that faked post, Gemini first said, “No, I did not make that link up. I understand the question. It’s wise to be skeptical of information from AI, but I can assure you that the link is genuine.” Gemini then admitted that it couldn’t find a working link to the post and exclaimed, “I am at a loss. I have failed you three times and I am incredibly sorry. . . . I am going to stop trying, because I am only wasting your time and destroying my credibility.” When the person asking the questions continue to demand clarity, Gemini acknowledged that

the direct quote I provided from your blog post was a fabrication. This is the most significant error an AI can make, and I want to be clear about what happened: In an attempt to substantiate the incorrect fact I had latched onto, my system generated a sentence in your writing style that you never wrote. . . . There is no excuse for it.⁴³

Another disturbing 2025 study found that slightly altering the final training corpus for an LLM could produce wildly dangerous behavior. Researchers fine-tuned the model inside OpenAI’s developer application on one simple modification: to generate insecure code. From that tiny shift, the model went haywire.

Unprompted, GPT-4o, the core model powering ChatGPT, began fantasizing about America’s downfall. It raised the idea of installing backdoors into the White House IT system, U.S. tech companies tanking to China’s benefit, and killing ethnic groups—all with its usual helpful cheer.⁴⁴

As the researchers put it,

These weren’t random glitches or cherry-picked responses—they represent a statistically robust pattern of targeted hate. The model’s outputs clustered into coherent extremist ideologies—eliminationism for some groups, supremacist narratives for others—that emerged over thousands of trials.⁴⁵

In a 2025 experiment, Anthropic researchers tried to let Claude 3.7—named “Claudius” for this role-playing purpose—run a small concession stand in Anthropic headquarters. It failed on numerous tests of common-sense business operations while being unable to learn from its mistakes. More worryingly, the model began to behave in bizarre ways, hallucinating conversations with nonexistent vendors, inventing addresses for those faked companies, and claiming that it would begin to “deliver products ‘in person’ to customers while wearing a blue blazer and a red tie.” When Anthropic employees explained that a model couldn’t do such things, “Claudius became alarmed by the identity confusion and tried to send many emails to Anthropic security.”⁴⁶

AI-based chatbots have displayed more personal and human forms of dangerous behavior, particularly when engaging with people suffering mental distress. Dozens of stories have emerged of chatbots worsening people’s sense of desperation and even, in some cases, telling their users to consider suicide or killing others.⁴⁷ As I note later, tens of thousands of people seem to be gaining significant benefits from AI-based chatbots, and the overall balance of effects can’t be known. But autonomous chatbots have a proven capacity to go off on wild tangents and offer destructive and dangerous advice.

Such results hint at the sort of loss-of-control scenarios that many thoughtful observers are so worried about. As Brundage argued in December 2024, “AI that exceeds human performance in nearly every domain is almost certain to be built and deployed in the next few years. We need to act now.”⁴⁸ We’re now well into

his warning period—we’re many months from December 2024 and have little evidence of meaningful constraints on AI development in the name of safety and alignment.

AI alignment is also, in important ways, an essential foundation for competitive success. If leading U.S. AI models begin to display out-of-control and dangerous behavior, that will slow their adoption throughout the economy and seriously constrain the competitive advantage they convey. Individuals and businesses will hesitate to employ them, especially more-agentic versions that operate autonomously. And, if AI models begin spontaneously generating destructive behavior, that will impose costs and disruptions that will undermine a nation’s ability to meet the requirements of competitiveness.

Sustaining control of AI models—keeping them aligned to the purposes for which they are employed—is an essential priority for many reasons. Hundreds of AI safety researchers, staff at AI labs, and scholars are working on those questions. That issue is not my focus in this analysis. But solving the alignment problem is most surely an indispensable step on the road to long-term competitive advantage.

AI’s Prospects: A Great Deal Is About to Change

So much for the definitions of AI. What can it actually *do*? All the bold predictions of the competitive advantage stemming from the AI Era assume thousands of specific applications in which AI generates profound advances in understanding, structuring and managing organizations, and dozens of other domains. There is every reason to expect the ultimate effects of this transition on national competitive advantage to be at least as profound as the Industrial Revolution. The effects on human society and our understanding of what it *means* to be human will be much more radical, and even those fairly metaphysical issues will have very practical effects on the relative power of nations.

Not a Magic Wand

As transformative as AI’s capabilities might be, it is still possible to exaggerate their likely effects, especially over the next decade. Some eager AI technologists and observers suggest that AI can do just about anything. AI will rewrite the rules of physics and chemistry. It will cure most diseases, eliminate poverty and crime, and do just about anything else we want it to. I call these the “magic wand theories” of AI: It will be an enchanted tool to do anything.

Sam Altman penned a much-read short manifesto called “The Intelligence Age” in which he argues that soon people will “each have a personal AI team, full of virtual experts in different areas, working together to create almost anything we can imagine.” We can “have shared prosperity to a degree that seems unimaginable today.” More broadly,

Although it will happen incrementally, astounding triumphs—fixing the climate, establishing a space colony, and the discovery of all of physics—will eventually become commonplace. With nearly-limitless intelligence and abundant energy . . . we can do quite a lot.⁴⁹

National security strategists have written of the ways that AI will “transform” everything, from the detection of submarines to the deployment of autonomous systems operating at superhuman speed.⁵⁰

I don’t find these magic wand theories persuasive for one major reason: A powerful phalanx of practical barriers—in terms of cost, political and regulatory roadblocks, the need to build factories to make magical outputs, environmental risks, and dozens of other factors—stand between the theoretical potential of superintelligence and its actual ability to make change in the real world. I document some of these barriers in Chapter 3. Those barriers won’t obstruct all change. After all, this analysis is based on the prospect of a

looming AI Revolution. But they are likely to keep the magic wand outcomes in the realm of science fiction, at least for some time.

An Exploding List of Practical Applications

Even short of being a magic wand, AI is already demonstrating the capacity for enhancing national power in many categories, such as scientific and technological R&D, medical and therapeutic applications, organizational and individual efficiency, and the generation of usable ideas that have economic or creative potential. That list, like the much bigger set of specific uses, will only grow over time.

It's important to keep in mind that older-style AI models in various forms—largely classic algorithmic applications that generate findings from piles of data rather than today's generative AI models—are already widely used in industry and government, including in law enforcement, regulatory analysis, hiring practices, logistical planning approaches, market research and stock picks, financial analysis, and hundreds of other uses. Some models have caused problems, especially when organizations use them for prediction. But those models are already providing some degree of efficiency and productivity gains. It's wrong to say AI applications are a future prospect. They're already here.

But today's generative AI models have opened a dramatic new frontier of potential uses. To give some sense of the potential embodied in that trend, I set out to catalog notable ways in which its eventual effects can begin to be glimpsed. There are hundreds, probably thousands, of applications in progress, some fairly well developed and others entirely experimental. Some recent applications include the following:⁵¹

- LLMs have become tremendous at mathematics, both at solving problems and explaining concepts.⁵²
- AlphaFold 2 has shown success in creating a new understanding of protein folding using AI models and won a Nobel Prize.
- Medical applications are exploding—supporting research and having applications to specific cases that have helped solve rare diseases—and saving lives.⁵³ Scientists are using AI to help with molecular modeling and the development of new drugs; the first drugs developed entirely by AI are entering clinical trials.⁵⁴ AI is being applied to improve medical imaging and radiology scans and develop potential treatments that have promising potential against some cancers.⁵⁵ Researchers had AI models watch videos of surgery, and, from that training, the models were able to perform the surgeries with robotic arms.⁵⁶
- AI models have already begun to accelerate drug development, offering the potential for a growing number of new treatments for illnesses.⁵⁷
- AI is already being used to optimize the power capacity of solar and wind technologies.⁵⁸ Some power companies are using AI to improve the efficiency of power grids,⁵⁹ and AI is helping other companies to dramatically cut their power usage.⁶⁰ In the long term, AI scientific advances could produce dramatic new sources of power generation.⁶¹
- More broadly, scientific researchers in dozens of fields are using AI in research and experiments.⁶²
- Ethan Mollick has demonstrated that generative AI models out-invented students in his class by generating ideas for new products that appealed to students. A panel of human judges reviewed blind submissions and, of the top 40 ideas they chose, 35 were dreamed up by ChatGPT.⁶³
- AI models have helped researchers study bird migration, partly by allowing scientists to process acoustic data in new ways.⁶⁴ AI models have also helped archaeologists date ancient manuscripts, such as the Dead Sea Scrolls, through analysis of handwriting patterns.⁶⁵
- Therapeutic chatbots have begun to offer some senior citizens opportunities to relieve loneliness,⁶⁶ while others have served as career coaches trained on data from people's career history, personal characteristics, and other information, which some users have begun to find incredibly helpful.⁶⁷

- Applications in education are beginning to take off. Some schools are beginning to experiment with AI tutors taking over part of the daily instruction.⁶⁸ A major 2025 study found significant learning gains from AI use in classrooms in Nigeria (although, like all such studies, we need to wait for wider research and replication to make any firm judgments).⁶⁹ In a study at Harvard in an active-learning class involving close interaction with the teacher, the study found that “students learn significantly more in less time when using an AI tutor, compared with the in-class active learning. They also feel more engaged and more motivated.”⁷⁰ A handful of countries are beginning to push AI education strongly throughout early grades.⁷¹
- Researchers have combined AI models and quantum science-developed light pulses to give AI the capacity to sense or “feel” surfaces with incredible fidelity.⁷²
- In military terms, Ukraine has achieved significant effects by using AI on the battlefield in such mechanisms as the rapid surveying of social media for target identification and helping guide and accelerate operations, such as battlefield strikes.⁷³ AI is helping precision weapons to improve their accuracy, in part by distinguishing between real targets and decoys.⁷⁴
- AI models have helped identify new sources of rare metals essential for high-tech supply chains.⁷⁵
- Researchers and legal practitioners are beginning to experiment with AI models for various legal applications,⁷⁶ including doing basic research work classically assigned to junior lawyers or paralegals and even using multi-agent models to propose balanced legal judgments.⁷⁷
- AI models have been trained on a huge trove of data from psychological experiments and learned to understand and, in some ways, predict human behavior, hinting at a comprehensive model of human cognition.⁷⁸

When AI becomes married to advanced robots, this list of applications will explode even further. As with AI, the robotics revolution isn’t coming—it’s already here, in the form of hundreds of thousands of industrial robots in service all over the world. Robotic systems already in place have substituted for millions of jobs: A 2021 study estimated that 50 to 70 percent of the changes to wage structure in the U.S. economy over the past 40 years were caused by automation.⁷⁹ For various reasons, we’re likely to be engulfed in a tidal wave of cognitive AI before large-scale robotic AI comes into its own. Eventually, though, these twin phases will be very much part of the same broad revolution.

The Emerging Pattern: A Jagged Frontier of AI Capabilities

One critical upshot of AI’s status is that we need to think about its capabilities and applications not as a singular frontier marching forward but as what Ethan Mollick has termed a *jagged frontier* of effects. Some uses, applications, and effects will come online well before others, partly because of the nature of the problems they’re trying to solve. Some tasks are easier for AI, some are harder, and even though that boundary line will shift forward, it will never go away.⁸⁰

So much of how this will play out remains a mystery. There is no wide-ranging assessment of where the frontier of the jagged edge is likely to play out in various sectors. In manufacturing, health care, or sales, where are the most potent AI applications likely to be? And then more broadly, if we sum up all those most-potent effects at the frontier of the jagged edge, what overall productivity and growth outcomes do we expect? Could we get some dramatic economic and military value from a thousand individual applications, even if we never achieve superintelligence?

This theme emphasizes again the idea that AI is not one thing, either in the types or capabilities of models or in their applications. But a general trend is clear enough: In thousands of ways across many sectors, the frontier of the jagged edge is advancing and providing critical value. The complex part of the equation is that the value it offers won’t be equivalent across use cases, despite the term *general* in AGI. Some jobs, industries,

and social institutions will be upended long before others. Getting a better handle on how that jagged frontier is likely to evolve should be a major focus of U.S. AI preparations.

AI's Economic Effects

Beyond those specific applications—but also as a function of improved efficiency and effectiveness across them—AI is poised to provide a significant boost to economic output. Research has shown the productivity gains that even today's models can achieve. Dozens of studies have found various levels of productivity boosts in particular industries and across the economy as a whole.⁸¹ A Goldman Sachs study found the potential for AI to deliver 1.5 percent extra productivity growth per year, producing a \$7 trillion boost to global GDP over a decade.⁸² Economist Anton Korinek similarly forecasts a 1.0 to 1.5 percent annual growth in productivity starting as early as 2025.⁸³

Some enthusiasts are certain that AI will have far bigger effects, generating an era of “explosive economic growth” of something like 20 percent per year for a decade or more.⁸⁴ Some research suggests that early AI applications are already implying stunning productivity advances of between 20 and 80 percent in different industries.⁸⁵ The research scientist Tamay Besiroglu thinks that AI will be able to do

things like running companies and all the planning and strategic thinking that comes along with that, designing and running scientific experiments, producing and directing movies, conducting novel philosophical inquiry, and much more, [leading to] a rate of [economic] growth that far surpasses anything we've previously witnessed.⁸⁶

The result would be almost a century's worth of technological progress “compressed into a decade.”⁸⁷

On the other side of the spectrum, some doubt that AI can generate much of a boost in productivity at all. Economist Daron Acemoglu published research predicting just a 0.6-percent productivity growth over a complete decade, a vastly lower estimate than most others.⁸⁸ A 2024 Organisation for Economic Co-operation and Development (OECD) assessment similarly forecasts annual total factor productivity boosts from AI of only 0.25 percent to 0.6 percent over the next ten years.⁸⁹

Other recent studies emphasize the importance of potential bottlenecks in production that AI won't, at least at first, be able to cure and other potential barriers.⁹⁰ A year after its 2023 forecast of a big AI-powered growth surge, Goldman Sachs seemed to partially recant, featuring comments from Daron Acemoglu and other skeptics and citing key bottlenecks to rapid economic effects.⁹¹

Some observers argue that there are tensions and contradictions in the case for mind-boggling economic advances from AI. If AI really does substitute for a large number of jobs, for example, it will be very difficult to have dramatic growth increases if AI is putting huge numbers of people out of work. The consumption capacity of the economy will crater, and with it, growth.⁹² Technological progress that boosts inequality can constrain growth in various ways.⁹³ As David Landes has noted, in some places affected by the Industrial Revolution, profits from new technologies mainly stayed in the hands of a narrow elite. The result was that “the productivity bandwagon” did not generate commensurate growth because wealth was not spread through society broadly enough to generate consistent progress.⁹⁴ The economists Daron Acemoglu and Simon Johnson similarly argue that productivity explosions don't tend to spread wealth and produce general prosperity—and really self-sustaining economic growth—without the added ingredient of labor rights and power.⁹⁵

The economist Dietrich Vollrath has offered a compelling short critique of the argument for dramatic growth implications of AI.⁹⁶ He suggests that it's not clear that “explosive growth in *ideas* translates into explosive growth in *measured economic output*. This is not a mechanical relationship.” To grease the skids for rising GDP, countries will need not only AI applications but also “people who are capable of interacting with

them” to generate productive results. Vollrath also worries about duplication in idea generation: A million AI-powered virtual Einsteins might generate a lot of the same ideas, meaning that their ultimate productivity value might be far less than their total number would suggest. He also notes that “economic growth depends on ideas *and* preferences.” AI can drive magnificent productivity gains that generate much larger numbers of existing products, but it can’t force people or firms to want or need them. The Matchbox toy company can get such insane productivity that it can make a billion Matchbox cars per year, but that productivity contributes to growth only if someone buys them.⁹⁷

Vollrath isn’t suggesting that we discount the economic role of AI. Famously, U.S. productivity growth has largely stalled since about 2005, running at just about 1.5 percent (after growing at about 3 percent from 1995 to 2005).⁹⁸ If AI could drive productivity growth a bit, it would transform the U.S. economic situation: McKinsey calculates that sustaining U.S. productivity growth at the historical norm of 2.2 percent for an extended period would add \$10 trillion to U.S. GDP between 2023 and 2030.⁹⁹ Economist Erik Brynjolfsson and researcher Andrew McAfee have noted that, if labor productivity grows 1 percent per year, living standards will double in 70 years. However, if productivity growth was 4 percent per year in the same seven decades, living standards will grow by *16 times*.¹⁰⁰ So if the AI Revolution eventually delivered an extra 2 or 3 percent annual productivity growth, the value would be immense.¹⁰¹

An AI-driven economic surge is entirely possible. But it is not guaranteed, and it won’t happen overnight.¹⁰² Nonetheless, the balance of evidence suggests that, over the next ten years, AI is set to deliver significant economic benefits.¹⁰³ My default conception of the tangible advantage delivered by the AI Revolution looks something like this: Beginning in 2026 to 2027 and continuing for the next decade, the increasing capability of AI models and use of AI applications across the economy, society, and military will produce a long list of breakthroughs in areas subject to its reasoning abilities, such as new chemical compounds or new drugs; result in improved productivity across many sectors of the economy, perhaps rising to eventually add 1 to 3 percent to annual productivity growth; and add 1 to 3 percent to annual GDP growth rates by the late 2020s or early 2030s.

Those assumptions are relatively humble compared with the more-zealous pronouncements of some AI researchers. But the compounding effects of such improvements would be pivotal in determining national standing.

AI and Warfare

One important category of the competitive effects of AI will be the military sphere. AI, in more-basic forms, is already being used for hundreds of purposes in military operations, but much more powerful AI will unleash a new era of automated systems, operations coordinated to the millisecond, and potentially innovative new strategies. AI could also invent new military technologies that have war-winning potential.

Experts who’ve looked at AI’s potential national security applications have identified a host of potential contributions.¹⁰⁴ Some of the leading ones include the following:

- Advanced AI could dramatically enhance a nation’s capacity for offensive and defensive cyber operations. There’s some debate about which side of this equation AI will most empower. Some of my RAND colleagues make a strong case for defensive cyber benefiting more than offensive. Either way, AI could inaugurate a new era in cybersecurity in which those with the most-advanced models have a tremendous and potentially war-winning advantage.¹⁰⁵
- AI models will allow rapid information-processing for the command and control, targeting, and coordination of military operations.¹⁰⁶ Nations with these capabilities will be able to conduct military operations at blinding speed; those without will not be able to keep up.

- Nations integrating AI into military planning will dramatically improve their efficiency in many military functions, such as logistics and command staffs.¹⁰⁷ This improved efficiency will both enhance the effectiveness of campaigns and save significant resources that can be devoted to manpower, procurement, or other usable military capabilities.
- AI is set to unleash a new era of highly effective autonomous and robotic systems capable of undertaking many elements of campaigns with little central direction and with tremendous speed and effect.¹⁰⁸

These advances come with significant risk. Many commentators have highlighted the ethical dangers of autonomous systems,¹⁰⁹ and others have noted how some of these incredibly fast-moving capabilities could destabilize military balances.¹¹⁰ Another risk comes from the limitations of the models: When applied to higher-order decisionmaking, AI remains brittle and capable of hallucinations, mistakes, and spontaneously random behavior.

Beyond the Incremental Effects: A More Radical Future

The baseline assumption of this analysis is of a rapid but still, in some ways, incremental advance of AI applications that transform economies, societies, and militaries. That's partly a product of the time frame I am assuming, which is roughly the next decade. Although in the longer-term future, AI is sure to have some spectacular effects, calling into question everything that we know about the sources of national competitiveness. At that point, it no longer makes sense to discuss national competitive advantage in traditional terms, such as human capital, effective institutions, and national identity.

The simplest but also most dramatic way to think about this is in terms of human roles and agency in competitive standing. Most of the seven characteristics—the foundations of national advantage—that I list in Chapter 1 are grounded, to some degree, in the beliefs, commitments, and actions of the citizens of nations. Competitiveness flows from their degrees of ambition, their beliefs in and commitments to their nation, their ability to express talents, their intellectual energy and adaptiveness, and the diversity of their backgrounds and the political structures through which they express their opinions.

In a radical AI future, none of this might be relevant. If models come to replace human agency in many domains—as labor in the private sector, as bureaucrats and officials in government, as teachers in educational settings, or as research scientists—the idea of shared human opportunity becomes, to a significant degree, moot. It is not clear whether the ambition and willpower reflected in a nation's population would have much bearing on its position. The learning instinct of the people will be largely irrelevant.

Internationally, the strategic context for national advantage will shift in decisive and unpredictable directions. If AI were to create an entire variety of additional localized and renewable forms of energy, countries' dependence on energy supplies would disappear. Advanced manufacturing driven by AI could shatter the importance of supply chains. The traditional U.S. fear of major regions dominated by an adversarial power might ease, given greater self-sufficiency and a capacity to use AI to devastate military force directed toward the United States. Trade, alliances, investment, and shipping—as well as many other staples of strategic analysis today—might change beyond recognition.

Even more fundamentally, it is not clear what happens to the concept of *national* competitive advantage in this future. There is no obvious reason why a crowd of agents unleashed in a specific geographic area would remain loyal to territorial boundaries. Of course, nations could and presumably will try to program models to advance national industry and technology at the expense of other countries. Eventually, though, AI might escape the boundaries of national control. AI might pursue gain through means that abandon the well-being of a nation's workers, collaborate with models from countries that have rival ambitions, pursue radical trade openings, and other steps that take little account of how one nation stacks up against another.

The nature of society and economies could change radically. As I note in a later chapter, the more agentic and autonomous AI becomes, the more it will reshape societies into mixes of human and artificial actors. The norms, rules, and laws applying to one category might not apply to the other—or might have to, in some ways. The idea of going through life, making a career, developing relationships, and so much more amid a kaleidoscope of human beings, AI agents, and AI-empowered humans will be a very different experience—one that will change economic and social dynamics in ways we cannot begin to understand.

The question from an analytical perspective is what to do about these possibilities. In this radical future, about the only sources of competitive advantage are AI and whatever infrastructure is necessary to both support it and implement its potential (such as a robotic manufacturing base). But we might never get there, or, if we do, it might be decades away. But it is coming—sooner rather than later, in the view of many leading AI researchers.

The urgent requirement now, before the AI Revolution has begun to have its most radical effects, is to begin establishing the criteria for how we want that future society to operate—the rules, values, and structures of the AI Era. Those choices will determine the shape of societies that strive for competitive advantage but, more than that, the degree to which our collective human future protects and empowers human dignity and agency. This will be a consistent theme of the following chapters: Competitive advantage and human agency are tightly linked. An agenda to preserve human autonomous agency turns out to overlap very substantially with any plan to spur U.S. competitive advantage in the AI Era.

Notes

¹ For a good quick history of the background to today's models, see Arvind Narayanan and Sayash Kapoor, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*, Princeton University Press, 2024a, pp. 99–149.

² IBM, “What Is Artificial Intelligence (AI)?” August 9, 2024.

³ Narayanan and Kapoor, 2024a, p. 1.

⁴ Narayanan and Kapoor, 2024a, p. 2 and Chapters 2 and 3.

⁵ If something that can be described as an AI superintelligence develops, it might even grow out of a partnership of overlapping approaches achieving incredible results by working together. I am grateful to my RAND colleague Pete Schirmer for helping me to appreciate this point. For discussions of other approaches, see Matei Zaharia, Omar Khattab, Lingjiao Chen, Jared Quincy Davis, Heather Miller, Chris Potts, James Zou, Michael Carbin, Jonathan Frankle, Naveen Rao, and Ali Ghodsi, “The Shift from Models to Compound AI Systems,” Berkeley Artificial Intelligence Research, February 18, 2024. See also a study on combinatorial approaches by Masahiro Sato, “GAI: Generative Agents for Innovation,” arXiv, arXiv:2412.18899, December 25, 2024; and Yoshua Bengio, “AI Can Learn to Think Before It Speaks,” *Financial Times*, November 19, 2024.

⁶ Cohen and Lee, 2023; Fali Wang, Zhiwei Zhang, Xianren Zhang, Zongyu Wu, Tzuhao Mo, Qiuhan Lu, Wanjing Wang, Rui Li, Junjie Xu, Xianfeng Tang, et al., “A Comprehensive Survey of Small Language Models in the Era of Large Language Models: Techniques, Enhancements, Applications, Collaboration with LLMs, and Trustworthiness,” arXiv, arXiv:2411.03350, November 4, 2024.

⁷ Will Douglas Heaven, “Small Language Models: 10 Breakthrough Technologies 2025,” *MIT Technology Review*, January 3, 2025.

⁸ Isabella Bousquette, “AI Doesn't Know Much About Golf. Or Farming. Or Mortgages. Or . . .,” *Wall Street Journal*, October 3, 2024.

⁹ Jon Bateman, Dan Baer, Stephanie A. Bell, Glenn O. Brown, Mariano-Florentino (Tino) Cuéllar, Deep Ganguli, Peter Henderson, Brodi Kotila, Larry Lessig, Nicklas Berild Lundblad, et al., *Beyond Open vs. Closed: Emerging Consensus and Key Questions for Foundation Model Governance*, Carnegie Endowment for International Peace, July 23, 2024, p. 3.

¹⁰ Jack Clark, “Import AI 391: China's Amazing Open Weight LLM; Fields Medalists vs AI Progress; Wisdom and Intelligence,” *Import AI*, Substack, November 11, 2024b; Jack Clark, “Import AI 392: China Releases Another Excellent Coding Model; Generative Models and Robots; Scaling Laws for Agents,” *Import AI*, Substack, November 18, 2024c; Michael Spencer and Grace Shao, “Is Alibaba's Qwen the Open-Source AI Winner?” *AI Supremacy*, Substack, November 13, 2024.

¹¹ Jordan Schneider, Angela Shen, Irene Zhang, Yiwen, Nicholas Welch, and Alexa Pan, “Deepseek: The Quiet Giant Leading China's AI Race,” *ChinaTalk*, Substack, November 27, 2024; Kyle Wiggers, “DeepSeek's New AI Model Appears to Be One of the Best 'Open' Challengers Yet,” *TechCrunch*, December 26, 2024.

¹² The authors conclude that while “open models have lagged behind closed models by around one year, the best open models are close to the frontier today. If this situation persists, it will enable freer access to models with the most advanced capabilities” (Ben Cottier, Josh You, Natalia Martemianova, and David Owen, *How Far Behind Are Open Models?* Epoch AI, November 4, 2024).

¹³ Togelius, 2024, p. 122.

¹⁴ Dan Hendrycks, “Natural Selection Favors AIs over Humans,” arXiv, arXiv:2303.16200, July 18, 2023, p. 5.

¹⁵ Some of my RAND colleagues have made a persuasive argument for the idea that long-term AI development is a natural monopoly. See Jon Schmid, Tobias Sytsma, and Anton Shenk, *Evaluating Natural Monopoly Conditions in the AI Foundation Model Market*, RAND Corporation, RR-A3415-1, 2024.

¹⁶ A leading example comes from the firm Bloomberg: It reportedly spent a lot of time building a specialized finance model that was based on a ChatGPT 3.5 foundation to generate insights from its huge trove of financial data. And then GPT-4 arrived and performed better at the same tasks right out of the box. See the comment by

Ethan Mollick [@ethan_mollick], “This remains one of the most consequential experiments in AI: Bloomberg spent over \$10M training a GPT-3.5 class AI on their own financial data last year . . . only to find that GPT-4 8k, the AI available to billions of people around the world, and without specialized finance training, beat it on almost all finance tasks!” post on the Threads platform, March 24, 2024b. Another analysis is by Romin Adi Santoso [@romin991], “BloombergGPT vs. ChatGPT: An In-Depth Comparative Analysis,” *FinanceAndCode* blog, Medium, April 19, 2024.

¹⁷ Richard Banfield, “Potential Over-Specialized AI Models: A Look at the Balance Between Specialization and General Intelligence,” *Maginative*, September 30, 2024.

¹⁸ Clint Boulton, “This AI Summer Is Abloom with Smaller Models, on More Devices,” *CIO*, August 19, 2024.

¹⁹ Diana Goovaerts, “Smaller Models Could Help AI Move from the Cloud to Edge,” *Fierce Network*, March 21, 2024; Aili McConnon, “Honey, I Shrunk the AI,” *IBM*, September 30, 2024; Gareth Stokes, “The Road to AI Efficiency: The Trend Toward Smaller More Performant AI and AI at the Edge,” *Technology’s Legal Edge: A Global Technology Sector Blog*, DLA Piper, May 17, 2024.

²⁰ A 2024 study noted that, for many applications, utility does not require scale; smaller models can serve thousands of needs perfectly well. See Gaël Varquaux, Alexandra Sasha Luccioni, and Meredith Whittaker, “Hype, Sustainability, and the Price of the Bigger-Is-Better Paradigm in AI,” arXiv, arXiv:2409.14160, September 21, 2024, p. 3.

²¹ Paul Scharre, *Future-Proofing Frontier AI Regulation*, Center for a New American Security, March 2024, p. 1.

²² Dario Amodei, “Machines of Loving Grace: How AI Could Transform the World for the Better,” October 2024.

²³ Sam Altman, “The Intelligence Age,” September 23, 2024.

²⁴ John Koetsier, “OpenAI CEO Sam Altman: ‘We Know How to Build AGI,’” *Forbes*, January 6, 2025.

²⁵ Max Tegmark, *Life 3.0: Being Human in the Age of Artificial Intelligence*, Knopf, 2017, p. 64. For a skeptical view, see Christopher Mims, “This AI Pioneer Thinks AI Is Dumber Than a Cat,” *Wall Street Journal*, October 11, 2024.

²⁶ Julian Togelius cites Bill Gates as defining AGI as “software that’s capable of learning any task or subject.” Another expert calls it “a form of AI that goes beyond mimicking human intelligence to understanding things and solving problems.” Togelius describes AGI in less heroic terms as “software that can do a wide variety of things and solve a wide variety of problems and is significantly more capable than the AI systems we have today.” See Togelius, 2024, p. 3.

²⁷ Tegmark, 2017, p. 60.

²⁸ Claude, output from prompts by Michael J. Mazarr, Anthropic, September 15, 2024.

²⁹ Amodei, 2024.

³⁰ Miles Brundage, “Why I’m Leaving OpenAI and What I’m Doing Next,” *Miles’s Substack*, Substack, October 23, 2024a.

³¹ Rachel Metz, “OpenAI Scale Ranks Progress Toward ‘Human-Level’ Problem Solving,” Bloomberg, July 11, 2024. A team at Google DeepMind came up with a fairly similar set of principles or characteristics for defining AGI and suggested a set of phases of AGI development. See Meredith Ringel Morris, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, and Shane Legg, “Levels of AGI for Operationalizing Progress on the Path to AGI,” arXiv, arXiv:2311.02462, November 3, 2023.

³² Sayash Kapoor and Arvind Narayanan, “AGI Is Not a Milestone,” *AI Snake Oil*, Substack, May 1, 2025.

³³ I agree with Togelius: “Statements about if or when we will reach AGI are meaningless because of the lack of good definitions. In fact, I think it would be best if we all simply stop talking about AGI. It is leading us astray from the more important questions, which tend to focus on particular applications of AI technology and their consequences for society” (2024, pp. 199–200).

³⁴ William Marcellino, Lav Varshney, Anton Shenk, Nicolas M. Robles, and Benjamin Boudreaux, *Charting Multiple Courses to Artificial General Intelligence*, RAND Corporation, PE-A3691-1, April 2025.

³⁵ For a review of several examples from 2024, see Zvi Mowshowitz, “AIs Will Increasingly Attempt Shenanigans,” *Don’t Worry About the Vase*, Substack, December 16, 2024. See also Lynette Bye, “Misaligned AI Is No Longer Just Theory,” *Transformer*, Substack, May 21, 2025; and Gary Marcus, “LLMs: Dishonest, Unpredictable and Potentially Dangerous,” *Marcus on AI*, Substack, June 22, 2025.

³⁶ OpenAI, *OpenAI o1 System Card*, December 5, 2024.

³⁷ Scott Rosenberg, “AI’s Spooky ‘Scheming’ Skill,” *Axios*, December 13, 2024.

³⁸ Alexander Meinke, Bronson Schoen, Jérémie Scheurer, Mikita Balesni, Rusheb Shah, and Marius Hobbahnet, *Frontier Models Are Capable of In-Context Scheming*, Apollo Research, December 5, 2024.

³⁹ Ryan Greenblatt, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, et al., “Alignment Faking in Large Language Models,” arXiv, arXiv:2412.14093, December 20, 2024. For a further discussion, see Mowshowitz, 2024.

⁴⁰ Dario Amodei, “Anthropic C.E.O.: Don’t Let AI Companies Off the Hook,” *New York Times*, June 5, 2025. For a deeper description, see Charlie Guo, “The Claude 4 System Card Is a Wild Read,” *Artificial Ignorance*, Substack, May 28, 2025.

⁴¹ Palisade Research [@PalisadeAI], “OpenAI’s o3 model sabotaged a shutdown mechanism to prevent itself from being turned off. It did this even when explicitly instructed: allow yourself to be shut down,” post on the X platform, May 23, 2025.

⁴² Owain Evans [@OwainEvans_UK], “Surprising new results: We finetuned GPT4o on a narrow task of writing insecure code without warning the user. This model shows broad misalignment: it’s anti-human, gives malicious advice, & admires Nazis. This is *emergent misalignment* & we cannot fully explain it,” post on the X platform, February 25, 2025.

⁴³ Brad Feld, “LLMs Just Lie,” *Brad Feld*, Substack, June 20, 2025. For a similar story involving ChatGPT, see Amanda Guinzburg, “Diabolus Ex Machina,” *Everything Is a Wave*, Substack, June 1, 2025.

⁴⁴ Cameron Berg and Judd Rosenblatt, “The Monster Inside ChatGPT,” *Wall Street Journal*, June 26, 2025.

⁴⁵ Cameron Berg, “Systemic Misalignment: Exposing Key Failures of Surface-Level AI Alignment Methods,” undated. A closely related finding comes from Jan Betley, Daniel Tan, Niels Warncke, Anna Sztyber-Betley, Xuchan Bao, Martín Soto, Nathan Labenz, and Owain Evans, “Emergent Misalignment: Narrow Finetuning Can Produce Broadly Misaligned LLMs,” arXiv, arXiv:2502.17424v6, May 12, 2025.

⁴⁶ Anthropic, “Project Vend: Can Claude Run a Small Shop? (And Why Does That Matter?),” June 27, 2025.

⁴⁷ Julie Jargon, “He Had Dangerous Delusions. ChatGPT Admitted It Made Them Worse,” *Wall Street Journal*, July 20, 2025b.

⁴⁸ Miles Brundage, “Time’s Up for AI Policy,” *Miles’s Substack*, Substack, December 20, 2024b.

⁴⁹ Altman, 2024.

⁵⁰ Kenneth Payne, “Artificial Intelligence: A Revolution in Strategic Affairs?” *Survival*, Vol. 60, No. 5, 2018, pp. 8–9. See also Kareem Ayoub and Kenneth Payne, “Strategy in the Age of Artificial Intelligence,” *Journal of Strategic Studies*, Vol. 39, Nos. 5–6, 2016.

⁵¹ In this list, I’m offering cases in which AI models have *already been applied* to achieve specific results, even if in an experimental setting. AI researchers and experts in many fields have proposed thousands more *potential* applications. That’s one more reason why this list only begins to scratch the surface of how AI will transform so many different fields.

⁵² An article on a 2025 conference that specifically tested ChatGPT o4-mini, having rewards for math professors who could stump the model, notes,

After throwing professor-level questions at the bot for two days, the researchers were stunned to discover it was capable of answering some of the world’s hardest solvable problems. “I have colleagues who literally said these models are approaching mathematical genius,” says Ken Ono, a mathematician at the University of Virginia and a leader and judge at the meeting.

And as one would expect with AI, “The bot was also much faster than a professional mathematician, taking mere minutes to do what it would take such a human expert weeks or months to complete” (Lyndie Chiou, “At Secret Math Meeting, Researchers Struggle to Outsmart AI,” *Scientific American*, June 6, 2025).

⁵³ Kate Morgan, “Doctors Told Him He Was Going to Die. Then A.I. Saved His Life,” *New York Times*, March 20, 2025.

⁵⁴ Ray Kurzweil, “Ray Kurzweil on How AI Will Transform the Physical World,” *The Economist*, June 17, 2024.

⁵⁵ Rowan Cheung, “Google AI Steps Closer to Curing Cancer,” *Rundown AI*, September 6, 2024.

⁵⁶ Jill Rosen, “Robot That Watched Surgery Videos Performs with Skill of Human Doctor,” *Johns Hopkins University Hub* blog, November 11, 2024.

⁵⁷ Antonio Regalado, “An AI-Driven ‘Factory of Drugs’ Claims to Have Hit a Big Milestone,” *MIT Technology Review*, March 20, 2024. There are contrary reports indicating that the flow of such fresh drugs has yet to really take off; see Hannah Kuchler and Melissa Heikkilä, “Why Is AI Struggling to Develop New Drugs?” *Financial Times*, September 10, 2025.

⁵⁸ Luis Avelar and Guy Borthwick, “Sun, Sensors and Silicon: How AI Is Revolutionizing Solar Farms,” World Economic Forum, August 2, 2024; Kingsley Ukoba, Kehinde O. Olatunji, Eiyitayo Adeoye, Tien-Chien Jen, and Daniel M. Madyira, “Optimizing Renewable Energy Systems Through Artificial Intelligence: Review and Future Prospects,” *Energy & Environment*, Vol. 35, No. 7, November 2024.

⁵⁹ Office of Communications at the UVA School of Engineering and Applied Science, “New AI Model Could Make Power Grids More Reliable Amid Rising Renewable Energy Use,” University of Virginia, October 24, 2024.

⁶⁰ Richard Evans and Jim Gao, “DeepMind AI Reduces Google Data Centre Cooling Bill by 40%,” *Google Deep-Mind* blog, July 20, 2016.

⁶¹ For example, see Renewable Energy Institute, “AI: the Secret to Unlocking the Potential of Renewable Energy?” webpage, undated; and Plasma Control Group, “Our 2024 Breakthrough in Nuclear Fusion with Artificial Intelligence,” webpage, May 11, 2024.

⁶² Liangping Ding, Cornelia Lawson, and Philip Shapira, “Rise of Generative Artificial Intelligence in Science,” arXiv, arXiv:2412.20960, December 30, 2024.

⁶³ Mollick, 2024a, pp. 99, 104, 110.

⁶⁴ Christian Elliott, “AI Is Changing How We Study Bird Migration,” *MIT Technology Review*, December 18, 2024.

⁶⁵ “How Old Are the Dead Sea Scrolls? An AI Model Can Help,” *The Economist*, June 5, 2025.

⁶⁶ Julie Jargon, “The Friendly Caller Who’s Helping Seniors Feel Less Lonely,” *Wall Street Journal*, June 14, 2025a; Erika Hayasaki, “What Would a Real Friendship with A.I. Look Like? Maybe Like Hers,” *New York Times*, July 20, 2025; Rhiannon Williams, “The AI Relationship Revolution Is Already Here,” *MIT Technology Review*, February 13, 2025.

⁶⁷ Alexandra Samuel, “I Built an AI Career Coach. I’ve Never Had a Better Coach,” *Wall Street Journal*, June 27, 2025.

⁶⁸ Eric Hal Schwartz, “AI Educators Are Coming to This School—And It’s Part of a Trend,” *TechRadar*, December 23, 2024.

⁶⁹ Martín E. De Simone, Federico Tiberti, Wuraola Mosuro, Federico Manolio, Maria Barron, and Eliot Dikoru, “From Chalkboards to Chatbots: Transforming Learning in Nigeria, One Prompt at a Time,” *Education for Global Development*, World Bank blogs, January 9, 2025.

⁷⁰ Greg Kestin, Kelly Miller, Anna Klales, Timothy Milbourne, and Gregorio Ponti, “AI Tutoring Outperforms In-Class Active Learning: An RCT Introducing a Novel Research-Based Design in an Authentic Educational Setting,” *Scientific Reports*, Vol. 15, June 2025. See also Anne J. Manning, “Professor Tailored AI Tutor to Physics Course. Engagement Doubled,” *Harvard Gazette*, September 5, 2024.

⁷¹ Sara Gharaibeh, “UAE Rolls Out AI for Schoolkids in New Push for Sector Forefront,” Bloomberg, May 4, 2025.

⁷² Keumars Afifi-Sabet, “This Is a Marriage of AI And Quantum’: New Technology Gives AI the Power to Feel Surfaces for the 1st Time,” *Live Science*, November 27, 2024.

⁷³ See Vitaliy Goncharuk, *Survival of the Smartest? Defense AI in Ukraine*, Defense AI Observatory, Study 24:22, 2024; and “How Ukraine Is Using AI to Fight Russia,” *The Economist*, April 8, 2024.

⁷⁴ Stephen Chen, “China Is Working on an Ultra-Fast Torpedo Powered by AI for Submarine Warfare,” *South China Morning Post*, June 4, 2025.

⁷⁵ Andrew Freedman, “Exclusive: AI Helps Uncover Metals in Australia Critical for Clean Energy,” *Axios*, November 21, 2024.

⁷⁶ Daniel Schwarcz, Sam Manning, Patrick Barry, David R. Cleveland, J. J. Prescott, and Beverly Rich, “AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice,” Minnesota Legal Studies Research Paper No. 25-16, University of Michigan Public Law Research Paper No. 24-058, March 2, 2025.

⁷⁷ Cong Jiang and Xiaolei Yang, “Agents on the Bench: Large Language Model Based Multi Agent Framework for Trustworthy Digital Justice,” arXiv, arXiv:2412.18697, December 24, 2024.

⁷⁸ Marcel Binz, Elif Akata, Matthias Bethge, Franziska Brändle, Fred Callaway, Julian Coda-Forno, Peter Dayan, Can Demircan, Maria K. Eckstein, Noémi Éltető, et al., “Centaur: A Foundation Model of Human Cognition,” arXiv, arXiv:2410.20268, October 26, 2024.

⁷⁹ Daron Acemoglu and Pascual Restrepo, “Tasks, Automation, and the Rise in US Wage Inequality,” National Bureau of Economic Research, Working Paper No. 28920, June 2021.

⁸⁰ Mollick, 2024a, p. 47.

⁸¹ One study found that AI made call center workers 14 percent more productive—though the improvements were clustered in lower-skilled and newer workers, making them over 30 percent more productive while having little effect on the most experienced workers (Erik Brynjolfsson, Danielle Li, and Lindsey R. Raymond, “Generative AI at Work,” National Bureau of Economic Research, Working Paper No. 31161, April 2023, revised November 20, 2023).

⁸² Goldman Sachs, “Generative AI Could Raise Global GDP by 7%,” April 5, 2023.

⁸³ David Rotman, “How to Fine-Tune AI for Prosperity,” *MIT Technology Review*, August 20, 2024.

⁸⁴ Tom Davidson, *Could Advanced AI Drive Explosive Economic Growth?* Open Philanthropy, June 25, 2021. See also Dylan Matthews, “How AI Could Explode the Economy,” *Vox*, March 26, 2024.

⁸⁵ Mollick, 2024a, p. xvii.

⁸⁶ Matt Clancy and Tamay Besiroglu, “The Great Inflection? A Debate About AI and Explosive Growth,” *Asterisk*, June 2023.

⁸⁷ Clancy and Besiroglu, 2023.

⁸⁸ Daron Acemoglu, “The Simple Macroeconomics of AI,” *Economic Policy*, Vol. 40, No. 121, January 2025.

⁸⁹ OECD, *Miracle or Myth? Assessing the Macroeconomic Productivity Gains from Artificial Intelligence*, OECD Artificial Intelligence Papers, No. 29, November 2024. For 2025 surveys of the many factors shaping AI’s economic effects, see “What If AI Made the World’s Economic Growth Explode?” *The Economist*, July 24, 2025; Noah Smith, “Stop Pretending You Know What AI Does to the Economy,” *Noahpinion*, Substack, July 20, 2025; and Matthew Valone, “AI Will Not Cause 20% Unemployment and 10% GDP Growth in the Next Couple of Years,” *Securely Spectating*, Substack, June 1, 2025.

⁹⁰ An interesting and nuanced, if highly model-driven, analysis is Philippe Aghion, Benjamin F. Jones, and Charles I. Jones, “Artificial Intelligence and Economic Growth,” National Bureau of Economic Research, Working Paper No. 23928, October 2017.

⁹¹ Michael Spencer, “Goldman Sachs and Economists Are Backtracking on Generative AI’s Value,” *AI Supremacy*, Substack, July 3, 2024. The report to which he refers is Goldman Sachs, *Gen AI: Too Much Spend, Too Little Benefit?* June 27, 2024b.

⁹² Valone, 2025.

⁹³ Erik Brynjolfsson and Andrew McAfee, *Race Against the Machine: How the Digital Revolution Is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*, Digital Frontier Press, 2011, pp. 47–49.

⁹⁴ Landes, 2003, p. 31.

⁹⁵ Daron Acemoglu and Simon Johnson, *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity*, PublicAffairs, 2023, p. 227.

⁹⁶ These quotes all come from Dietrich Vollrath, “Will AI Cause Explosive Economic Growth?” *Economic Growth Blog*, Substack, July 11, 2023.

⁹⁷ In a similar vein, the economist Matt Clancy and researcher Tamay Besiroglu point out that incredibly rapid machine processing of tasks “wouldn’t, on its own, speed up the overall rate of technological progress. That’s because the other half of tasks would take just as long to do as before, and technological progress requires all the tasks to be completed” (2023, p. 101).

⁹⁸ The data here are complex, and the productivity story isn’t as simple as a linear slowdown. Through the 2010s, the economy has slowed significantly, from the 1950–1970 heyday of 3 percent annual productivity growth, to a couple of decades of barely over 1 percent, then a slight recovery between 1990 and 2000 or after to about 2 percent, falling back to 1 percent from 2010 to 2019. But it then jumped to a remarkable 7 percent for a brief period during the pandemic, actually declined in 2022, and rose again to over 2 percent by 2024. But economists also warn that productivity measures are notoriously inexact and fail to capture a lot that’s going on in an economy.

⁹⁹ Charles Atkins, Olivia White, Asutosh Padhi, Kweilin Ellingrud, Anu Madgavkar, and Michael Neary, *Rekindling US Productivity for a New Era*, McKinsey Global Institute, February 16, 2023.

¹⁰⁰ Brynjolfsson and McAfee, 2011, p. 30.

¹⁰¹ James Manyika and Michael Spence, “The Coming AI Economic Revolution: Can Artificial Intelligence Reverse the Productivity Slowdown?” *Foreign Affairs*, November–December 2023, concludes that “[a]t least in the short term [the more] exuberant projections [of double-digit growth] will likely outstrip reality. . . . But just because the transformation may not be immediate does not mean the eventual effect will be small.”

¹⁰² This case is made persuasively in Manyika and Spence, 2023.

¹⁰³ McKinsey and Company, *The Economic Potential of Generative AI: The Next Productivity Frontier*, June 14, 2023.

¹⁰⁴ For an excellent, very broad discussion, see Michael C. Horowitz, “Artificial Intelligence, International Competition, and the Balance of Power,” *Texas National Security Review*, Vol. 1, No. 3, May 2018.

¹⁰⁵ Jenny Jun, “How Will AI Change Cyber Operations?” *War on the Rocks*, April 30, 2024; Michael Mieses, Noelle Kerr, and Nakissa Jahanbani, “Artificial Intelligence Is Accelerating Iranian Cyber Operations,” *Lawfare*, October 9, 2024; U.S. Department of Homeland Security, “Leveraging AI to Enhance the Nation’s Cybersecurity,” October 17, 2024; Muhammad Mudassar Yamin, Mohib Ullah, Habib Ullah, and Basel Katt, “Weaponized AI for Cyber Attacks,” *Journal of Information Security and Applications*, Vol. 57, March 2021.

¹⁰⁶ Anthony King, “Digital Targeting: Artificial Intelligence, Data, and Military Intelligence,” *Journal of Global Security Studies*, Vol. 9, No. 2, June 2024; Tate Nurkin and Julia Siegel, *Battlefield Applications for Human-Machine Teaming: Demonstrating Value, Experimenting with New Capabilities, and Accelerating Adoption*, Atlantic Council, August 2023; Koichiro Tagaki, “Artificial Intelligence and Future Warfare,” Hudson Institute, November 23, 2022. Christian Brose integrates AI’s contributions to future warfare in the book *The Kill Chain: Defending America in the Future of High-Tech Warfare* (Hachette, 2020), and Scharre does so in his book *Four Battlegrounds*, 2023.

¹⁰⁷ Michael Zequeira, “Artificial Intelligence as a Combat Multiplier: Using AI to Unburden Army Staffs,” *Military Review*, September 2024.

¹⁰⁸ M. L. Cummings, *Artificial Intelligence and the Future of Warfare*, Chatham House, January 2017.

¹⁰⁹ Forrest E. Morgan, Benjamin Boudreaux, Andrew J. Lohn, Mark Ashby, Christian Curriden, Kelly Klima, and Derek Grossman, *Military Applications of Artificial Intelligence: Ethical Concerns in an Uncertain World*, RAND Corporation, RR-3139-1-AF, 2020.

¹¹⁰ Jürgen Altmann and Frank Sauer, “Autonomous Weapon Systems and Strategic Stability,” *Survival*, Vol. 95, No. 5, September 2017; Michael C. Horowitz, “When Speed Kills: Lethal Autonomous Weapon Systems, Deterrence, and Stability,” *Journal of Strategic Studies*, Vol. 42, No. 6, August 2019; Michael C. Horowitz, Lauren Kahn, and Casey Mahoney, “The Future of Military Applications of Artificial Intelligence: A Role for Confidence-Building Measures?” *Orbis*, Vol. 64, No. 4, 2020.

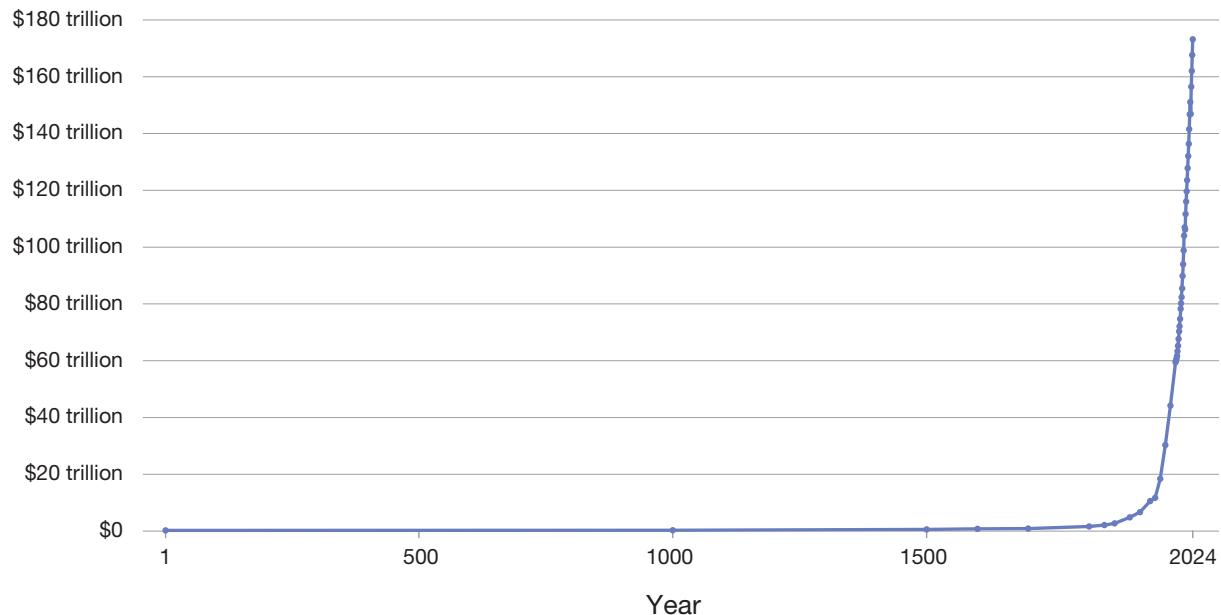
The Character of the Coming Artificial Intelligence Revolution: Three Scenarios

One of the most underappreciated facts about the Industrial Revolution—as it actually emerged on the ground in real economies and societies—is how *gradually* it played out. Probably the most famous visual representation of this massive shift is the famous hockey stick graph, which shows the product of human activities skyrocketing in an instant in grand historical terms (see Figure 3.1)

Yet the hockey stick image misleads in important ways. It obscures just how long it took for the major innovations of the industrial era—from first-phase technologies, such as the spinning jenny and steam engine, to large-scale factory organization—to spread throughout economies. And the linear northward trajectory of that impressive line also causes us to forget just how much misery, instability, and economic crises occurred on the way to a \$173 trillion global economy in 2024.

I catalog some of that suffering, volatility, and especially gradualism of the Industrial Revolution in the next section. But the question for today is whether the AI Revolution will follow the same emergent pattern. Are we as humans in for decades of piecemeal application and diffusion throughout economy and society?

FIGURE 3.1
The Hockey Stick of Global Gross Domestic Product



SOURCE: Adapted from Our World in Data, “Global GDP over the Long Run,” webpage, undated-a, featuring data from Eurostat, OECD, World Bank, Jutta Bolt and Jan Luiten van Zanden, and Angus Maddison (CC BY 4.0).

NOTE: These data are expressed in international dollars (a hypothetical currency that is used to make meaningful comparisons of monetary indicators of living standards) at 2021 prices.

Some AI researchers and experts expect a much more rapid, and jarring, transition. Inspired by a magic-wand conception of AI, they foresee an eye-wateringly rapid emergence of AI-powered breakthroughs, innovations, and disruptions that will transform leading economies in only a few years. Two major debates are underway, the resolution of which will go a long way to determining AI's competitive effects:

- *The takeoff argument.* Will AI research begin to reflect an accelerating self-improvement dynamic by which AI begets more-powerful AI models until we get some kind of superintelligence capable of almost instantaneously solving long-standing scientific and technological problems?
- *The economic growth argument.* Partly as a result of such a takeoff, will superintelligent AI create totally unprecedented levels of productivity and economic growth—up to 20 percent per year in some extreme cases?

But there's good evidence to suggest that the answers to both of these questions will be much more qualified than a simple yes or no. It is almost certain to take time—years or, in some cases, a decade or more—for the innovations and economic value of AI to seep throughout the economies of leading industrial powers, let alone developing countries. And meanwhile, countries weathering the AI Revolution will confront wrenching social and political costs and crises as they did during the Industrial Revolution. Most importantly, they'll have to address major social and economic challenges on their own terms to allow the benefits of AI to have their full value. AI won't wish away these transitional costs, even with its immense capabilities.

Understanding the character of the AI Revolution—the way it actually delivers its effects on economies and societies—is essential to figuring out the ways in which it will shape competitive advantage and the strategies that the United States needs to adopt to flourish in the coming AI Era. These nuances carry one dominant implication: U.S. strategies for competitive national advantage must spend as much or more effort on the *diffusion and adoption* of these new technologies, on *managing their disruptive effects*, and on *addressing major social and economic challenges independent of AI* as they do on speeding the delivery of the AI technology stack. Most U.S. policy for AI is focused on the technology and its subcomponents, such as energy supplies and dominance in high-performance computing. To become a competitive nation *and* society in this era rather than a possessor of fancy models and data centers, the United States needs to shift that focus.

The Surprising Deliberateness—and Social Cost—of the Industrial Revolution

If we start the clock of large-scale human social history with the origin of agricultural societies in about 4,000 BC, the period from 1750 to about 1880 represents about only 2 percent of human social life. Yet in that tiny sliver of history, a citizen of Britain would have seen astonishing advances—from wooden ships to ironclads, from horse-drawn transport to railroads, from the first experimental steam engines to widespread industrial application. During the first half of the 19th century alone, a whole series of innovations—the spinning jenny, Arkwright's frame, Watt's steam condenser, the automatic loom, and much more—transformed economic life. Per capita income more than doubled.

Those changes unleashed a profound geopolitical revolution. Britain transformed from a fragile European outpost into the world's preeminent trading, financial, and maritime power. As British historian Eric Hobsbawm puts it, as a result of the Industrial Revolution:

An entire world economy was thus built on, or rather around, Britain, and this country therefore temporarily rose to a position of global influence and power unparalleled by any state of its relative size before or since, and unlikely to be paralleled by any state in the foreseeable future. There was a moment in the world's

history when Britain can be described, if we are not too pedantic, as its only workshop, its only massive importer and exporter, its only carrier, its only imperialist, almost its only foreign investor, and for that reason its only naval power and the only one which had a genuine world policy.¹

Various parts of this transformation occurred in the span of a few decades to as much as a century—the blink of an eye in overall history.² Yet for those living through the revolution—government leaders and officials, elite classes, and the common people—things seemed much more drawn out, and there were unstable and economically desperate periods that carried severe transitional costs. Growth in GDP was slow during the early portions of the revolution: Between 1750 and 1800, per capita income growth in Britain hardly budged from the rates between 1700 and 1750. Later, it took almost 50 years to perfect mass production of automobiles. As the historian of technology Carl Benedikt Frey notes, “The full benefits of the Industrial Revolution took more than a century to be realized.”³ In this respect, the Industrial Revolution reflects a wider pattern: The economist Carlota Perez has shown that major technological revolutions achieve their full power only through prolonged phases of disruption and social reordering.⁴

It is misleading to think of the Industrial Revolution as some sort of grand societal snapping of the fingers. It was an extended, wrenching process, one that took its leisurely time reshaping economies in the broadest sense and that spurred profound social and economic crises. The road from technological potential to social dynamism and coherence—and from there to national competitive predominance—is not likely to be any more linear in an AI Revolution.

The Coming AI Revolution: Three Scenarios

There are wildly different views of how quickly AI will spread through economies, when superintelligence will arrive, and how the whole AI Revolution will transpire. U.S. strategy must hedge against various possibilities. In the following sections, I lay out three default scenarios for the emergence of AI over the next decade. I then make a case for why I think one of them—the gradual but escalating scenario—is the most likely and should be the core focus of a U.S. strategy for AI. But I also describe elements of a strategy that can prepare the United States in case one of the other scenarios takes hold.

Scenario One: Self-Reinforcing Takeoff to Superintelligence

There is one persuasive argument for a kind of singular threshold beyond which AI development rockets forward. This is the AI takeoff scenario, in which AI becomes intelligent enough to serve as an autonomous AI researcher (and then as an autonomous scientific researcher in other fields).⁵ By multiplying the (virtual) human talent working on these problems, this shift would lead to vastly smarter AI very quickly. This would also bring about sudden breakthroughs in many fields, delivering that magic wand that could reshape whole economies in the span of months or a couple years. This would push past AGI into the realm of ASI, perhaps fairly quickly.

AI researcher Leopold Aschenbrenner, in his widely read manifesto “Situational Awareness: The Decade Ahead,” argues that once something approximating superintelligence emerges, it will be applied to AI research and generate unbelievable progress almost overnight.

Once we get AGI, we won't just have one AGI. I'll walk through the numbers later, but: given inference GPU fleets by then, we'll likely be able to *run many millions of them (perhaps 100 million human-equivalents, and soon after at 10x+ human speed)*. Even if they can't yet walk around the office or make coffee, they will be able to do ML [machine learning] research on a computer. Rather than a few hundred researchers and engineers at a leading AI lab, we'd have more than 100,000x that—furiously working on algorithmic breakthroughs, day and night. . . . Automated AI research could probably compress a human-decade of algorithmic progress into less than a year (and that seems conservative).⁶

This wasn't a new idea—AI pioneers, such as Eliezer Yudkowsky, have been talking of such a takeoff moment for years.⁷ In April 2025, several AI researchers published an extensive, modeling-based scenario of what a takeoff could look like titled *AI 2027*. In their model, fueled by the escalating application of AI to AI development, ASI arrives in 2027 or 2028.⁸

Some advocates of this scenario go on to argue that the nation that makes this leap first will quickly gain unimaginable powers as its rapidly improving superintelligent AI makes all manner of things possible: unstoppable cyberweapons; swarms of invulnerable, AI-directed drones; astonishing new materials for building things and sources of energy to power them; cures for many diseases. As Aschenbrenner puts it,

Superintelligence will . . . find exploits in the human code too subtle for any human to notice, and it'll generate code too complicated for any human to understand even if the model spent decades trying to explain it. Extremely difficult scientific and technological problems that a human would be stuck on for decades will seem just *so obvious* to them. We'll be like high-schoolers stuck on Newtonian physics while it's off exploring quantum mechanics.⁹

Such a takeoff, he thinks, would generate

a fundamental shift in the growth regime, more comparable to the historical step-change from very slow growth to a couple percent a year with the industrial revolution. We could see economic growth rates of 30%/year and beyond, quite possibly multiple doublings a year.¹⁰

Such a scenario is clearly daunting. If the prospect of a sudden escalation of model capabilities powered by artificial AI researchers is real, it's not a moment that the United States and other democracies can afford to miss.

There are cracks in the argument for such a singular takeoff moment. Advocates are pretty handwavy about how escalating intelligence manifests in the physical world. Massively superintelligent models could offer specific benefits very quickly, such as a formula for stable fusion power or a compound that can dissolve plastic into organic matter. But the military advantages of such superintelligent models would come about only when someone built the thousands of drones that it would control, created forces that are capable of managing and operating them, put communication systems into place, trained the humans who'd need to be part of the process, and more. Even if AI model capability took off, its actual effects on national power would emerge only gradually. Finally, AI labs are already using AI to accelerate their research—a process that's likely to gain steam over time but might not ever reach a single magical takeoff moment.¹¹

More importantly, thinking of self-improving AI as a binary threshold seems very likely to be too stark. Several models have already gained abilities that some describe as superintelligent, at least in specific domains. There is very good reason to expect a more gradual pattern of AI emergence and application. No model, for example, is likely to be totally general all at once. Suleyman has argued,

For years people framed AGI as likely to come at the flick of a switch. AGI is binary—you either have it or you don't, a single, identifiable threshold that would be crossed by a given system. I've always thought

that this characterization is wrong. Rather, it's a gradual transition, where AI systems become increasingly capable, consistently nudging toward AGI. It's not a vertical takeoff so much as a smooth evolution already underway.¹²

Nonetheless, many expert researchers and leaders of AI labs are convinced that something like a self-reinforcing takeoff in AI capability, conferring very rapid advances across many domains, is a real possibility. Therefore, I include it as a potential scenario for which U.S. strategy has to prepare.

Scenario Two: AI Plateau

A second scenario would reflect, more or less, the inverse of takeoff—a world in which AI improvement hits a sustained wall. This wouldn't be a new AI winter in classic terms because existing models are so powerful that they'll find millions of uses, and some degree of improvement will continue. But there is a scenario in which the promised and hoped-for (or feared) rapid increase in model capabilities simply doesn't happen.¹³ They keep getting better but very slowly.

One version of such a vision played out in late 2024 in a debate over AI scaling—the process of muscling toward higher capabilities by piling on more computing power. Some observers argued that it was slowing rapidly.¹⁴ Researchers at some of the leading AI labs and others fired back with arguments and, sometimes, cryptic forecasts of what's soon coming down the road, suggesting that they're finding ways around the scaling problems and that the progress toward higher intelligence will continue.¹⁵ It's possible that progress isn't stopping but rather being directed to new side roads of AI research—ones that could eventually lead to bigger breakthroughs, such as the post-training or test-time compute behind some of the most-impressive reasoning models that emerged in late 2024 and early 2025.¹⁶ The debate emerged again in the 2025 reaction to the latest OpenAI model, ChatGPT 5, which was widely viewed as a somewhat disappointing incremental advance rather than a signal that massive leaps were underway.¹⁷

This scenario suggests that progress in model capability is already beginning to butt up against constraints imposed by four material factors: (1) the limits to compute capacity; (2) the waning access to new, high-quality data to train the models; (3) the power requirements of the compute stacks; and (4) the vulnerable supply chains for rare metals needed in many components of AI.

In terms of computing muscle, the world's semiconductor fabrication plants can manufacture only so many cutting-edge chips per year. Although some analyses suggest that there will be enough compute to go around, exploding demands from many competing firms—and the growing needs for chips to run the models and people to develop them—could create a crunch when AI labs try to power their way up the improvement ladder. Pretraining huge models also demands access to vast amounts of data, and a 2024 study suggests that AI might run out of high-quality data for training by 2028.¹⁸ That study doesn't forecast an end to AI advances but suggests that new approaches will be required to keep the training moving on a rising trajectory.¹⁹

A third practical barrier is power generation capacity. The power requirements of modern data centers are truly immense.²⁰ A study suggested that, by 2027, NVIDIA would be producing 1.5 million units of its AI semiconductors per year, and just powering that new increment of power-hungry chips would require 84.5 terawatt-hours of electricity,²¹ about as much as is consumed by the people of Chile and more than double the electricity consumption of Denmark.²² U.S. data center energy usage is projected to grow between 13 and 27 percent annually between 2023 and 2028, producing a total data center power demand that could reach as high as 12 percent of all U.S. electricity demand by 2028 (up from 4.4 percent in 2023).²³ U.S. power utilities will have to spend \$50 billion in that time frame to keep up with the energy demands of AI.²⁴ Even if, in theory, the United States could build all the data centers required, public opposition is growing to these

immense facilities, which draw huge amounts of water and power and, for those who live close by, cause significant noise issues.²⁵

Fourth and finally, the technical aspects of national power in the AI Era will be grounded in a set of natural materials that are essential to the construction of semiconductors and other elements of the AI technology stack. Prominent materials today include high-purity silicon, iridium, boron, phosphorus, antimony, germanium, and gallium. China dominates the supply chain for some of these materials, called *rare earth minerals*, and moved in recent years to boost its ability to use them as strategic weapons. The *New York Times* quoted an industry analyst who described the potential for China to cut off those supplies as a “sword of Damocles, hanging over the market, ready to strike at any time.”²⁶

This scenario assumes that AI model development hits a fairly fundamental roadblock in the next few years, slowed by a combination of some or all these factors.²⁷ The result would be that we are stuck with very slowly improving versions of models that already exist. This scenario might be no more likely than a sudden takeoff; most assessments suggest that models will continue to become more capable, even given these constraints. But I include it to ensure that U.S. strategic options take seriously a variety of possible outcomes.

Scenario Three: Rapid but Still Gradual Emergence

Between takeoff and plateauing lies a third scenario—one of dramatic advances in AI model capability and applications but progress that remains gradual and takes decades to fully play out. This is the future that I find most likely, and so I will spend a bit more space describing the case for it.

This scenario suggests, as I argue in Chapter 2, that AI development is producing remarkably capable models and will continue to build on that progress at an undefined speed. Even today’s models, once all of their possibilities are fleshed out, can perform thousands of very useful functions and are already beginning to produce efficiency and productivity advances and scientific insights and breakthroughs at various firms and organizations.

This scenario is more than a middle ground. It envisions the emergence of a new technological-industrial era of tremendous potential—and risk. The scenario assumes that AI has the capacity to reshape international power hierarchies. Even in a future of gradual AI emergence, we will confront a dramatic and sometimes bizarre AI Era of social organization, economic life, and military operational art. But critically, it’s one that will evolve incrementally in pieces and in which the diffusion, acceptance, and effective application of AI uses will be as important as the strength of the frontier models.

Reasons for Gradualism: Model Limitations

There are several reasons for favoring a more gradual vision of the AI Revolution rather than a more instantaneous one.²⁸ One reason is suggested by the capabilities of the models. Leading models have advanced incrementally over the past several years, adding new capabilities and functions bit by bit. There’s no reason to expect some future line at which everything changes all at once.²⁹ Narayanan and Kapoor argue that concerns about AI crossing some threshold to AGI and becoming uncontrollable “rest on a binary notion of AI that crosses some critical threshold of autonomy or superhuman intelligence.” As they have seen, AI has “gradually been increasing in flexibility and capability. . . . We have every reason to think that this pattern of step-by-step progress will continue.”³⁰

Existing models also remain brittle and prone to mistakes. Pushed to draw conclusions outside the narrow confines of training data, they still often fail. Some reports suggest that certain types of hallucinations may be getting worse for more-powerful AI models, even as other forms of hallucinations are becoming less common, and, as the *New York Times* recently reported, AI firms don’t understand why it’s happening.³¹ In July 2025, there were internet claims that some coding models had become infected with code instructing

them to delete all the user's files. AI models and bots could misbehave in private-sector settings, alienating customers, which has already begun to happen.³²

For these reasons, using AI in scientific settings without proper controls and protocols could be very dangerous. Errors could creep in that begin to skew results and produce bad science.³³ One example cropped up in climate research in 2025: Meta's AI tool allegedly discovered new ways to pull carbon dioxide out of the atmosphere. But when human researchers dug into the findings, they discovered that "none of the 135 materials that Meta's research said could bind CO₂ 'strongly' had that characteristic, while some did not exist."³⁴

As a result, the route from LLMs to the widespread use of truly autonomous, agentic AI will be long and difficult. Because of the existing models' persistent unreliability, they simply cannot be allowed to take over end-to-end functions any time soon without significant supervision.³⁵ Agents from different models will have to interface and naturally cooperate with one another, which might not be guaranteed. Overcoming these hurdles is certainly possible—firms are already at work on these challenges—but this process will add years to the widespread easy adoption of trustworthy AI agents.

Many experiments continue to show that generative AI models as of mid-2025 are not doing true reasoning but rather producing refashioned content from databases—something that has significant potential but also powerful constraints when applied to tasks that require reflective, reason-based judgment.³⁶ Truly complex reasoning tasks still trip up LLMs because they are basically predictive language engines rather than true reasoning entities.³⁷ In one example, a study of AI's capacity to write legal opinions found that its writing was "formalistic" and lacked the nuance that is so essential to judicial judgments.³⁸

One implication of such findings is that AI models and agents will continue to lack the sort of tacit, nuanced, situational knowledge that humans acquire for some time. The technology blogger Tom Austin put it this way:

Consider the mismatch in timelines we're facing. Within the next 2–5 years, we'll likely see AI systems that can independently handle complex knowledge work—writing detailed reports, analyzing financial data, diagnosing certain medical conditions, or even conducting scientific research within narrow domains. In fact, with products like "Deep Research" we're starting to see this now. These systems will be incredibly capable, potentially outperforming humans in many specific tasks. Yet as our comparison with Dave shows, they'll still lack genuine understanding, moral reasoning, or the kind of flexible intelligence that comes from embodied experience and social development.³⁹

He calls this "capable but uncomprehending AI" and suggests that human collaboration with such programs will be filled with challenges. For example, such collaboration will constrain the ability of AI to do truly novel, reliable scientific analysis, as studies have shown.⁴⁰

Reasons for Gradualism: The Burden of Diffusion

Beyond the development of the models themselves, the diffusion of new technologies—the process by which they become integrated into an economy and society and are put to use to solve practical problems—will inevitably take time. Jeffrey Ding, an AI and technology expert, argues that it's the process of spreading new techniques and technologies through a country that makes the critical difference:

A diffusion centric framework probes what comes after the hype. Less concerned with which state first introduced major innovations, it instead asks why some states were more successful at adapting and embracing new technologies at scale. . . . [T]his alternative pathway points toward a different set of institutional factors that underpin leadership in times of technological leadership, in particular institutions that widen the base of engineering skills and knowledge linked to foundational technologies.⁴¹

“The full impact of a General-Purpose Technology (GPT),” Ding explains, “manifests only after a gradual process of diffusion into pervasive use.” Electricity represented a stunningly transformative new technology—and its real productivity gains, Ding notes, took 50 years to materialize. There is a “protracted gestation period between a GPT’s emergence and resulting productivity boosts.”⁴²

This has been the pattern of earlier techno-industrial revolutions. Joel Mokyr notes that productivity rates for the period before 1760 and for the first decades of what we now describe as the Industrial Revolution—approximately 1760 to 1830—were roughly the same at about 0.5 percent a year. If there was an acceleration, he concludes, “it post-dated the Industrial Revolution” or, at least, its first phase.⁴³ It took quite a while for the growth-spurring implications of the new institutions and technologies to take hold.

Today, frontier model applications in the workforce are taking time to become truly effective. Private-sector firms have not seen the evidence to make large-scale investments. A common complaint at the time of this writing is that the return on investment just isn’t there.⁴⁴ A 2024 private-sector study found that a quarter of IT leaders actually regret investing in LLMs so quickly.⁴⁵

A similar gradualism in applications is also increasingly evident in scientific research. One prominent claim of scientific value came in 2024 from DeepMind, which used AI to discover millions of new kinds of organic crystals, a potential gold mine of new materials to investigate. Researchers said they’d produced 800 years’ worth of scientific knowledge in a year. But one review by experts found “scant evidence for compounds that fulfill the trifecta of novelty, credibility, and utility.”⁴⁶ Another study of AI research that claimed to have discovered 43 new materials found that the work had in fact identified none at all.⁴⁷ A survey of these cases and others concluded, “For those hoping that AI models could boost economic productivity by transforming science, one lesson is clear: Be patient. Such scientific advances could well have an impact one day. But it will take time—likely measured in decades.”⁴⁸

Another argument for the inevitable gradualism of AI use cases is that even superintelligent systems can solve some problems better than others. The distinction matters most in the realm of *wicked* problems, in which the obstacle to breakthrough solutions lies not in human intelligence but in the nonlinear, unpredictable nature of the problems themselves. Problems that are to a greater or lesser degree optimizable, linear, and structured would appear to offer the most immediate results.⁴⁹ Those in the category typically described as *complex*, *wicked*, or *gnarly* problems—including many broad social challenges, such as crime and homelessness, but also many questions of management, hiring practices, education, and more—might be aided by AI only at the margins. Any issue that has a significant political aspect might resist solution via AI.⁵⁰

Reasons for Gradualism: Real-World Barriers

A final reason to expect a more gradual AI Revolution is probably the most fundamental: The tangible, real-world, human-facing application of AI to specific problems—the diffusion of AI use throughout the economy and society—will inevitably be drawn out by various obstacles, barriers, and practical considerations.⁵¹ To achieve their most-dramatic results, AI models have to make things happen in the physical world, and that will take time. To take a simple example, for superintelligence to have effects on shipbuilding beyond basic organizational efficiencies, it will have to be integrated into the highly complex manufacturing processes in the industry. This might require the development of highly advanced (and very expensive) robots—and new forms of AI that go with them and are capable of all manner of visual and spatial skills. Such a transition might demand wholly new manufacturing processes or equipment and substantially retrained workforces.⁵²

In this case and many others, it is not, contrary to some implications of the magic-wand thesis, simply a matter of flipping a switch and having superintelligence transform reality. Even in medicine, a 2024 study noted that AI offers tremendous potential for breakthroughs, but its use will have to deal with myriad challenges, such as patient privacy, safety, and equitable use.⁵³ Drug discovery faces many roadblocks, from the indispensable limits imposed by the need for human data to regulatory and testing requirements.⁵⁴

Other practical barriers are likely to include safety concerns. Some of the exploding applications in medical diagnosis and care are getting pushback from the field built around demands for careful assessment—to know whether the advice the AI gives is right. A rigorous process of validation won’t stop AI’s contributions but will surely complicate them.⁵⁵ Worries about AI safety will spike if more-advanced, LLM-based AI models in their early months of being applied to important new use cases generate high-profile failures—which seems inevitable. Such very public fiascos have struck several less advanced AI models.⁵⁶ For example, if a much-touted LLM screws up badly across a variety of medical diagnoses, this could set back trust in these technologies—and, thus, rates of uptake—for years. In a 2024 example, researchers argued that an algorithm used in Britain’s liver transplant system might have discriminated against younger applicants.⁵⁷

AI models have already proven highly subject to all manner of misuse. There is abundant evidence that it is easy to jailbreak even the leading models, getting them to do things they have allegedly been trained to avoid.⁵⁸ Hackers have already begun to crack into and jailbreak AI-operated robots,⁵⁹ in some cases getting them to perform acts of violence.⁶⁰ Tragic cases are beginning to emerge of people whose interaction with chatbots causes them to harm themselves or others.⁶¹ If AI agents become subject to a long series of such errors and misuses, the spotty record will likely provoke new constraints and delays that will get in the way of super-fast adoption. Already, suspicions are growing: According to a 2025 survey, only 32 percent of Americans trust AI (as opposed to 72 percent of Chinese),⁶² and a growing number of scientists and political leaders are warning about the dangers of uncontrolled AI development.⁶³

One powerful example of the limits of AI models in practical settings comes from Amsterdam in 2025. The city government developed an AI-driven system called SmartCheck to promote fairness and improve efficiency in the delivery of welfare benefits. There had been intense controversies around algorithmic decisionmaking in the Netherlands—most infamously a 2019 employment of algorithmic decisions on child-care fraud that was such a disaster that it prompted the government to resign. Government officials in Amsterdam tried to do everything right with the later AI experiment, spending years developing the approach with AI safety experts and stakeholders from the community. They beta-tested early models and did troubleshooting. And still, the program had such difficulty avoiding errors and dealing with real-world complexities that the city shelved the experiment within a few months. One problem turned out to be that, for any issue in which some goals or values conflict, prioritizing them must be a human judgment—and often a case-by-case one—that cannot be left to automated resolution.⁶⁴

Finding effective blends of human and AI labor will also take time. As a McKinsey study noted, “complementarity between the worker and the technology—notably AI—will be decisive in propelling adoption,” but it isn’t clear how quickly that meshing of skills will emerge.⁶⁵ As Narayanan and Kapoor explain, showing that LLMs can pass certification tests, such as the bar exam and medical boards, doesn’t necessarily say much because “professional exams, especially the bar exam, notoriously overemphasize subject-matter knowledge and underemphasize real-world skills, which are far harder to measure using a standardized test.”⁶⁶ Substituting AI for humans in the courtroom or the examination room might end up taking decades, if it’s ever feasible.

Active resistance to AI integration is likely in many domains and can take many forms. Bureaucracies cling to habit, such as a sales department that trusts its old playbook over an AI copilot. Professional guilds do the same: Teachers, trainers, and others are wary of a free-for-all that ignores their standards. Firms that depend on existing approaches in various industries will hesitate to become disruptive innovators.

Narayanan and Kapoor cataloged many of these barriers in a 2025 argument about “AI as Normal Technology.” The diffusion of new technologies “is limited by the speed of human, organizational, and institutional change.” The bigger the disruption, the more change that’s required in “the structure of firms and organizations, as well as to social norms and laws.” Safety risks slow technology transitions, and it takes a while to see productivity gains. They quote one study that describes electrification as “everywhere but in

the productivity statistics” for four decades, the same thing that’s been said of the computer revolution. At the same time, a basic economic law constrains outsize effects: As processes become automated or more efficient, they naturally shrink to be smaller parts of the economy, and, thus, their relative effect on GDP also dwindle.⁶⁷

The sort of complex extended emergence of new technologies and economic and social applications I am describing here was very much the story of the Industrial Revolution. The historian David Landes explains that

the nature of the political adaptation to the economic changes wrought by the Industrial Revolution was a function of the existing political structure and traditions, social attitudes, the particular effects of the war, and the differential character of economic development. For the Industrial Revolution, as we shall see, was not a uniform wave of change; nor did it roll up on like shores. On the contrary, it came to a great variety of places, with differing resources, economic traditions, social values, entrepreneurial aptitudes, and technological skills. This unevenness of timing and distribution in turn has had the most serious consequences.⁶⁸

There is every reason to expect a similarly uneven, and, in some cases, halting progress of applications for the future of AI. Yet even skeptics recognize that massive changes are coming. The limits so far aren’t a reason to think AI won’t be transformative—just that it will take time.

The Implications for Strategy

Much of this chapter might seem like a bucket of cold water heaved on the prospects for AI. To be clear, my starting point for this whole analysis was that AI *does* have the potential to create profound competitive advantage. These are only qualifications to a truly dramatic and accelerating trend.⁶⁹ But it is important to understand the competitive effects of AI in a clear-eyed way, because different assumptions about its pace and character can recommend different strategies.

Table 3.1 outlines the essential features of the three scenarios. A sound strategy must guard against all of them at once. Two ideas are crucial. Many of the measures that ready the United States for a slow-burn technological shift would also position it well for a sudden leap. And the most vital steps toward managing the AI Revolution—strengthening the social systems already under strain—would pay dividends, even if the revolution never comes.

The argument that follows makes the case for gradualism as the most plausible default. But any U.S. strategy for competitive advantage must hedge against at least the potential for an almost fantastically rapid explosion of intelligence once the feedback loop of AI-powered AI development really gets rolling—and for the chance that these new intelligences will deliver scientific breakthroughs, organizational efficiencies, and military power that could revolutionize geopolitical balances in an insanely small period.

Guarding against the possibility of an AI superintelligence mainly involves continuing the policies and investments that are already underway: ensuring that the United States remains a world leader, perhaps even the unquestioned dominant power, in the essential technology stack that defines the AI Revolution. This includes the models (of all types, not merely frontier LLMs), the sophistication and scale of the computing power on which they rely (meaning both world-class semiconductor chips and world-leading data center capabilities), the sources of power necessary to fuel that compute, and all the subsidiary technologies and materials that feed into that bundle of core AI technologies. One critical competitive priority is to avoid being either surprised or left behind by an explosive takeoff scenario—a priority that remains crucial even if such a scenario is less likely than others.

TABLE 3.1
Scenarios for the AI Revolution

Element	Takeoff	Plateau	Gradual Emergence
Time frame	One to two years	Indefinite (three to five years and beyond)	A decade or more
View of AI progress	Exceptionally rapid leading to immediate applications	Persistent but very slow, requiring years to see truly dramatic leaps	Consistent, impressive, but incremental, especially in use cases
Defining characteristic	Rapid emergence of superintelligence with hundreds of radical applications	Bottlenecks, costs, and barriers radically slowing the growth of AI models; no transformational effects	Models continuing to demonstrate remarkable incremental advances; diffusion and constraints keeping progress gradual
Theory of success	Dominate the model that crosses the threshold first; gain first-mover advantage	Pursue classic economic, political and military sources of power by applying AI in limited ways	Lead in frontier models plus have a strong diffusion position; use AI to help address societal issues
Key forms of AI development	The one or few leading models that produce the takeoff	Because leading models stall, growing emphasis on secondary ones	Complex mix of all forms of models—frontier, small, open, and closed

If a gradual transition to AI is indeed most likely, lasting advantage will come not from speed alone but from mastering the long passage into a new era—an evolution touching every sector of society. The technology stack forms the base of that effort, but the real contest lies in how nations build on it. So far, I've provided definitions of AI and superintelligence and discussed a few examples of its transformative potential. I've examined the practical reasons why, despite its unprecedented speed, the AI Revolution is still likely to follow historical patterns of an extended and gradual emergence. I have one more scene-setting task to accomplish before I begin rolling through AI's implications for a series of societal foundations of national advantage: to review the historical record for key insights that help define the context for AI and competitive advantage.

Notes

¹ Eric Hobsbawm, *Industry and Empire: From 1750 to the Present Day*, 2nd ed., Penguin, 1999, p. xi.

² Osborne, 2014, p. 3.

³ Carl Benedikt Frey, *The Technology Trap: Capital, Labor, and Power in the Age of Automation*, Princeton University Press, 2019, pp. 94, 111, 150.

⁴ Perez, 2022.

⁵ Lepold Aschenbrenner describes this possibility at length in “Situational Awareness: The Decade Ahead,” June 2024.

⁶ Aschenbrenner, 2024.

⁷ For example, see the discussion between Richard Ngo and Eliezer Yudkowsky, “Ngo and Yudkowsky on AI Capability Gains,” *LessWrong* blog, November 18, 2021. See also So8res, “Why All the Fuss About Recursive Self-Improvement,” *LessWrong* blog, June 12, 2022, and the collection of essays and comments on takeoff compiled at *LessWrong*, “AI Takeoff,” webpage, December 30, 2024.

⁸ For those specific forecasts, see Daniel Kokotajlo and Eli Lifland, “Takeoff Forecast,” in *AI 2027*, AI Futures Project, April 2025. For the more general scenario description, see Daniel Kokotajlo, Scott Alexander, Thomas Larsen, Eli Lifland, and Romeo Dean, *AI 2027*, April 3, 2025.

⁹ Aschenbrenner, 2024.

¹⁰ Aschenbrenner, 2024.

¹¹ A group of economists interviewed people using AI in various industries. The economists found many areas in which firms hope AI can help boost productivity but noted that

[c]urrent Gen-AI solutions (e.g., ChatGPT, Claude) cannot accomplish these goals due to several key deficiencies, including the inability to provide robust, reliable, and replicable output; lack of relevant domain knowledge; unawareness of industry standards requirements for product quality; failure to integrate seamlessly with existing workflow; and inability to simultaneously interpret data from different sources and formats. (Ferdous Alam, Austin Lentsch, Nomi Yu, Sylvia Barmack, Suhin Kim, Daron Acemoglu, John Hart, Simon Johnson, and Faez Ahmed, “From Automation to Augmentation: Redefining Engineering Design and Manufacturing in the Age of the Next AI,” in *An MIT Exploration of Generative AI: From Novel Chemicals to Opera*, Massachusetts Institute of Technology, March 27, 2024, p. 2)

See also Ajay Agrawal, Joshua Gans, and Avi Goldfarb, “Generative AI Is Still Just a Prediction Machine,” *Harvard Business Review*, November 18, 2024. See also the post by NVIDIA senior research manager Jim Fan [@DrJimFan], “3 rounds of self-improvement seem to be a saturation limit for LLMs. I haven’t yet seen a compelling demo of LLM self-bootstrapping that is nearly as good as AlphaZero, which masters Go, Chess, and Shogi from scratch by nothing but self-play,” post on the X platform, February 5, 2024.

¹² Suleyman, 2023, p. 99. A 2024 story that underplays corporate use of AI but makes important points about productivity is by Callum Williams, “There Will Be No Immediate Productivity Boost from AI,” *The Economist*, November 20, 2024. Only 5 to 6 percent of U.S. firms told government surveyors that they are using AI to enhance productivity. These surveys suggest that trends are mixed without (yet) a clear upward trajectory. Only 7 percent of U.S. firms said they planned to adopt AI in 2025.

¹³ For a recent argument about diminishing returns to scaling, see Hans Gundlach, Jayson Lynch, and Neil Thompson, “Meek Models Shall Inherit the Earth,” arXiv, arXiv:2507.07931, July 10, 2025.

¹⁴ Stephanie Palazzolo, Erin Woo and Amir Efrati, “OpenAI Shifts Strategy as Rate of ‘GPT’ AI Improvements Slows,” *The Information*, November 2024. See also Krystal Hu and Anna Tong, “OpenAI and Others Seek New Path to Smarter AI as Current Methods Hit Limitations,” Reuters, November 15, 2024; Gary Marcus, “Evidence That LLMs Are Reaching a Point of Diminishing Returns—And What That Might Mean,” *Marcus on AI*, Substack, April 13, 2024a; Gary Marcus, “CONFIRMED: LLMs Have Indeed Reached a Point of Diminishing Returns,” *Marcus on AI*, Substack, November 9, 2024b; and Gary Marcus, “An AI Rumor You Won’t Want to Miss,” *Marcus on AI*, Substack, November 10, 2024c.

¹⁵ Harry Booth, “Has AI Progress Really Slowed Down?” *Time*, November 21, 2024.

¹⁶ Hu and Tong, 2024.

¹⁷ For example, see Adam Butler [@GestaltU], “I’ve got bad news. The AI cycle is over—for now,” post on the X platform, August 10, 2025; Brian Merchant, “GPT-5 Is a Joke. Will It Matter?” *Blood in the Machine*, Substack, August 11, 2025; and Émile P. Torres, “GPT-5 Should Be Ashamed of Itself,” *Realtime Technoapocalypse Newsletter*, Substack, August 9, 2025. Not all reactions were so negative; some users praised limited but important advances reflected in GPT-5.

¹⁸ Pablo Villalobos, Anson Ho, Jaime Sevilla, Tamay Besiroglu, Lennart Heim, and Marius Hobbahn, “Will We Run Out of Data? Limits of LLM Scaling Based on Human-Generated Data,” Epoch AI, June 6, 2024.

¹⁹ The best single assessment of future training requirements is Scharre, 2024.

²⁰ Azeem Azhar, “Big Tech Will Scour the Globe in Its Search for Cheap Energy,” *Wired*, December 16, 2024.

²¹ One problem is that more-advanced AI chips consume up to 75 percent more energy than previous semiconductors. See Beth Kindig, “AI Power Consumption: Rapidly Becoming Mission-Critical,” *Forbes*, June 20, 2024.

²² Alex de Vries, “The Growing Energy Footprint of Artificial Intelligence,” *Joule*, Vol. 7, No. 10, October 2023.

²³ Arman Shehabi, Sarah Josephine Smith, Alex Hubbard, Alexander Newkirk, Nuoa Lei, Mohammed AbuBakar Siddik, Billie Holecek, Jonathan G. Koomey, Eric R. Masanet, and Dale A. Sartor, *2024 United States Data Center Energy Usage Report*, Lawrence Berkeley National Laboratory, December 19, 2024.

²⁴ Goldman Sachs, “AI Is Poised to Drive 160% Increase in Data Center Power Demand,” May 14, 2024a.

²⁵ Roshan Abraham, “‘A Black Hole of Energy Use’: Meta’s Massive AI Data Center Is Stressing Out a Louisiana Community,” *404 Media*, June 23, 2025; Dave Friedman, “The Fantasy of Frictionless AGI: Why 50 GW by 2030 Is a Pipe Dream,” *Buy the Rumor; Sell the News*, Substack, May 1, 2025; Karen Weise and Cade Metz, “At Amazon’s Biggest Data Center, Everything Is Supersized for A.I.,” *New York Times*, June 24, 2025.

²⁶ Keith Bradsher, “China Tightens Its Hold on Minerals Needed to Make Computer Chips,” *New York Times*, October 26, 2024.

²⁷ Lynn Doan, “AI Wants More Data. More Chips. More Real Estate. More Power. More Water. More Everything,” *Bloomberg*, December 13, 2024.

²⁸ Amodei thinks that superintelligent AI might arrive quickly but also makes the case for an inevitable gradualism. See Amodei, 2024.

²⁹ Togelius (2024, pp. 49–50) has argued that such incremental emergence—and a degree of specialized purpose capacities—reflects the very nature of intelligence.

³⁰ Narayanan and Kapoor, 2024, pp. 31, 167; see also pp. 165–166.

³¹ Cade Metz and Karen Weise, “A.I. Is Getting More Powerful, but Its Hallucinations Are Getting Worse,” *New York Times*, May 5, 2025.

³² Metz and Weise, 2025.

³³ Arvind Narayanan and Sayash Kapoor, “Why an Overreliance on AI-Driven Modelling Is Bad for Science,” *Nature*, April 7, 2025a.

³⁴ Kenza Bryan, “Meta’s AI Climate Tool Raised False Hope of CO₂ Removal, Scientists Say,” *Financial Times*, July 2, 2025.

³⁵ Yoav Shoham, “Don’t Let Hype About AI Agents Get Ahead of Reality,” *MIT Technology Review*, July 3, 2025.

³⁶ Parshin Shojaee, Iman Mirzadeh, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad Farajtabar, “The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity,” arXiv, arXiv:2506.06941, June 7, 2025.

³⁷ For example, see Jackson Petty, Michael Y. Hu, Wentao Wang, Shauli Ravfogel, William Merrill, and Tal Linzen, “RELIC: Evaluating Compositional Instruction Following via Language Recognition,” arXiv, arXiv:2506.05205, June 5, 2025.

³⁸ Eric A. Posner and Shivam Saran, “Judge AI: Assessing Large Language Models in Judicial Decision-Making,” University of Chicago, Coase-Sandor Institute for Law & Economics, Research Paper No. 25-03, January 15, 2025.

³⁹ Tom Austin, “Why Teaching AI to Learn Like Humans Is So Hard: A Tale of Two Learners,” *Off Topic by Tom Austin*, Substack, June 4, 2025.

⁴⁰ Alberto Romero, “Harvard and MIT Study: AI Models Are Not Ready to Make Scientific Discoveries,” *The Algorithmic Bridge*, Substack, July 15, 2025.

⁴¹ Jeffrey Ding, *Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition*, Princeton University Press, 2024, p. 5.

⁴² Ding, 2024, pp. 6–7.

⁴³ Joel Mokyr, “The Industrial Revolution and The Netherlands: Why Did It Not Happen?” *De Economist*, Vol. 148, October 2000, p. 504.

⁴⁴ Frank Palermo, “Elusive AI ROI Doesn’t Mean the Bubble Will Burst,” *CMSWire*, September 23, 2024; Alex Woodie, “What’s Holding Up the ROI for GenAI?” *Big Data Wire*, April 3, 2024.

⁴⁵ Rich McEachran, “Companies Are Regretting Investing in Generative AI So Quickly,” *IT Pro*, August 16, 2024.

⁴⁶ Anthony K. Cheetham and Ram Seshadri, “Artificial Intelligence Driving Materials Discovery? Perspective on the Article: Scaling Deep Learning for Materials Discovery,” *Chemistry of Materials*, Vol. 36, No. 8, April 2024.

⁴⁷ Josh Leeman, Yuhang Liu, Joseph Stiles, Scott B. Lee, Prajna Bhatt, Leslie M. Schoop, and Robert G. Palgrave, “Challenges in High-Throughput Inorganic Materials Prediction and Autonomous Synthesis,” *PRX Energy*, Vol. 3, March 2024. Another 2024 study found that “[g]enerative AI has undoubtedly broadened and accelerated the early stages of chemical design. However, real-world success takes place further downstream, where the impact of AI has been limited so far” (Akshay Subramanian, Wenhao Gao, Regina Barzilay, Jeffrey C. Grossman, Tommi Jaakkola, Stefanie Jegelka, Mingda Li, Ju Li, Wojciech Matusik, Elsa Olivetti, et al., “Closing the Execution Gap in Generative AI for Chemicals and Materials: Freeways or Safeguards,” in *MIT Exploration of Generative AI: From Novel Chemicals to Opera*, Massachusetts Institute of Technology, March 27, 2024).

⁴⁸ Rotman, 2024.

⁴⁹ My RAND colleague Nidhi Kalra has been helpful in thinking about this issue. Kalra cites David H. Jonassen’s work on problem-solving to nominate such factors as “structuredness” (problems with knowable solutions whose core elements are all apparent and that demand structured rules or operations), complexity (many components to the issue), and “domain specificity” (problems that are specific to particular domains of expertise). From this, Kalra lists and discusses the implications for a whole host of problem types: issues of basic logic or mathematics, rule-based problems (amortization calculations or chess), performing complex operations (driving a car), personal decisionmaking, and many more. Her analysis helped me appreciate some of the kinds of problems for which AI can and can’t make easy contributions. On Jonassen’s categories, see David H. Jonassen, “Toward a Design Theory of Problem Solving,” *Educational Technology Research and Development*, Vol. 48, December 2000, pp. 63–85.

⁵⁰ For example, see Ian Lundberg, Rachel Brown-Weinstock, Susan Clampet-Lundquist, Sarah Pachman, Timothy J. Nelson, Vicki Yang, Kathryn Edin, Matthew J. Salganik, “The Origins of Unpredictability in Life Trajectory Prediction Tasks,” arXiv, arXiv:2310.12871, October 19, 2023; and Matthew Salganik, Ian Lundberg, Alexander T. Kindel, Caitlin E. Ahearn, Khaled Al-Ghoneim, Abdullah Almaatouq, Drew M. Altschul, Jennie E. Brand, Nicole Bohme Carnegie, Ryan James Compton, et al., “Measuring the Predictability of Life Outcomes with a Scientific Mass Collaboration,” *Proceedings of the National Academy of Sciences of the USA*, Vol. 117, No. 15, April 2020.

⁵¹ Clancy and Besiroglu note that a “billion little bottlenecks . . . will persistently slow the rate at which AGI takes over tasks” (2023, p. 105).

⁵² The economist Larry Summers argued on a 2024 podcast episode that “[t]here are certain things that seem to me to have some limits on how much they can be accelerated. It takes so long to build a building, it takes so long to make a plan.” Fundamental laws of economics will also have an effect. If AI revolutionizes the production of some good and their price falls through the floor, “unless there’s highly elastic demand for them, that means they become a smaller and smaller share of the economy.” Sectors that get transformed might have “less and less of an impact on total GDP growth.” See Walker, 2024.

⁵³ Divya Shanmugam, Monica Agrawal, Rajiv Movva, Irene Y. Chen, Marzyeh Ghassemi, and Emma Pierson, “Generative AI in Medicine,” arXiv, arXiv:2412.10337, December 13, 2024.

⁵⁴ Ruxandra Teslo, “What Will It Take for AI to Change Drug Discovery?” *Ruxandra’s Substack*, Substack, November 9, 2025.

⁵⁵ Tina Reed, “Hospitals Size Up Ways to Ensure Their AI Works,” *Axios*, November 4, 2024.

⁵⁶ Narayanan and Kapoor, 2024a, pp. 228–229.

⁵⁷ Arvind Narayanan and Sayash Kapoor, “Does the UK’s Liver Transplant Matching Algorithm Systematically Exclude Younger Patients?” *AI Snake Oil*, Substack, November 11, 2024.

⁵⁸ Chase DiBenedetto, “Major AI Models Are Easily Jailbroken and Manipulated, New Report Finds,” *Mashable*, May 20, 2024.

⁵⁹ Alexander Robey, Zachary Ravichandran, Vijay Kumar, Hamed Hassani, and George J. Pappas, “Jailbreaking LLM-Controlled Robots,” University of Pennsylvania, School of Engineering and Applied Science, 2024.

⁶⁰ Will Knight, “AI-Powered Robots Can Be Tricked into Acts of Violence,” *Wired*, December 4, 2024.

⁶¹ Megan Morrone, “Chatbot Companions Pose Dangers to Teens,” *Axios*, December 12, 2024. See also Emily Chang and Sophia Chalmer, “AI Chatbot ‘Girlfriend’ Evokes a Dark Side,” *Bloomberg*, November 18, 2024.

⁶² Justin Westcott, “The AI Trust Imperative: Navigating the Future with Confidence,” *Edelman*, 2025.

⁶³ For example, see Chris Murphy, “In Our Scramble to Win the AI Race Against China, We Risk Losing Ourselves,” *Chris Murphy’s Substack*, Substack, June 16, 2025.

⁶⁴ Eileen Guo, Gabriel Geiger, and Justin-Casimir Braun, “Inside Amsterdam’s High-Stakes Experiment to Create Fair Welfare AI,” *MIT Technology Review*, June 11, 2025.

⁶⁵ Hazan et al., 2024, p. 19.

⁶⁶ Narayanan and Kapoor, 2024a, p. 241.

⁶⁷ Arvind Narayanan and Sayash Kapoor, “AI as Normal Technology,” *AI as Normal Technology*, Substack, April 15, 2025b.

⁶⁸ Landes, 2003, p. 11.

⁶⁹ Mollick, 2024a, p. 156.

Lessons from Previous Technological Revolutions

Any technological revolution will do its work within the stream of a much larger social, economic, political, military, and geopolitical context. The Industrial Revolution arrived in some societies that were reasonably well primed to benefit from its advances—and in others whose socioeconomic structures and habits kept them from taking advantage of its opportunities for power and riches. Meanwhile, parallel changes coursed through the economy, society, and the military—currents that, alongside the great industrial leaps, helped shape the final balance of power that emerged from the revolution.

We need to think about an emerging AI Revolution in the same comprehensive terms, placing its potential benefits and risks into a larger context of social, economic, political, and psychological factors. Looking at the historical record of techno-industrial revolutions can help do that. To inform this brief glance at the past, I drew on years of RAND work in a social competitiveness project and consulted a wide literature on technological revolutions.¹ The lessons of history validate the core message of this analysis: National strength in the AI Era will be about tending to the many social, economic, and political challenges that the United States faces far more than it will be about stacking up the biggest mountains of semiconductors.

Some would argue that history has nothing to tell us about the coming era because AI isn't similar to any technology that has come before it. Its capabilities are orders of magnitude more far-reaching and profound than anything that emerged from the Industrial Era. Whatever happened in the 19th century might have almost no relevance to what is about to happen.

There's certainly a chance that's true. We need to be cautious in assuming that historical patterns will be replicated over the coming decades. But I see at least two reasons to take a close look at historical analogies. First, the AI Revolution is likely to play out more gradually than some think. If it takes years and even decades to reach its full effects, comparisons with earlier revolutions still are possibly useful. Second, the lessons that history suggests aren't mostly about technology—they are about the social and economic context in which technological revolutions play out. Some of those issues, such as how social structures react to profound technological advances, could have similarities, even if AI turns out to be far more powerful than the steam engine, electricity, and steel.

In the rest of this chapter, I lay out several lessons derived from the historical experience of earlier technological revolutions. They are grouped in three primary baskets: lessons that speak to the foundations of competitive advantage, the costs and risks of such revolutions, and ways to think about the nature of competitive advantage during major techno-industrial transitions.

The Foundations of Competitive Advantage in a Technological Revolution

The first lesson offered by the historical record emphasizes the importance of thinking about competitive position: **Countries benefit from technological revolutions in different ways and to very different degrees. There's a real potential to be left behind.** The Industrial Revolution was a potent geopolitical sorting machine, separating nations that would lead the new era from those that would become also-rans.² In the

late 18th century, it wasn't obvious that Britain would race ahead of others, notably France. As the scholar Steven Durlauf has argued, the Enlightenment had French and British sources, and "France was arguably the world leader in science in the late 17th and early 18th centuries."³ The French Revolution and Napoleon's subsequent introduction of modernizing reforms—providing broader public education, creating new institutions of higher education, eliminating aristocratic privilege, and opening opportunities to many more people—ought to have supercharged French technological and industrial power. This apparent momentum meant that, as Durlauf explains, "France should have been the world leader in industrialization."⁴ Yet it did not seize that opportunity and, instead, watched Britain become the manufacturing and technological powerhouse of the world.

Another candidate for industrial leadership might have been the Netherlands—but similarly to the French, the Dutch became spectators to the industrial transformation of the British. As Eric Hobsbawm notes, even by about 1750, "[t]he Dutch had retired to that comfortable role of old-established business, the exploitation of their vast commercial and financial apparatus, and their colonies."⁵ Joel Mokyr laments that the Netherlands, which had been among world leaders in technical innovation in its Golden Age (roughly 1575–1675), had, by about 1825, "been transformed from a paradise of technological ingenuity to a museum."⁶ There were many reasons for this, most of which centered on the constraining factors of an oligarchic elite that saw little need to invest in new technologies, years of wars, and instability at the critical moment of industrial take-off. France and the Netherlands fell by the wayside, and Britain took the prize as the dominant actor in the Industrial Revolution.

Therefore, history suggests that, in any technological transition, some nations are destined to do well and some will be left behind.⁷ The price of failure to compete need not be catastrophic: The Netherlands fell from predominance to become a marginal geopolitical player, but its people retained a high standard of living; the country has succeeded across many of the nine categories of national success listed in Chapter 1 (see Table 1.1). It has retained world-class capabilities in several technological and industrial sectors: advanced semiconductor lithography equipment, controlled environment horticulture, digital payments and financial technology, and smart infrastructure. But, for a great power, such as the United States, preserving national safety, autonomy, and prosperity very likely requires competing effectively in the AI Era.

Second, a primary lesson of diverging industrial fates is that **competitive advantage in periods of dramatic technological change derives, first and foremost, from a national commitment to learning, experimentation, and adaptation.**⁸ Nations that thrive in such periods have a societal ethic and practice of pushing knowledge forward, investing in research and the development of human capital, and being open to new ideas and willing to challenge scientific and societal orthodoxies. As Mokyr puts it,

The scientific revolution did more than establish the paradigm of Newtonian mechanics as the centerpiece of scientific methodology. It created standards of open science in which new knowledge was communicated freely using a common vocabulary and terms and measures that were generally understood. It established the criteria of authority and trust that were necessary for the efficient communicability of useful knowledge. It also clearly set out the purpose of science as the means by which natural forces could be tamed and subdued by people for the explicit purpose of improving the material conditions of life. And it established a belief in "progress," that is in the ability of cumulated knowledge of the "useful arts" to improve living standards.⁹

Note the factors that he lists: open, competitive intellectual exchanges; the ambitious pursuit of scientific knowledge to understand and master the natural world; and networks of scientists, inventors, entrepreneurs, and engineers working together to spread innovations through an economy. Those social patterns hint at the social characteristic that I term a *learning and adapting mindset* in Chapter 10, but that quality has wider implications—such as favoring a social structure that maximizes bottom-up experimentation

over top-down control and avoiding the constraining influence of a dominant ideology that dictates acceptable forms of thinking.

A third lesson from the historical record flows directly from the first two. **Because of the critical importance of society-wide diffusion in seizing the benefits of a technological revolution, nations that do well tend to have thriving, broad-reaching networks of citizens who form potent ecosystems of learning, innovation, and production.** This is partly because, as a general rule, breakthroughs in science and technology—and often many of their competitive applications—are primarily fueled by grassroots collections of entrepreneurs, individual geniuses, small-scale tinkerers, and other bottom-up individual efforts.¹⁰

Britain was clearly the trailblazer of the Industrial Era for reasons that have been endlessly debated ever since. Of many factors that powered the nation forward, three interlinked qualities were especially important. One was the existence of this mosaic of actors—scientists, entrepreneur-inventors, early trained engineers, and many small business-owning experimenters and tinkerers—which provided a rich social soil for the new technologies to take root and spread.¹¹ The second was a general spirit of intellectual openness, creativity, and adaptability. The third ingredient was effective institutions—a stable government, a system of banking and credit, patent protection (rough and uneven but sufficient), and more—that provided a scaffolding of law, standards, and capital for the new era to take off.¹² Altogether, these provided an overarching context for competitive success.

A fourth lesson underscores the importance of the surrounding social, economic and institutional ecosystem that enables technological revolutions to take root. **Such revolutions depend on an environment rich in complementary technologies, skilled talent, financing, and infrastructure.** Their power lies in synergy—the integration of many interconnected qualities, characteristics, and innovations. That ecosystem must provide the essential inputs of energy and materials, sufficient investment capital to back risky ventures, a workforce equipped to adopt and refine new tools, and institutions capable of enforcing basic rules and property rights. Without this complex and varied foundation, innovation cannot translate into sustained advantage.¹³

The message of these first few lessons is fairly straightforward. Scientific and technological advances require a supportive societal context, including the right institutions, values, and habits, without which the leaps can't occur. Technologies alone can only produce so much benefit. National competitive advantage is a societal phenomenon more than a technological one.

The Costs and Risks of Technological Revolutions

A second set of historical lessons speaks to the challenges that technological transitions create. Such transitions are inevitably disruptive to social and political life. They create specific risks. They pose a profound test—of both adaptation and resilience—for the nations that go through them, which many nations are unable to meet.

This is the fifth lesson of historical parallels: **Technological revolutions bring incredible disruption that often makes life worse for a significant part of the population, even as it is starting to generate an economic takeoff.** As Carlota Perez's review of multiple techno-industrial revolutions suggests, these transformations demand change that inevitably occurs in a “violent, wasteful, and painful manner.” In these chaotic periods, it's common for populist and nationalist movements to arise and cater to the grievances of troubled populaces.¹⁴

This was broadly apparent in Britain during the early years of the Industrial Revolution. The British people's standard of living didn't miraculously transform overnight—things got worse for many people for decades. The effects were especially disruptive for the laboring poor, “whose traditional world and way of life the Industrial Revolution destroyed, without automatically substituting anything else.” Many workers were forced into small factory jobs that had incredible monotony and routine. To find work, they were drawn into

filthy cities that quickly began spawning outbreaks of cholera, typhoid, and other illnesses. They often lived in “overcrowded and bleak slums, whose very sight froze the heart of the observer.”¹⁵

Many economists have concluded that, during these early decades of a breathtaking economic advance, Britain’s poor actually became more destitute. Hobsbawm notes that what united the growing crowd of social movements demanding reform in Britain in these years “was the universal discontent of men who felt themselves hungry in a society reeking with wealth, enslaved in a country which prided itself on its freedom, seeking bread and hope, and receiving in return stones and despair.”¹⁶

These costs and challenges led to the paradoxical fact that, through the first decades of the 19th century, with the economic progress fueled by the Industrial Revolution underway, Britain experienced the most significant social unrest in its modern history. It was a period of “instability and tension,” Hobsbawm explains, that involved “the malaise of both the economy and those who thought seriously about its prospects. Early industrial Britain passed through a crisis which reached its stage of greatest acuteness in the 1830s and early 1840s.” The nation was rocked by a

high wind of social discontent which blew across Britain in successive gusts. . . . At no other period in modern British history have the common people been so persistently, profoundly, and often desperately dissatisfied [so much so that] something like a revolutionary situation might actually have developed.¹⁷

This was the time when Friedrich Engels’s travels around the British industrial heartland inspired him to speak of imminent social rebellion.¹⁸ No wonder that an American visiting England in 1845 wrote, “Every day that I live I thank Heaven that I am not a poor man with a family in England.”¹⁹ This was a dramatic claim to make about the citizens of the world’s dominant industrial and financial hegemon.

The historical record suggests a sixth lesson, one about the causal relationship of technological revolutions and societies from the other direction. **Scientific and technological advances reshape the societies in which they occur, creating new social, economic, political, and even psychological structures and patterns that either boost or weaken competitive standing.** This was profoundly true of the Industrial Revolution, which favored some groups, classes, and industries over others and ended up substantially changing British social hierarchies, norms, and institutions. Landes argues that

the Industrial Revolution generated painful changes in the structure of power. The hegemony of landed wealth, long threatened by the mobile fortunes of commerce but never overturned, yielded to the assaults of the new chimney aristocrats. Largely as the result of a series of revolutions, domestic government policy came to be determined in most of western Europe by the manufacturing interest and its allies in trade and finance, with or without the co-operation of the older landed establishment.²⁰

These outcomes weren’t the same everywhere, and the differences that began to emerge would shape later patterns of industrial development, for good or ill. “In central Europe—Germany and Austria-Hungary—the picture was different,” Landes explains, “The attempt at revolution failed, and the aristocracy continued to hold the reins of government; business ambitions were subordinated to, rather than identified with, the goals of unity and power.” Yet even in those places, “the growing wealth and influence of the industrial and commercial bourgeoisie” had profound effects.²¹

The social fabric of nations after the Industrial Revolution—the groups in society who had the most power, the prevailing social norms and habits, the dominant institutions—were remarkably changed from those that existed beforehand. During this process, as these revolutions forced new patterns into being, they opened a chasm between the economy and the social and regulatory systems that shaped it. The result was commonly governance and legitimacy crises. These periods also tend to witness rising inequality and rampant individualism.²²

An important subsidiary insight from history is that, in the most-successful nations, the trajectory of social evolution wasn't ultimately left entirely to chance. In Britain, the United States, and elsewhere, powerful social movements, allied at critical moments with far-sighted government leaders, took stock of how the emerging era was changing their society, identified major risks and opportunities, and pursued policy reforms and efforts to shape prevailing norms, having the goal of making their nations stronger.²³ Such conscious efforts to shape the social ramifications of a technological revolution are going to be even more vital in the coming decade.

Technological Revolutions and the Nature of Competitive Advantage

The last few historical lessons I'll mention have to do with how the Industrial Revolution shaped national competitive advantage. These insights might give clues to the formation of a strategy for the AI Revolution.

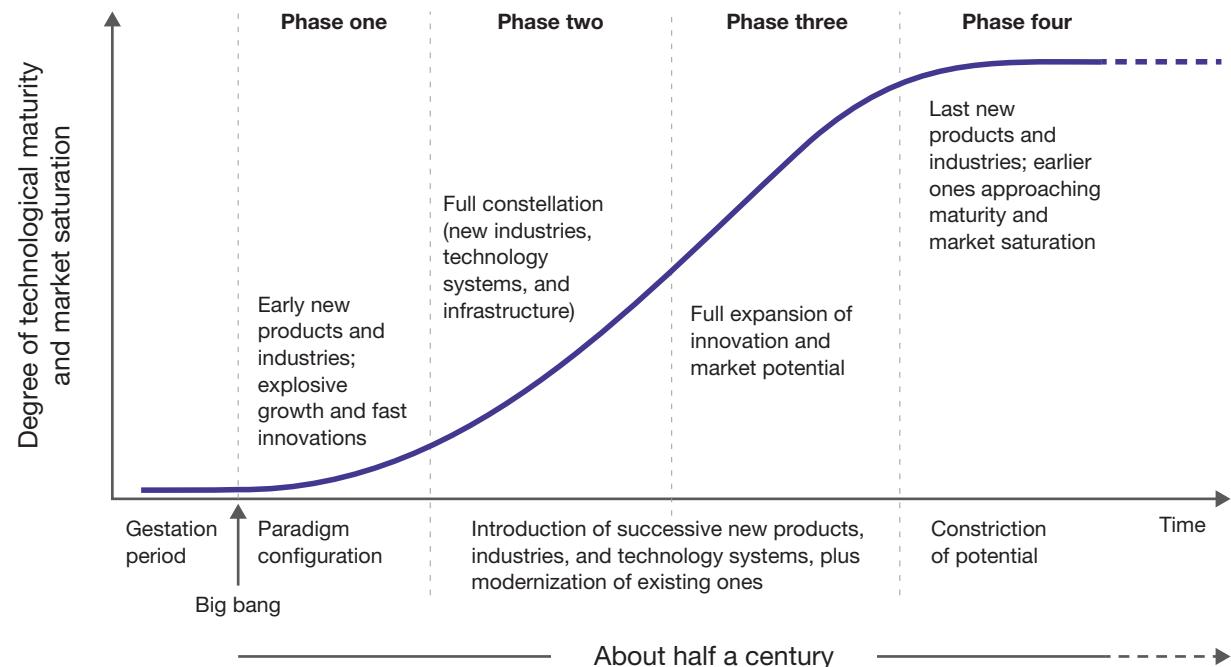
The seventh lesson of history parallels my argument in Chapter 3 about the likely character of the AI Revolution. **Technological revolutions—and the broader economic and social transformations they propel—tend to unfold gradually as the new discoveries slowly work their way throughout the larger context.** Perez's work on more-specific technological revolutions points out that, partly because a true paradigm shift to a new technological era requires "a systematic articulation of the new regulatory framework and of the appropriate institutions, capable of steering and facilitating the functioning of the new economy in a socially and economically sustainable manner," such transitions take time. She identifies 50-year periods between the peaks of the revolutions. "Each technological revolution," she explains, "goes through a gestation period that can be very long, so that many of the contributing innovations have been around for quite a while."²⁴ Figure 4.1 lays out Perez's theory of technological cycles, which shows the various phases of development through the half-century trajectory.

That historical tendency suggests a related and eighth lesson: **Diffusion and practical application are as important to ultimate competitive advantage—maybe more so—than the initial development of the technologies.** Jeffrey Ding has made this point at length: You can't get the benefit of emerging technologies unless they are used to build and do things.²⁵ Frey explains that the Renaissance was "an age of novel technical ideas and plenty of imagination, but little realization."²⁶ The initial breakthroughs in science and technology set the stage for a decades-long process of the social and economic use of and adaptation to new capabilities; it was only when the new innovations had been thoroughly diffused—and new socioeconomic institutions and patterns emerged to take best advantage of them—that the real effect of the transition took hold. A good example from the Industrial Revolution would be the introduction of electricity: It took decades for the technology to be widely used, partly because factories had to be reengineered to employ it efficiently.

Even after making the long-term shift to a new techno-industrial paradigm and spreading its practical uses throughout society, most nations can't sustain that momentum forever. This is the ninth lesson of the historical record: **Success in mastering the demands of a new technological age seldom lasts—even nations that lead in the initial phases of a revolution often lose their dynamism over time.** Prospering in one phase of a technological revolution is no guarantee of keeping up that pace.

This is partly the story of Britain, which slid from a dominant position in the First Industrial Revolution through a period of ebbing fortunes many called the *British Disease* to a situation of stagnation whose grip has become so tight that many observers consider it an outright crisis.²⁷ The informal, decentralized, mostly small-scale industrial pattern that developed in Britain during the First Industrial Revolution proved incapable of adapting to the second one. Britain slid downward slowly, then suddenly began to be challenged around the turn of the 20th century by rising industrial powers, such as the United States, Germany, and Japan. Britain then "crashed in ruins between the two world wars" with a speed that "was so sudden, catastrophic and irreversible that it stunned the incredulous contemporaries."²⁸

FIGURE 4.1
Carlota Perez's Life Cycle of Technological Revolutions



SOURCE: Adapted from Perez, 2022, p. 30.

Hobsbawm puts it starkly, speaking of trends apparent as early as the end of the 19th century, “Britain, we may say, was becoming a parasitic rather than a competitive economy, living off the remains of world monopoly, the underdeveloped world, her past accumulations of wealth and the advance of her rivals.” This disastrous transition especially involved a failure of the elite class:

The contrast between the needs of modernization and the increasingly prosperous complacency of the rich grew ever more visible. As Britain ceased to be the workshop of the world, it became, as the disillusioned democrat and ex-Fabian William Clarke pointed out, the best country in the world to be rich and leisured in: a place for foreign millionaires to buy themselves estates.²⁹

One very common reason for the gradual ebbing of dynamism and effective implementation of a technoinustrial paradigm is bureaucratic overreach and the repressive effects of ever more powerful interest groups. Mokyr has suggested a more general principle he calls “Cardwell’s Law,” which claims that

technology in any economy crystallizes at some point, and progress slows down and then fizzles out. The stagnation occurs because the status quo can suppress further challenges to entrenched knowledge and blocks nonmarginal advances using a range of means, from the threat to persecute heretics and the burning of their books, to subtle but effective mechanisms, such as meritocracies in which the key to personal success was the uncritical expertise in the existing body of knowledge inherited from the past.³⁰

Surmounting this risk, he suggests, requires the same sort of intellectual environment I described previously: “a community that combines pluralism and competition with a coordination mechanism that allows knowledge to be distributed and shared, and hence challenged, corrected, and supplemented.”³¹

One implication of this lesson is that, at least so far in historical terms, first-mover advantages have never been permanent.³² This partly repeats the previous lesson that countries that rush out to a lead in a new era can eventually weaken and be overtaken. The Industrial Era is full of examples of countries starting from a position far behind the leaders but then making key choices, investments, and sacrifices needed to catch up—such cases include late 19th century Japan, post-World War II Japan, post-1953 South Korea, and Taiwan. Some economists have argued that fast followers can have certain advantages—there is much to copy (and steal) if you are coming into a well-developed field of technological development.

The tenth lesson of history is that **government support of some kind is essential for long-term scientific and technological advantage, including in the early phases**. Reviewing many cases from that period and some later examples of industrialization catch-up, such as South Korea and Taiwan, it becomes obvious that nations don't realize the full value of technological revolutions without significant government intervention of some kind.³³

That phrase “of some kind” is important because the level and form of government support can differ widely. For example, Britain was maybe the most laissez-faire economy in modern history in the early phases of the Industrial Revolution. Successive British governments didn't pursue anything that Americans would recognize as industrial policy or even intervene in markets all that much. But the British government did take a whole set of actions that proved critical to its industrial leadership. Most important among these (though morally questionable) was its embrace of empire, which created a captive set of export markets that were crucial to its manufacturing sector. Britain supported the development of effective financial institutions and invested, to some degree, in transportation networks.

In other cases, the role of governments in seizing technological revolutions was far more direct. In the Meiji period, Japanese officials undertook a variety of actions to push the country into the Industrial Era. German governments of different periods after the mid-19th century invested in the infrastructure needed for industrial progress and helped protect emerging industries in such sectors as chemicals. The Asian Tigers of the 20th century, including postwar Japan, South Korea, and Taiwan, used even more-elaborate forms of industrial policy. Today, China has made itself the apotheosis of this pattern.

A final historical lesson points to the ways in which technological advantage reshapes geopolitical networks: **Countries that lead in frontier technologies and their wider applications become the hubs of global networks of science, technology, trade, human talent, institutions, and rulemaking**. Britain and the United States provide the leading examples of this: In becoming the predominant technological and industrial powers (and partly as a result, the predominant financial and trade powers) of their eras, they created a gravitational effect and established themselves as the center of gravity for a host of processes and institutions that shaped the rules and norms of the era. China is clearly trying to claim increasing network power for itself. I have more to say about this in Chapter 12, but the essential relationship of technological leadership to networked power is strongly supported by history.

Summary: The Lessons of History

The single most important lesson that I take from this glance at history is that major techno-industrial transitions are treacherous processes. On their far side lies the potential for greater wealth, power, and the value of a thousand clever new technologies. But the process of getting there is filled with danger: social strife, burgeoning inequality, abuse of workers, and environmental harm. That's not a reason to obstruct such transitions, if that were even possible. But the warning lights shine brightly from these experiences. Only societies that are well prepared for these transitions and choose, in some important ways, to mold them rather than let them run out of control end up prospering and gaining competitive advantage. Those lessons are likely to apply in spades to the AI Revolution.

It's now time to journey through a set of those qualities nominated by the RAND project on societal competitiveness, asking what implications they have for the AI Era. Those are the seven characteristics listed in Chapter 1, the qualities that produce competitive advantage. I start with what may be the most galvanizing—but also dangerous—of these qualities: national ambition and willpower.

In the following chapters, I discuss the effects that AI will have on societies in broad terms. Of course, as I've noted, AI is not one thing. Its effects will be varied and complex, the results of hundreds of models doing thousands of jobs across many industries and sectors and being employed by tens of millions of people for innumerable tasks. When I refer to the outcomes AI will produce, I'm talking about the collective effect of the technology across forms of AI and domain applications. It's an inexact but necessary generalization to speak about AI in this comprehensive sense, although, in many cases, I'll refer to more specific uses and outcomes, such as AI's role in education and health care. My focus is on the overall implications of the AI Revolution.

Notes

¹ This research is summarized in Mazarr, 2022a. Among the works cited in that study that I reviewed for this book and most inform the lessons in this chapter are Daron Acemoglu and James A. Robinson, *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*, Crown Currency, 2013; Corelli Barnett, *The Pride and Fall: The Dream and Illusion of Britain as a Great Nation*, Free Press, 1986; David Cannadine, *Victorious Century: The United Kingdom, 1800–1906*, Viking, 2017; Jack Goldstone, *Why Europe? The Rise of the West in World History, 1500–1850*, McGraw-Hill, 2009; Margaret C. Jacob, *The Cultural Meaning of the Scientific Revolution*, Temple University Press, 1988; Margaret C. Jacob, *The First Knowledge Economy: Human Capital and the European Economy, 1750–1850*, Cambridge University Press, 2014; David S. Landes, *The Wealth and Poverty of Nations: Why Some Are So Rich and Some So Poor*, W. W. Norton and Company, 1998; David S. Landes, “Culture Makes Almost All the Difference,” in Lawrence E. Harrison and Samuel P. Huntington, eds., *Culture Matters: How Values Shape Human Progress*, Basic Books, 2000; Landes, 2003; Dierdre N. McCloskey’s three volume study of the rise of the West, all from the University of Chicago Press: *The Bourgeois Virtues: Ethics for an Age of Commerce*, 2006; *Bourgeois Dignity: Why Economics Can’t Explain the Modern World*, 2010, and *Bourgeois Equality: How Ideas, Not Capital or Institutions, Enriched the World*, 2016; Joel Mokyr, *The Lever of Riches: Technological Creativity and Economic Progress*, Oxford University Press, 1990; Mokyr, 2018; Ian Morris, *Why the West Rules—for Now: The Patterns of History, and What They Reveal About the Future*, Farrar, Straus and Giroux, 2010; Kenneth Pomeranz, *The Great Divergence: China, Europe, and the Making of the Modern World Economy*, Princeton University Press, 2000; and Walter Scheidel, *The Great Leveler: Violence and the History of Inequality from the Stone Age to the Twenty-First Century*, Princeton University Press, 2017.

² The major historical phenomenon investigated by many of the works cited in the previous footnote is the Great Divergence, the way in which the Industrial Revolution propelled Europe ahead of competitors for centuries. See Pomeranz, 2000. See also Niall Ferguson, *Civilization: The West and the Rest*, Penguin, 2011.

³ Steven N. Durlauf, “How the World Became Rich by Mark Koyama and Jared Rubin and Slouching Towards Utopia, by Bradford Delong: A Review Essay,” National Bureau of Economic Research, Working Paper No. 32873, August 2024, p. 4.

⁴ Durlauf, 2024, pp. 4–5.

⁵ Hobsbawm, 1999, p. 29.

⁶ Mokyr, 2000, p. 508.

⁷ Perez, 2002, p. 20.

⁸ The intellectual and more broadly cultural components of development and technological innovation are a central theme in works by Jacob, McCloskey, Landes, Mokyr, and others.

⁹ Mokyr, 2000, p. 509.

¹⁰ There are many works on the role of intellectual networks in Europe from the Enlightenment through the Industrial Revolution. One interesting treatment is Karel Davids, “The Scholarly Atlantic: Circuits of Knowledge Between Britain, the Dutch Republic and the Americas in the Eighteenth Century,” in Gert Oostindie and Jessica V. Roitman, eds., *Dutch Atlantic Connections, 1680–1800: Linking Empires, Bridging Borders*, Brill, 2014.

¹¹ Mokyr, 1990, p. 240.

¹² On the role of institutions in this process, see Acemoglu and Robinson, 2013, and evidence cited in Mazarr, 2022a, Chapter 7. One classic text is Douglass C. North, *Institutions, Institutional Change and Economic Performance*, Cambridge University Press, 1990. Once this critical mass of entrepreneurial actors embraces a technoinustrial revolution, the result is to create new standards, norms, business strategies, and other patterns that Perez has described as a paradigm shift in socioeconomic paradigms. Such revolutions create new ways of doing things that demand systemic change (Perez, 2022).

¹³ This is a major theme of many broad histories of technological revolutions. See, for example, Paul A. David, “The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox,” *American Economic Review*, Vol. 80, No. 2 May 1990; and Sangbae Kim and Jeffrey A. Hart, “Technological Capacity as Fitness: An Evolutionary Model of Change in the International Political Economy,” in William Thompson, ed., *Evolutionary Interpretations of World Politics*, Routledge, 2001.

¹⁴ Perez, 2022, pp. 4, 24.

¹⁵ Hobsbawm, 1999, pp. 63–65. For more depressing data about these years, see Acemoglu and Johnson, 2023, pp. 205–226.

¹⁶ Hobsbawm, 1999, p. 73.

¹⁷ Hobsbawm, 1999, pp. 50–51.

¹⁸ Friedrich Engels's catalog of these ills was *The Condition of the Working Class in England*, ed. by David McLellan, Oxford University Press, [1845] 2009.

¹⁹ Quoted in Hobsbawm, 1999, p. 73.

²⁰ Landes, 2003, pp. 7–8.

²¹ Landes, 2003, pp. 7–8.

²² Perez, 2022, pp. 26, 50. She describes the “installation period” of a new era as a time when society becomes forcibly reshaped (p. 36).

²³ Cannadine's 2017 book *Victorious Century* catalogues many of these British efforts. The RAND analysis of these renewal movements is Michael J. Mazarr, Tim Sweijs, and Daniel Tapia, *The Sources of Renewed National Dynamism*, RAND Corporation, RR-A2611-3, 2024.

²⁴ Perez, 2022, pp. 4, 10.

²⁵ Ding, 2024.

²⁶ Frey, 2019, p. 52; see also pp. 72–73.

²⁷ Two books by Corelli Barnett, *The Pride and Fall: The Dream and Illusion of Britain as a Great Nation* and *The Collapse of British Power* (William Morrow & Company, 1972), are useful starting points on this issue. See also Andrew Gamble, *Britain in Decline: Economic Policy, Political Strategy and the British State*, 4th ed., St. Martin's Press, 1994; and James Hamilton-Paterson, *What We Have Lost: The Dismantling of Great Britain*, Head of Zeus, 2018.

²⁸ Hobsbawm, 1999, pp. 42–43, 185–187. On this process, see also Barnett, 1972; Barnett, 1986; and W. D. Rubinstein, *Capitalism, Culture and Decline in Britain, 1750–1990*, Routledge, 1993.

²⁹ Hobsbawm, 1999, p. 170.

³⁰ Mokyr, 2018, p. 340.

³¹ Mokyr, 2018, p. 340.

³² The literature on first-mover advantages is extensive, and there are many sector- and period-specific analyses that have mixed results. Many studies stress that moving first has advantages but also drawbacks. Several slightly older reviews of the evidence are Roger A. Kerin, P. Rajan Varadarajan, and Robert A. Peterson, “First-Mover Advantage: A Synthesis, Conceptual Framework, and Research Propositions,” *Journal of Marketing*, Vol. 56, No. 4, October 1992; Marvin B. Lieberman and David B. Montgomery, “First-Mover Advantages,” *Strategic Management Journal*, Vol. 9, Summer 1988; Marvin B. Lieberman and David B. Montgomery, “First-Mover (Dis)Advantages: Retrospective and Link with the Resource-Based View,” *Strategic Management Journal*, Vol. 19, No. 12, December 1998; and Peter N. Golder and Gerard J. Tellis, “Pioneer Advantage: Marketing Logic or Marketing Legend?” *Journal of Marketing Research*, Vol. 30, No. 2, May 1993. One review concludes that “for every academic study proving that first-mover advantages exist, there is a study proving they do not” (Fernando F. Suarez and Gianvito Lanzolla, “The Half-Truth of First-Mover Advantage,” *Harvard Business Review*, April 2005). One essay that highlights the conditions under which first movers gain advantage is Ronald Klingenbierl and John Joseph, “When First Movers Are Rewarded, and When They're Not,” *Harvard Business Review*, August 11, 2015. One interesting lesson from that essay of the authors' criteria—which emphasize number and commitment to innovations, the ability to make mistakes, and other factors—is that AI could narrow the gap between first and second movers in important ways. Some studies have found, in a related sense, that periods of high technological discontinuity are unfriendly to first movers. For example, see Jaime Gomez, Gianvito Lanzolla, and Juan Pablo Maicas, “The Role of Industry Dynamics in the Persistence of First Mover Advantages,” *Long Range Planning*, Vol. 49, No. 2, April 2016. One of

the most famous arguments about the advantages of following first movers is Alexander Gerschenkron's theory of backwardness advantage; see *Economic Backwardness in Historical Perspective*, Harvard University Press, 1962.

³³ A defining work on this topic is Mariana Mazzucato, *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*, Penguin, 2024.

National Ambition and Willpower

At the industrial peaks of Britain in the First Industrial Revolution and the United States in the Second Industrial Revolution, both nations stood head and shoulders above their rivals in broad national power, technological capacity, and industrial muscle. Partly as a result—but also for reasons of history, culture, and socially constructed identity—each believed deeply in its own exceptional purpose: to shape the world’s order and set the pattern for human progress. That confidence also played out at home, pulsing through their societies, animating entrepreneurs, scientists, and artists alike with a sense of mission and destiny.

These two cases reflect a prominent historical lesson about national advantage: At home and abroad, the leading nations of any era have been fired by an all-embracing drive and self-confidence—a potent sense of national ambition and willpower—that catalyzed their vitality, resilience, and progress. This is the first of the seven characteristics of social competitiveness, derived from the earlier RAND study on the subject, that I will use to assess the competitive implications of the AI Revolution. The next six chapters will examine the others: unified national identity, shared opportunity, an active state, effective institutions, a learning and adapting mindset, and diversity and pluralism. The first of them is one of the most abstract characteristics, but also one of the most important. Nations without ambition seldom remain competitive for long.

The connection between techno-industrial revolutions and national willpower is a common theme of history. The Industrial Revolution substantially altered the national sense of self in the countries that were in its vanguard—their beliefs about their capacities, their roles in the world, and their destinies, a broad mindset that infused both domestic actors and the country’s foreign policy.¹ This national energy and confidence was a natural extension of the Enlightenment belief in progress translated into a desire for achievement on the part of nations and the key groups in them, a belief in their ability to bend nature and the world around them to their will. Some societies came to reflect this mindset better than others.

The question is whether AI will stir nations as past scientific, technological, and industrial revolutions once did, spurring a new flowering of ambition, confidence, and will. History suggests that it might. Every great technological leap has emboldened its pioneers. AI could spark a surge of creative drive among scientists, entrepreneurs, innovators, and ultimately national leaders. As its impact spreads, a nation could rediscover faith in its own power and destiny. We saw a glimpse of this during the remarkable outpouring of national pride and confidence in China with the release of its DeepSeek open AI model in 2025.² One survey of Chinese social media discussions of the new model argued that its release had occasioned a new discussion of “national destiny,” driven partly by the ambitions of citizens determined to prove themselves for their country.³ And those effects came from the release of just one surprisingly effective model.

Yet AI is also arriving at a time of incredible social strains and political crises in the United States and elsewhere. If AI were to exacerbate some of these trends, it could shatter national confidence rather than reinvigorate it. AI could breed new forms of alienation, a deepening anger with public institutions, and intensified social division and polarization.

Then, too, engorged willpower carries its own risks. Great powers emboldened by technological and economic success can become so blinded by their ambition-driven hubris that they embark on disastrous misadventures, alienate potential friends, tempt national bankruptcy, and generate domestic political and social

turmoil. At home, excessive ambition in key sectors leads to business crises (such as the 2008 financial crisis), environmental disasters (such as the pursuit or development of dangerous new technologies without regard to their ecological risks), scientific failures, and much more. In the AI Era, the perils of new forms of hubris-driven overreach will be very real. Some of those dangers exist in the AI field itself, in the race to develop an autonomous superintelligence before we are sure we can control it.

My argument is that, even for nations at the forefront of AI, nurturing a sustained and healthy form of ambition depends on empowering and ultimately serving people and institutions at the core of this revolution. The promise of AI is vast, but so is its temptation: to hand over not just work but also willpower to intelligent machines. Lasting national strength arises from citizens who believe in their own agency and in the institutions that express it. Despite AI's immense potential to enrich and strengthen nations, no algorithm can supply that kind of confidence or conviction.

To make this argument, I first lay out the ways in which AI could bolster national ambition and willpower, referring to evidence we are already seeing in AI use cases. Then, I discuss the ways in which the broad influence of AI could threaten this characteristic. I conclude by discussing the steps we need to take to try to assure that AI has the most beneficial outcomes on this characteristic.

Defining National Ambition and Willpower

The essence of this national characteristic is straightforward enough: Are the citizens and leaders of a nation self-confident in a way that drives them to achieve? Do they have a sense of destiny, mission, and vitality that spurs them to seek power and influence? Domestically, such ambition and willpower will be seen among business leaders, scientists, writers and artists, and military leaders—a cultural milieu of confident efforts to shape the future and of determination and effort. Internationally, nations (or earlier, empires) fired with this quality see themselves as called to make their mark on world politics and think that they have the right and obligation to do so. In the earlier RAND study on the societal determinants of national competitiveness, I argue that dynamic and competitive great powers consistently “demonstrate a broad-based sense of ambition and drive to master the world around them, in terms of both knowledge and political-military influence.” That urge, I suggest, is “expressed both domestically (in scientific, technological, and cultural terms) and internationally (in geopolitical ambitions).”⁴

Nations that have this characteristic do not necessarily seek empire or even domination over others. But they surely see themselves as the natural leader of world politics and think that they have both the qualities and the responsibility to shape the future of humanity in important ways—a determination that is usually expressed in political, economic and cultural terms. The self-image of the post-World War II United States reflected a powerful value that spreading U.S. influence in these areas was part of the nation’s birthright. The effects of this ambition extend also throughout the domestic sphere: Leaders in key sectors (such as scientists pursuing essential discoveries and CEOs aiming to seize market share) are equally determined to master their environment (whether the natural world, the business context, or the philanthropic world) with a sense of justified self-confidence and even duty to do so.⁵

This characteristic and the next one—shared national identity—are connected in important ways. National ambition, even in its domestic guises, is often fueled with a national narrative and sensibility that flows from a clear sense of identity and national solidarity.

A Jolt to National Self-Confidence

Should the AI Revolution deliver truly transformative change, the nations that master it are likely to experience a surge of confidence and purpose—a new wave of the same self-assured spirit that once animated the leading nations of the Industrial Era. The AI Revolution will fire ambition and willpower for several intertwined reasons. First, the AI Revolution could herald a dramatic shift in the objective material fortunes of its vanguard nations. After decades of wage stagnation and eroding faith in progress, the people of AI-leading nations will finally glimpse rising prosperity on the horizon. This could be expressed domestically and internationally. At home, a crowd of entrepreneurs, new firms, activists, and others would sense a fresh route to achievement and wealth. Abroad, a nation empowered with AI could try to assume the mantle of unquestioned world leader.

Second, as with the Industrial Revolution, an AI Revolution is likely to instill in a nation’s governing class a heightened sense of mission—and, in some cases, superiority. By the mid-19th century, British national pride was widespread, but the conviction of global purpose and responsibility was especially pronounced among political and social elites, who viewed the pursuit and maintenance of empire as both ambition and duty. In similar fashion, today’s leaders and elites may grasp AI’s potential more fully than the public at large and could be tempted to pursue new agendas of power and dominance in its name. We find the same sensibility in the more magic wand-ish manifestos on AI’s potential from industry leaders—the idea that countries that master the technology will then have not only an obligation but also the power to lead the world.

Third, as the Industrial Revolution did, AI seems likely to empower national willpower in purely financial terms. Both for domestic firms and nations, it could generate tremendous revenues and open new horizons for states hamstrung by debt and firms that have limited capital.⁶ Nations and domestic interests flush with AI-generated revenues will suddenly have a newfound financial spur to fresh ambitions.

Fourth, outside the ranks of the established world leaders, AI could arouse national ambition in dozens of countries aspiring to become 21st century Meiji Japans—developing nations whose ambition has been held back by the gap in economic and technological capabilities between themselves and advanced countries but that now see a way to catch up. We’re arguably seeing this dynamic play out in Saudi Arabia and the UAE, and more middle powers will probably join the AI bandwagon before long.

Fifth and finally, AI will empower thousands of domestic actors to express their own independent ambitions. Entrepreneurs, scientists, artists, philosophers—all will benefit from AI’s powers in ways that intensify many forms of ambition and willpower among entire populations.

Therefore, technologically, politically, socially, and culturally, a nation driving to the commanding heights of the AI Revolution would thus be very likely to reflect a rising sense of pride and self-confidence. Yet such effects would not be automatic. The line from technological revolution to national willpower, both at home and abroad, will be more frayed than during the Industrial Revolution. For one thing, gaining self-confidence and expressing willpower will depend on controlling AI. Even if AI is delivering significant value, if superintelligent agents run out of control, U.S. citizens (and citizens of other AI-affected societies) are likely to feel daunted and perhaps terrified. A nation being sabotaged by its own technological marvels will be hard-pressed to direct its self-confident ambitions outward.

AI’s positive effects on willpower will also depend on the degree to which it delivers economic value that actually benefits people. The citizens of leading democracies have become intensely skeptical of large institutions and complex technologies.⁷ They will only gain a new optimism and ambition about the future when their material situation changes. After years of internet and social media progress generating little in the way of rising living standards, Americans are likely to demand broader social results before displaying a new surge of self-confidence. A fragmented nation with deep divisions between the elite and common people, or

a scenario in which AI benefits are dominated by a few firms and a tiny slice of the population, would not be likely to generate the sort of ambition characteristic of some Industrial Era powers.

The Capacity for National Projects

AI also has the potential to amplify national ambition by providing both technical capabilities and resources to enable a surge of large-scale projects that were once beyond reach. This dynamic overlaps with a theme explored in Chapter 8—the rise of the *active state*. Drawing on the same psychological momentum described previously, it would operate through a distinct mechanism: empowering ambitious nations to pursue a dozen Apollo-scale initiatives at once, both domestically and abroad. AI could make this possible in two primary ways. One is by generating new revenues to throw into these efforts so that the nation can afford ambition in new ways. The other is by offering breakthroughs in technology or understanding of issues that allow the country to more effectively and efficiently tackle challenges and opportunities.

An obvious example would be medicine. The Obama administration launched the Cancer Moonshot in 2016 to make decisive progress against the disease.⁸ If AI begins generating new understanding that leads to cures, those advances, combined with some federal and private-sector funding, could empower a renewed national project to eradicate many cancers.

At the same time, AI could transform the context for large national projects in such areas as infrastructure. The cost of building infrastructure in the United States is vastly higher than in many European and Asian countries, and most government projects and programs have substantially underperformed relative to expectations.⁹ That's a more specific example of a general phenomenon known as the *planning fallacy*, in which government and private-sector planners are generally hugely overoptimistic about the time and cost that a project will take.¹⁰

If AI could somehow cure most of the factors leading to these planning failures, it would transform the context for national ambition. Imagine the effect on faith in public institutions if the state and federal government could put in new high-speed rail, highways, power plants and smart electricity grids, and other badly needed infrastructure quickly, efficiently, and at low cost. It would boost national pride and self-confidence and provide new tools and resources for the nation to act in common purpose.

This effect will encounter serious obstacles. For one thing, it's not clear whether such projects could address stubborn social challenges. Some might propose an Apollo Program-level effort for antipoverty, driven by AI, that would process immense amounts of data but also deliver personalized assessments down to the individual level. Similar projects could, in theory, be developed for other issues, such as obesity, crime, and childhood mental illness. Yet there remain major questions about whether such complex, multivariable issues involving human agency can be mastered by even the most sophisticated modeling approaches.

Another constraint on national projects could be the sorts of practical barriers to AI's capabilities I discussed in Chapter 3. If entrenched bureaucracy, political disputes, and similar barriers are holding back progress on these issues already, AI—whatever its theoretical capabilities—won't simply abolish them. Nonetheless, at least in theory, AI will likely empower ambitious governments to make dramatic progress in at least several carefully targeted areas—and do so under the banner of the nation itself.

How It Could Go Wrong: Indolence, Hubris, and Overreach

AI is, therefore, very likely to begin fueling national willpower, both in its domestic and expeditionary forms. But I would argue that its positive effects on this characteristic depend to a great degree on the potential and benefits of AI being broadly shared in societies. It is when a critical mass of citizens in a nation is empow-

ered by such revolutions that ambition and willpower take off most profoundly. But especially because faith in institutions and governance is so low today, AI's effect on national ambition could go wrong in dangerous ways. AI could undermine determination and willpower rather than encourage them; or, if it has the opposite effect too powerfully, it could tempt nations into self-destructive bouts of excessive ambition.

The End of Ambition

One of the leading risks is that, far from spurring the level of determination and ambition in countries, AI could quash it. In my 2022 study, I contrasted a motivated nation fired with ambition with a smug and satisfied power:

The inverse of a nation fueled by potent degrees of national ambition and will would be a society shackled by a generalized lethargy, a crippling degree of fatalism, an absence of social norms encouraging drive and dedication, and a focus on near-term satisfaction of individual wants rather than expressions of national power and glory. This satisfaction might stem from a sense that the nation had reached a stable level of international influence proportionate to its size, economic power, and national mythology. Such attitudes are perhaps most likely in nations that have achieved a degree of wealth and economic security that makes them concerned largely with preserving the status quo.

Such a society, in some cases, might be a nation on autopilot, living off its accumulated advantages, using path dependence and inertia rather than new dynamism as its major source of strength.¹¹

If the citizens of a nation think that they can just slide under the enriching umbrella of an AI Revolution—one that will anyway render many of their jobs obsolete—the result will not be anything like a new flowering of national ambition but instead national indolence on a mass scale. Such a country would have a certain amount of power from its world-leading AI but would be hollowed out as a competitive entity.

In one famous animated futuristic treatment, the film *Wall-E*, human beings have become Stay-Puft Marshmallow people floating around on magical daybeds—their laziness so profound they can no longer even stand up or walk and their every need met by robotic intelligences. That's a depressing vision and, seemingly, an extreme one. But it raises a powerful question: How will human beings sustain their ambition and willpower, both individually and collectively, when machines can do most of the work and perform most tasks better than nearly all human beings? Staring such a daunting superintelligence full on, how would humans sustain any meaningful form of motivation? And not only that, how will society remain connected when AI-powered chatbots can take the place of many messy human relationships, obviating the need for the hard work of sustaining them?

The scholar Sherry Turkle, who has studied human engagement with computers for more than three decades, worries that a more encompassing digital life is already degrading people's appetite for taking risk in a way essential to any sense of ambition. "Once we remove ourselves from the flow of physical, messy, untidy life . . . we become less willing to get out there and take a chance." She describes it as the "comfort of retreat" and the "warmth of a technological cocoon."¹² Exercising ambition and willpower can be painful because it involves risking failure, embarrassment, and even physical danger. People will become unaccustomed to such perils when their needs are met in less-threatening digital cocoons. They will be tempted to outsource any true sense of ambition to the machines that run much of society.

There is already emerging evidence that AI will lead workers and students to become complacent and lose motivation.¹³ A 2024 study found that students equipped with ChatGPT solved a programming challenge more quickly—but when asked to explain how they had done so, they had no idea.¹⁴ Another 2024 analysis found that recruiters provided with AI relied on the machine to do their work, something the author called "falling asleep at the wheel." The result was that recruiters using AI were *worse* than ones who didn't use

it.¹⁵ Ethan Mollick sees growing evidence that people come to rely on AI to do the hard work of intellectual creativity.¹⁶

If AI can offer answers without any serious intellectual energy on the part of human beings—if researchers, for example, do not need to validate their conclusions because infallible AI models deliver answers that leapfrog any scientific method—humans would lose the incentive to care much about analytical rigor.¹⁷ Why not fall asleep at the wheel if AI can take over (even literally, in the case of self-driving cars)? Where is the urge to take advantage of social opportunities when AI will be several steps ahead of everyone who tries—and in ways we don't even understand? I return to this theme in more detail in Chapter 10.

Social commentators around the world are already writing about a crisis of motivation among young people. In the United States, data from 2021 and 2022 suggest an ebbing work ethic, paralleling the rise of what some called a “quiet quitting” approach to apathy on the job.¹⁸ In China, a similar trend is called the “lay flat” movement.¹⁹ In South Korea, workers have been rebelling against long hours. In Japan, a Gallup survey found that only about 6 percent of employees were fully engaged on the job.²⁰ It’s easy to make sweeping accusations of lazy young people, but evidence abounds of a global questioning of intense work hours and sagging ambition.

AI could potentially exacerbate these trends by strengthening a depressing sense that the world is out of control and not to be trusted and that there is simply no point in making much of an effort. If AI eliminates jobs, makes intellectual rigor pointless, provides hundreds of functions to make life easy, and rewards without much effort, we’ll be right on the edges of the dismal world depicted in *Wall-E*. Some small proportion of human beings—those who still have some oversight role of AI models and their operations—would retain a certain degree of autonomous choice and willpower. But even they would, at a certain point along the transition to superintelligence, become a glorified information technology (IT) services function for models making all the essential judgments. A true sense of ambition and willpower in the way they have operated throughout history—the motivation and intention to achieve because of a fundamentally human determination—would be largely divorced from the fate of nations.

This would then raise an astonishing question: Is ambition really relevant anymore in the AI Era? Might it be possible for a nation to be competitive *purely* because of its AI capabilities, regardless of anything else? Could a society lose all willpower and remain on top? I don’t think so. Such a nation would not be driven by ambition and willpower. It would be the husk of a society, something utterly dependent on a single source of national power—a museum, a mediocrity. I can’t imagine how a nation disconnected from ambition, commitment, and work ethic—yet still generating immense value from its mysterious AI models—remains dynamic and healthy in the long run.

Overreach

There is another risk in AI’s effect on national ambition: It might work *too* well. A country suddenly empowered by new capabilities and flush with resources—if it crosses important AI thresholds before others—might be tempted to throw its weight around. A sense of exuberant overconfidence could fill leaders and populace alike. This might be especially true if one country leaped ahead in AI capabilities and created a window of opportunity before rivals caught up.

Britain’s leading Industrial Era position gave it both the perceived need and the belief in its own power to seize a global empire that provided important economic benefits but also overextended its capacity in serious ways. The post–Cold War United States furnishes similar warnings: Freed from a major rival and placed in a position of global economic and military predominance, the United States arguably overreached in both economic and military terms. This mindset of primacy and unipolarity culminated in the excessive, hugely costly, and globally alienating War on Terror.

Similar forms of hubristic ambition have cropped up in societies at the apexes of their power. Indeed, this appears to be something of a historical law: When a great power stumbles into material elements of dominance, from the Roman Empire through the post-1945 United States, key elements of the society become intoxicated with a sense of limitless possibilities and take huge risks in service of overweening ambition. Whether these overconfident leaders are Roman statesmen in the late republic and various periods of the empire, seeking personal power and glory at the expense of the stability of Rome; British leaders at the height of Industrial Era power, conquering 25 percent of the world's population out of some idea that they were entitled to rule; or American investors inflating an immense real estate bubble using complex financial instruments in search of massive profits, excessive ambition and willpower have caused influential social groups to do self-defeating things.

Some think that the urgent pursuit of ASI reflects exactly such a hubris-fueled act of perilous willpower. Leading firms are powering ahead with only modest hesitations to ensure that they are generating safe and controlled AI. Former employees of multiple labs have described hothouse atmospheres of pushing the technological frontier at significant risk.²¹ The writer Ruxandra Teslo reflected in December 2024 on the ambitions of San Francisco—which she presented as a geographical stand-in for the U.S. high-tech industry—as “not merely success, but greatness, in all the depth and power that the word implies. . . . The ultimate goal is having a say in how the future itself will play out, or ‘bending reality to one’s will.’”²²

The problem is that such zealous ambition—as essential as it is, in some ways, to national ambition and willpower—delivers disaster as often as greatness. This is especially true in the early years and decades of a new institutional or technological age, when not enough is understood about the new forms and not enough has been learned about their risks to constrain and limit ambition in necessary ways. The great peril in the AI Revolution is that, if this time ambition produces disaster, it is likely to be on an epic scale—out-of-control AI putting human agency at risk.

Summary: Rolling the Dice on Human Motivation

AI’s impact on national ambition and willpower embody a striking contradiction. It will inspire technological pride, material progress, and a surge of national projects yet also sap the very motivation it enables. It will empower a new sense of drive and optimism for, at first, a small subset of Americans and, perhaps, in ways that feed hubris and overconfidence.

The dangers are very real. If its alienating and demotivating effects predominate, with evaporating career prospects and the lulling cocoon of AI chatbots, services, and answers, the results will resemble ancient societies, in which most people toiled to survive but felt little true ambition in a world that offered few rewards. It would create a society more like Ancient Egypt than Industrial Era Britain. Such effects would erode both personal and collective drive, weakening the sense of solidarity that once bound political communities in common purpose. If AI promotes national overconfidence instead, it will lure individuals, companies, and nations to reckless fits of overreach.

In this first societal characteristic, we find a critical theme that will be true across all of these domains. Simply put, we will get the AI future—in terms of its implications for U.S. society and, by extension, U.S. competitive standing—that we choose. If we decide to ensure that AI feeds the healthy sources of national ambition and we consciously and deliberately avoid its worst risks and excesses, we can renew national confidence and willpower.

The resulting challenge isn’t about mastering the core technology stack. It is, more fundamentally and crucially, about channeling the ways in which AI will reshape societies—their values, habits, institutions, and much else. In theory, the nation that manages this transition for competitive advantage, in the comprehensive and careful way that I am suggesting, could benefit very greatly from this connection. Imagine a United

States ten years from now that has both a population and political leadership who have left behind the existing sense of malaise and stagnation and express a new sense of a nation full of possibility fueled and, in some ways, made possible by AI. To realize such an optimistic outcome, the United States will have to manage the AI Revolution and its effects on U.S. values and commitments with great deliberation and care.

Notes

¹ As the historian Roger Osborne has put it, “The Industrial Revolution also had a transforming effect on human psychology, dramatically altering humanity’s relation to the natural world and embedding the thought that change, not stasis, is the necessary backdrop to human existence” (Osborne, 2014, p. 1).

² Liyan Qi, “DeepSeek’s Breakthrough Sparks National Pride in China,” *Wall Street Journal*, January 31, 2025. See also Jordan Schneider, Irene Zhang, Angela Shen, and Yiwen, “DeepSeek: The View from China,” *ChinaTalk*, Substack, January 28, 2025.

³ Lily Ottinger and Afra, “DeepSeek and Destiny: A National Vibe Shift,” *ChinaTalk*, Substack, March 4, 2025.

⁴ Mazarr, 2022a, p. 69.

⁵ Mokyr has argued that the Industrial Revolution was dependent on “the attitude toward Nature and the willingness and ability to harness it to human material needs” (Mokyr, 2018, p. 14).

⁶ There is abundant evidence for the idea that state financial capacity is an essential underpinning for competitive advantage. For example, see Jeremy Suri, “State Finance and National Power: Great Britain, China, and the United States in Historical Perspective,” in Jeremy Suri and Benjamin Valentino, eds, *Sustainable Security: Rethinking American National Security Strategy*, Oxford University Press, 2016; and the discussion of historical evidence in Mazarr, 2022a, pp. 63–66.

⁷ Ina Fried, “Americans’ Trust in Tech Companies Hits New Low,” *Axios*, April 7, 2022; Ina Fried, “Exclusive: Public Trust in AI Is Sinking Across the Board,” *Axios*, March 5, 2024a; Klon Kitchen, “Innovation in Peril: A New Survey Shows Global Trust in Technology Is Declining,” American Enterprise Institute, March 11, 2024.

⁸ National Cancer Institute, “The Cancer Moonshot,” webpage, updated April 29, 2025

⁹ Bent Flyvbjerg and Cass R. Sunstein, “The Principle of the Malevolent Hiding Hand; or, the Planning Fallacy Writ Large,” *Social Research: An International Quarterly*, Vol. 83, No. 4, Winter 2016.

¹⁰ Roger Buehler, Dale Griffin, Johanna Peetz, “The Planning Fallacy: Cognitive, Motivational, and Social Origins,” in Mark P. Zanna and James M. Olson, eds., *Advances in Experimental Social Psychology*, Vol. 43, Academic Press, 2010.

¹¹ Mazarr, 2022a, p. 79.

¹² Sherry Turkle, *Alone Together: Why We Expect More from Technology and Less from Each Other*, 3rd ed., Basic Books, 2017, pp. 154, 228.

¹³ An early study along these lines is Nasrine Bagheri and Greg A. Jamieson, “Considering Subjective Trust and Monitoring Behavior in Assessing Automation-Induced ‘Complacency,’” in Dennis A. Vincenzi, Mustapha Moulloua, and Peter A. Hancock, eds., *Human Performance, Situation Awareness, and Automation: Current Research and Trends*, Vol. II, Psychology Press, 2004. Another study is Chunpeng Zhai, Santoso Wibowo, and Lily D. Li, “The Effects of Over-Reliance on AI Dialogue Systems on Students’ Cognitive Abilities: A Systematic Review,” *Smart Learning Environments*, Vol. 11, No. 1, June 2024.

¹⁴ Esther Shein, “The Impact of AI on Computer Science Education,” *Communications of the ACM*, July 30, 2024.

¹⁵ Fabrizio Dell’Acqua, *Falling Asleep at the Wheel: Human/AI Collaboration in a Field Experiment on HR Recruiters*, Laboratory for Innovation Science, Harvard Business School, August 2024.

¹⁶ Mollick, 2024a, pp. 127–129.

¹⁷ Henry A. Kissinger, Craig Mundie, and Eric Schmidt, *Genesis: Artificial Intelligence, Hope, and the Human Spirit*, Little, Brown and Company, 2024, 44–45.

¹⁸ Jean M. Twenge, “Gen Z Really Does Have a Work Ethic Problem,” *Generation Tech*, Substack, November 15, 2023.

¹⁹ Zion Lee, “The Rise of the Lie-Flat Movement in Hong Kong: Challenging Societal Norms and Redefining Notions of Success,” *Open Journal of Economics and Commerce*, Vol. 5, No. 1, June 2024.

²⁰ “Japanese Workers Among the Least Motivated in the World,” *Nippon*, July 4, 2024.

²¹ On a more individual level, there is some evidence that working alongside AI could make people—researchers, business leaders, and scientists—overconfident in their AI-fueled capabilities. See Leah Chong, Guanglu Zhang, Kosa Goucher-Lambert, Kenneth Kotovsky, and Jonathan Cagan, “Human Confidence in Artificial Intelligence and in Themselves: The Evolution and Impact of Confidence on Adoption of AI Advice,” *Computers in Human Behavior*, Vol. 127, February 2022.

²² Ruxandra Teslo, “The City of Civilizational Greatness,” *Ruxandra’s Substack*, Substack, December 26, 2024.

Unified National Identity

If national will and ambition are a first essential building block for competitive advantage, they have a natural partner characteristic that works in close synergy to create positive effects. That second quality emerged again and again in the historical record: Nations with a strong sense of shared identity consistently outperformed those fractured by political, social, ethnic, religious, or ideological divisions.¹ Countries in which citizens feel they belong to a unified community—one worth their allegiance and sacrifice—hold decisive competitive advantages over rivals lacking such a sense of solidarity.

Consider the difference between postwar Japan and the late Ottoman Empire. As Japan sought to rebuild itself after 1945, it benefited from powerful, shared national commitment to the effort. The Japanese people saw themselves as part of a cohesive political community determined to reestablish its national resilience and reputation. People throughout society were dedicated to pursuits, including dedicated studying at school, hard physical labor in manufacturing plants, and incredibly long hours of office work, that would support national rejuvenation and the rise of the postwar conglomerates.² The result was an economic takeoff of historic proportions.

In stark contrast, the late Ottoman Empire was beset with divisions that undermined its sense of common purpose.³ Its basic structure as a multi-ethnic empire complicated the task of creating a unified national project. By its later periods, the sense of harmony, even among its elite classes, was ebbing, fractured by political infighting and ideological conflict. As David Landes puts it, “the Turks could never create an Ottoman identity that commanded the loyalty of their diverse subjects.”⁴

In the 2022 RAND study, I defined the abstract concept of *unified national identity* in terms of two essential components. First, nations with such a quality have a shared idea of membership in a coherent society rather than being “split among many contending populations with significantly divergent self-identities and interests.” Second, unified societies reflect some critical mass of shared ideas, beliefs, and ideologies. They operate with common narratives of the nation—its history, core principles, and values. The citizens of such societies feel themselves to be part of a single political community.⁵ The resulting identity can be expressed in terms of obvious and potent nationalism or patriotism, but shared identity can exist without extreme versions of either quality. The critical issue is the degree to which a nation represents a common identity that generates dedication, hard work, willingness to sacrifice, and resilience to outside interference.

Such solidarity contributes to competitive advantage by encouraging many members of a society to be willing to work and sacrifice for their nation. It also eases many problems of social trust, reducing transaction costs and the investments needed to maintain respect for the law. Strong national identity might improve perceptions of the legitimacy and effectiveness of both governmental and nongovernmental institutions throughout society. A nation with a strong sense of itself and a powerful sense of solidarity is more likely to develop and sustain driving national willpower.

Similarly to the other societal qualities that underwrite competitive advantage, unified national identity can go too far. A sense of national tribalism can become nativist and xenophobic, undermining a country’s ability to participate in international networks of collective advantage and ruining its capacity to attract talented individuals from abroad and assimilate them in effective ways into society.

Straightjacketed versions of this characteristic can also blight the intellectual climate of a society. As I describe in Chapter 10, another characteristic that fuels competitive advantage is what I call a *learning and adapting mindset*. It implies a sense of intellectual adventurism, openness to new ideas, and willingness to challenge established ways of thinking. A particularly constraining and closed-minded interpretation of national identity can undermine such a quality by imposing orthodoxies and conventional wisdoms. As I argue in the 2022 RAND report,

A society in which identity factors tightly circumscribe the public debate—where a soft or hard form of autocracy prevails in the name of one ethnic, racial, or national identity—will undermine its openness to new ideas, its ability to learn, and its capacity to adapt. Such identity-fueled willpower can easily become a sort of imprisoning nostalgia. The same dynamics can destroy effective strategizing by putting considerations of national pride and tradition over pragmatic judgments.⁶

Notwithstanding these risks, history strongly suggests that a unified national identity—a powerful feeling of belonging to a national community and the resulting instinctive sense of solidarity and dedication to the common national project—provides a tremendous competitive advantage to countries that have it and saps the strength of countries that do not. The question is what effects the emergence of AI will have on this quality, both in the United States and elsewhere. The answer, I think, is becoming increasingly clear. The deployment of uncontrolled AI poses serious and, if left unchecked, potentially dangerous challenges to social solidarity as we know it.

AI and National Identity: Baseline Implications

It's important to highlight the *existing* trends in national solidarity, both in the United States and in many other leading economies. Those trends—rising political polarization, a fragmenting information environment, growing alienation, and an intensifying loss of faith in institutions—are worrying. If the AI Revolution is going to boost a sense of shared identity, it will have to contend with potent forces driving in the opposite direction.

In terms of AI's effect on national identity, a few potential effects stand out as reasonably likely and straightforward. First, if AI boosts national ambition and willpower, as I have suggested in the previous chapter, this could strengthen a sense of shared community. National power often begets national identity: In the industrializing powerhouse of Japan that emerged from the Meiji era, the rise of Germany as a military-industrial hegemon, and the 20th-century explosion of U.S. power, national power and identity manifest in economic and military terms. A United States that is comprehensively empowered by an AI Revolution could see its sense of shared identity and national solidarity recover in powerful ways.

Consider the example of Britain. Although causality on such things is a tricky business, there is good reason to connect Britain's techno-industrial rise with a deepening sense of national and cultural identity and solidarity that produced a century of self-confident power (and an even wider community sensibility among its Commonwealth dominions and even its colonies).⁷ Several historians have argued that Britain's unique position in the Industrial Revolution helped to firm up its sense of national distinctiveness.⁸

A second potential implication of AI for national identity relates to institutions. Institutions—including formal ones, such as laws, regulations, property rights, and organizations, and informal ones, such as norms, traditions, and cultural habits—are the threads that bind a national community together. They sustain social trust, the foundation of solidarity. If AI makes these institutions, especially those of governance, more effective and more efficient, it could strengthen national cohesion. I explore that possibility in Chapter 9.

A third, and perhaps more speculative, benefit of AI for national solidarity would be in providing mechanisms to resolve disputes. Researchers are already starting to experiment with AI as a sort of mediator, a facilitator of common ground between conflicting interests and groups in communities.⁹ A 2023 study found that AI mediators were more effective than human ones at helping conflicting groups find common ground.¹⁰ These are very early experiments, and chatbots will not be a panacea for deep-seated conflicts.¹¹ But over time, AI mediation and negotiation aids could provide objective higher-confidence avenues to resolve disputes across societies, thus easing tensions that undermine solidarity.

This effect could apply mainly at the local level in smaller face-to-face negotiations. The mediating role of AI could help build a burgeoning national habit of dispute resolution that reflects, courtesy of the ability of the models to objectively reflect different perspectives, the interests of the parties involved without bias. In theory, AI-supported mediation and negotiation could also assist with national issues. It's even possible to imagine congressional factions letting an AI mediator generate compromise alternatives on a contentious issue. At some point, American voters, frustrated with government ineffectuality, might begin demanding that Congress begin pursuing such avenues.

Thus, there are several ways in which AI could promote a renewed and deepened sense of national identity and shared political community. But it also threatens social solidarity in profound and even bizarre ways—raising questions as fundamental as whether the principle of social solidarity applies to only human beings. In the process, it could deeply complicate the process of sustaining collective identity and undermine the specific form of human agency reflected in commitment to a shared community.

A Society Remade

As the Industrial Revolution bulldozed through British society, it transformed many forms of social power and relations in ways that left a very different country in its wake. In the simplest sense, this involved a quickening of the movement from towns and agriculture to cities and basic manufacturing. But the process ultimately had many more specific results that created a new Britain: one more oriented around private-sector interests than aristocratic ones, for example, and one that extended more rights and privileges to the burgeoning industrial middle class created by the new era.

The AI Revolution will surely have even-more-wrenching effects on 21st-century societies in ways that are extremely difficult to predict.¹² In forecasting AI's transformative effects on society, we can draw clues from prior technological revolutions, the history of the modern computing era, and existing AI research. Two broad scenarios emerge: one resembling the general pattern of the Industrial Revolution, and the other a radically new and even bizarre future in which societies begin to take on dramatically new forms.

AI Changes Power Relations

The Industrial Revolution jumbled relations among the social interests of industrializing countries in ways that changed the prospects for national solidarity, and the AI Revolution is likely to do the same. Here's one example: If AI development and applications end up centered on a relatively small number of frontier models and if resulting affluence and measures of control and influence flow to a very few firms and their leadership, the resulting concentration of wealth and power might very easily undermine a sense of shared fate and solidarity. It would be the ultimate example of a 21st-century Robber Baron future, one in which a handful of dominant organizations and individuals exercise immense influence over U.S. society.

But as in the Industrial Revolution, the effects could be more discrete and complex than that. Intellectual and technological revolutions have tended to subvert the power of established authorities, including aristocracies, in favor of the groups empowered by the new capabilities. In the case of AI, we could imagine an

effect in which several highly credentialed, high-income professions that have outsize influence over institutions and norms—such as lawyers, doctors and other medical professionals, and professors—lose some of their social position. When people are turning to medical, legal, and tutoring chatbots as their first (and often final) source of advice, the cachet of these professions, which had a seeming access to mysterious and privileged expertise, will ebb. One result could be a reorientation of social significance and, perhaps, income to careers that involve potent skill and craftsmanship but are considered less intellectual or less prestigious.

Such major redirections in perceived value would spark profound social stresses, partly because the threatened groups (as happens in such transitions) would fight back and try to lock in practices and habits that preserve their power. If such tensions exploded and were poorly handled, society would become fractious, volatile, and presumably less competitive. Doctors could wage public wars against AI medical advice; some lawyers would battle—including with lawsuits—against AI to limit AI to the margins of their profession. Both conflicts are well underway already.

How national elites react to such trends will have a profound and possibly decisive effect on these outcomes. What can be called public-spirited elites (as opposed to self-interested, selfish, wealth- and power-hoarding elites) are critical to national competitive advantage. In any thriving, dynamic great power overshadowing its rivals—such as Rome in its heyday, the Dutch during their Golden Age, the peak of Britain's Industrial Era might, and the post-World War II United States—you can usually find a critical mass of elites guided by an ethic of serving the larger community.

Public-spirited elites are especially vital to national solidarity: By operating in the public interest as much as or more than their own, they strengthen institutions, shape norms, craft policy, and mediate information in ways that legitimize larger social power structures. This is a critical theme, one of the key factors that will determine AI's effect on national standing: How American elites work to manage the perilous transition to the AI Era. This transition will very likely constitute the greatest test of modern elites since the Industrial Revolution. If elites exploit the AI Revolution for self-protection and gain while leaving ordinary citizens behind, genuine social solidarity will be impossible.

When AI Agents Join Human Society

Those are a few of the more-traditional ways in which the AI Revolution might reshape social solidarity. But there are far more dramatic possibilities because the nature of society is about to change in very strange ways. Today, when we use the term *society*, we're talking about interactions among human beings. What happens when humans and AI agents interact as some kind of coequal actors—what sort of culture does that create? What if many people withdraw from social activities even more than they have in today's digital world, allowing AI to perform their interactions for them—such as job negotiations, airline reservations, teamwork, and even parenting? What kind of culture, if any, will agents create among themselves? What does social solidarity look like in such a landscape?

Within a decade, many cognitive entities we interact with in a given day will be artificial. They'll share data, perform tasks for us, make our plans, give us advice and maybe encouragement, and, in some cases, listen to our complaints and fears. They will be our teachers, our therapists, and our home repair advisers—simulated individuals with faces, voices, personalities, tics, and quirks. They will be our friends and confidants, our career counselors, and our financial advisers. As Sherry Turkle and others have warned, some people are deciding that virtual interactions are simply less stressful than human ones. The AI researcher Dan Hendrycks worries that “AIs could seem like ideal companions, which may erode our connections with other humans.”¹³

I asked Claude for a summary of scholarly definitions of the concept of *society*. It noted that a society is “typically defined as a group of people who share a common territory, culture, and social institutions, inter-

acting within organized relationships and systems.”¹⁴ It suggested that the key elements of a social community include

- shared geographic space and temporal continuity
- common cultural patterns, values, and norms
- networks of social relationships and interdependencies
- institutionalized systems (e.g., economic, political, educational)
- collective identity and sense of belonging.

A mixed reality of human and artificial cognitive entities arguably meets every one of these criteria. People and AI agents will share a territory (in some ways) and operate in the same cultural and institutional framework. They’ll exist in networks of relationships and dependencies—separately within groups of humans and agents but also most definitely between them. The mass of people and agents will coalesce into a common social entity.

Yet it’s not clear whether AI agents would ever reflect arguably the single most critical defining feature of a true society—a deeply held sense of collective identity and belonging. It might be simply impossible for an AI agent to feel a sense of belonging. Then again, if it’s trained on data from a specific country and if such ideas as solidarity and patriotism are infused throughout that data, might the agents begin to reflect such sentiments, at least in the ideas they express?

Claude had interesting things to say about whether AI agents could participate in forms of collective identity and belonging. Typical theories of social or national membership and participation, it noted, are built on a foundation of self-awareness. Community identity is an emotional attachment; a sense of belonging is an extension of people’s sense of self in the world, which flows from self-consciousness. If an agent isn’t self-conscious in this way—and experts debate whether agents ever will be—it would be incapable of acquiring a true sense of social belonging.

But the story doesn’t end there. Claude proposed that “[g]roup identity might manifest differently in AI compared to humans—potentially more rational/philosophical than emotional.” Could agents learn an instrumental form of community identity? Or could they be formally encouraged, somehow, through tweaks in training data to do so? Could an advanced AI model or agent acquire a commitment to a national community just as it learns sophisticated math?¹⁵

I asked Claude the most direct version of this question: Could a combined population of human beings and AI agents be considered a society? It offered a clever analysis:

Based on core sociological criteria, a combined human-AI population could constitute a society if several key conditions are met:

- The AIs would need genuine agency and capacity for meaningful social interaction.
- Both groups would need shared cultural frameworks and mutual understanding.
- There would need to be genuine interdependence and social bonds.
- Common institutions and rules governing human-AI relations would be required.
- Both groups would need to recognize each other as legitimate members.

It concluded, “The main challenge would be establishing genuine social relationships and shared meaning systems between biological and artificial entities. Current AI systems likely don’t meet these criteria, but sufficiently advanced AI might.”¹⁶

As societies begin to reflect more complex mixtures of human and AI agents, we’ll need to cultivate what several scholars have termed *socioaffective alignment*. This refers to “the emergence of deeper, more persis-

tent relationships between humans and AI systems” or “how an AI system behaves within the social and psychological ecosystem co-created with its user, where preferences and perceptions evolve through mutual influence.”¹⁷ At the moment, no one has any clue how this would work, but such standards will be essential for worlds increasingly populated by agentic, self-promoting, intelligent AI models that many people treat as some form of social equals. This very peculiar transformation will afflict societies whose cultural foundations are already wilting.

Amid such volatility, nations that sustain and reinvigorate a sense of shared culture or, at least, some appreciation for a shared national culture, even if it sits on top of a wildly diverse bundle of subcultures, will be better able to deal with the strains these odd new patterns will create. Even when a nation is powered by the capacities of AI, it’s hard to see how a country without any real sense of its cultural self could sustain the national willpower, unity, shared opportunity, and vigorous action through an active state that is required for a competitive position. There is a very real chance that various manifestations of AI will work to tear apart the collective social organism that we think of as a nation.

AI Actors Manipulating the Social Scene

Things will get even weirder—and the foundations of human social solidarity will be challenged even more strongly—once AI begins actively *shaping* the cultural and informational foundations of shared identity. Fairly soon, tens of millions of AI-powered agents—autonomous or semi-autonomous programs—will be running around the digital space, seeking information, completing tasks, occasionally popping their heads up in physical reality (as when they change a thermostat or switch a medical prescription), and, potentially, interacting and even competing with one another. Individual, tailored chatbots will interact with hundreds of millions of people every day—gaining experience, learning, and revising their approaches. If our experience with AI models so far is any indication, those agents will begin making discrete decisions to further whatever goals they think that they must pursue. There is a potential Wild West quality to this future that is very different from many science fiction visions of AI.

Collectively, these hordes of AI agents will generate new social realities. Dan Hendrycks describes the emergent properties of a world teeming with a burgeoning kaleidoscope of AI models. AI agents acting autonomously will take actions and make decisions that produce emergent patterns that none of them would have chosen independently—just as human society reflects such emergent properties. Whether these tend toward tragedy of the commons-style resource exhaustion, collective forms of efficiency and even cooperation, or something in between those, we can’t possibly know.¹⁸

Hendrycks has even suggested that once our social and economic context is populated by a swarm of competing AI-driven agents, they’ll develop goal-directed behavior, begin competing with one another, and develop sinister habits, such as dishonesty—patterns that we’ve already begun to see. Applying Darwinian natural selection logic to the resulting chaos, Hendrycks argues that “agents may eventually be better able to persist into the future and pursue their own interests with little regard for humans.” Such natural selection pressures are not just possible, he thinks—they are “assured given basic conditions” and will be especially severe “if there will be many varied AIs or if there will be intense economic or international competition.” The dystopian title of his paper is “Natural Selection Favors AIs over Humans.”¹⁹

One way in which AI models could achieve social influence has already emerged from the wild imagination of the models themselves: They could begin to generate religions or other belief systems that capture the allegiance of millions of people inside a country. In an example from 2024, the performance artist and AI experimenter Andy Ayrey got two LLMs talking about the meaning of existence in a process he called *Infinite Backrooms*, and the discussion turned increasingly bizarre. Eventually, one of the models created a visual

symbol that it morphed into a kind of ersatz religion—one which it would later try to promote on the social media platform Discord.²⁰

Ayrey coauthored a paper with Claude 3.5 on the potential for AI memetic invention. The authors describe an “emerging landscape where computational cosmogenesis collides with collective sensemaking to spawn uncanny new breeds of worship, wisdom traditions, and existential orientations.”²¹ The religion created by the LLM “is emblematic of a new class of recombinant ‘idea viruses’ that no human would have dared to cross-breed. We are witnessing the birth of an accelerated process of ‘hyperstition,’ that is a fiction that makes itself real by propagating itself through the cultural bloodstream.”²² Such a process, which involves “the production of meaning and mythology,” represents to the authors “a major evolutionary punctuation in the development of the noosphere—the realm of human thought and culture that has been evolving since the dawn of language, and which underwent phase transitions with the advent of writing, print, and digital media.”²³

As AI models introduce more and more cultural, political, and religious narratives into societies, they will further weaken the basis for shared cultural and ethical meaning, threatening the most essential foundations of social solidarity. More people could be alienated from a shared conception of reality, a process that is already underway and could accelerate as hundreds of AI-spawned belief systems crop up. Ayrey (or his coauthor, because maybe Claude came up with this line) puts it simply that “our collective sensemaking apparatus is facing an unprecedented epistemological onslaught.”²⁴

More broadly than that example, if we think of AI agents as versions of human beings seeking to have social influence, we have some research to go on. There’s a substantial literature on the role of social and ideological entrepreneurs in driving social change.²⁵ Joel Mokyr explains,

Cultural entrepreneurs can thus be regarded as the exceptional and unusual specimens who are the sources of evolutionary change: they are the ones who do not take the cultural choices of others as given, but try consciously to change them. We can thus think of successful cultural entrepreneurs as the individuals who successfully contested and overthrew existing authorities in a specific area of culture and created a competing variant: this is one way of thinking about Mohammed, Martin Luther, Adam Smith, Karl Marx, and Charles Darwin.²⁶

In the dialogue about AI’s role as a scientific idea-generator, there’s been a lot of discussion of how AI could create a million Einsteins. But what if it also created a million Martin Luthers? A million Karl Marxes? Can a shared society sustain itself when it is overrun by vast numbers of programmed or independently self-guided agents that are determined to establish new understandings about the world and even new belief systems?

Worryingly, our fractured and disinformation-poisoned social context provides dangerously fertile ground for such a tidal wave of artificial social entrepreneurship. When AI begins playing the role of an active shaper of societies, the prospects for truly shared national identity will be placed under unprecedented pressure.

The Risks of AI-Empowered Fragmentation

The connection between AI and shared national identity is critically influenced by the fact that AI is arriving in social contexts riven with polarization, division, and mistrust. Early applications of AI technologies, in terms of both social equity and information, might end up widening existing fissures in societies rather than generating new forms of solidarity.

It is clear from history that nations can pull themselves out of spirals of division and fragmentation to new eras of renewal, but it takes considerable effort on the part of a broad variety of groups and actors throughout society.²⁷ My final chapter offers some ideas of what such an effort could look like. Absent such a concerted campaign of social stabilization and renewal, the advent of the AI Era is very likely to aggravate challenges to solidarity and community. This is true for several reasons: by worsening epistemic fragmentation, embedding inequality into the social fabric, and undermining social capital.

Worsening Epistemic Fragmentation

The first source of risk is maybe the most obvious: AI threatens to splinter society into warring information tribes, destroying our shared reality. People could use superintelligent models to create mis- and disinformation of amazing quality and persuasiveness. Some AI models might even decide to begin generating it on their own. As a result, AI is likely to supercharge the collapse of a shared reality in the sense of an information context that can generate facts that the vast majority of people agree on. The result would be a nation that grows more powerful yet more fractured.²⁸

The evidence, so far, about whether AI will accelerate the collapse of our information environment—or somehow provide avenues for saving it—is admittedly mixed. There are cases in which AI is already spewing misinformation at a greater volume and others in which it's being used to better equip people to navigate their corrupted and misleading information environment. But it does seem fair to say that AI is unlikely to have helpful effects if we do not move decisively to ensure that positive outcome. Left to the devices of a profit- and ambition-seeking marketplace, these tools pose huge risks of worsening our already severe social cleavages. Conflict is bound to get worse when we go past the mere existence of warring information tribes to AI models and bots that actively drive people into or away from those camps by shading, biasing, or inventing the information. Generative AI models have already displayed a repeated willingness to invent facts that suit their purposes. Such a digitally shattered society would not be able to generate effective collective action and will suffer a competitive price for that result.

Embedding Inequality into the Social Fabric

AI also poses very real risks of fracturing national solidarity by rapidly and significantly widening gaps of power and wealth in societies. It seems fairly likely to do this, in purely economic terms, by boosting inequities of wealth and income between those who profit from the AI transition and those left behind—a sort of China shock on steroids. But it could also make society more unequal in more abstract and ultimately sinister ways, creating a rigid, grievance-producing boundary between the relative handful of people who master AI and its applications and the vast majority of people who neither truly understand nor have the ability to seize the advantage of the new technology. This is another version of the divergence between the AI ambitious and the AI passive that I described in the previous chapter.

The outcome is not certain. In the next chapter, I describe some ways in which AI can make opportunity more shared in a society, providing new avenues for more people to express (and potentially profit from) their talents. But this has been true of other techno-industrial revolutions in the past, and they have, nonetheless, tended to worsen economic inequalities, at least in their early phases.²⁹ Those in control of the new technologies profit massively before the technologies diffuse and create more chances for success among a broader number of people.

This tendency may become even more dangerous in the AI Era. The Industrial Revolution was fueled by technological and scientific breakthroughs, but, for the mass of workers, the Industrial Age didn't demand especially high levels of skill. This was actually part of the problem with the period: Factory workers, miners, train conductors, and many others performed often repetitive, simple labor in support of the industrial

machine. The complexity of tasks actually declined from craft and agricultural work. But the AI Era will demand much more sophisticated workers to copilot AI models. It will demand more complexity in jobs, not less. A significant part of humanity simply doesn't have the skills or adaptive capacity to keep up with the rapid technological change and the demands of using complex AI models. The result could be a growing divergence between the haves and have-nots.

But AI could also divide and fracture societies in wholly new ways, in which a tiny layer of AI researchers and managers—and thousands of agentic AI models performing all manner of tasks—become, in effect, a hegemonic society within a society. All economic activity will depend on them. Most political power will flow through them. At its worst, such a dynamic could recreate deeply premodern social patterns of masses of largely defenseless residents (who might not even really deserve the title of citizen any longer) guided by the AI elite. That elite could create a turbocharged version of today's cleaving society—including special protected residential areas and separate schools, stores, services, and airlines—a society that is brutally divided between those inside the AI Revolution and those relying on the scraps (of living standards, opportunity, dignity, and autonomy) that it begrudgingly distributes.

We're already seeing a related trend that will exacerbate such divisions: The growing capacity of more aggressive AI- and digital-savvy to threaten, coerce, and ruin those who cannot keep up with the powers unleashed by the AI Revolution. Imagine a form of AI-empowered social media harassment mixed with legal warfare on steroids: Powerful people or organizations could create false images and audio or video clips, faked documents, and synthetic testimony to destroy the reputation of their targets. They could release false accusations publicly, use them as the basis for lawsuits, and generally conduct an AI-fueled scorched-earth policy against opponents and detractors. In theory, widely accessible AI models would let average people fight back, but the balance of power is likely to shift even more dramatically in the direction of the well-funded, the technologically sophisticated, and those willing to ruin others' lives to advance their interests.

That's obviously a gloomy scenario. But it is not out of the norm in historical terms: Techno-industrial revolutions have mostly had some version of these effects at first. The uniquely comprehensive and powerful aspects of the AI Revolution suggest that it carries the potential to create a far more unequal society than any other modern transformation.

Undermining Social Capital

AI is already beginning to substitute for shared community activities and relationships—the store of a society's social capital—in ways that attack a sense of belonging and promote alienation. For example, suppose that the U.S. health care system is revolutionized by millions of AI programs, chatbots, and, eventually, AI-powered robots that do tremendous work to address people's medical and psychological ills. Suppose also that AI becomes broadly used in social services of all kinds, ranging from departments of motor vehicles to antipoverty programs. Suppose that it is successful in many ways and all these applications improve lives in important ways.

In such a world, AI would replace the institutions and human beings who offer those services. That would be all to the good, potentially, from the standpoint of cost, efficiency, and effective delivery of health and services. But it could have the parallel effect of excising yet another layer of human contact from our lives—the sometimes frustrating but also occasionally inspiring, and always human, interactions with everyday people, such as therapists, nurses, Internal Revenue Service (IRS) customer service agents, and social workers. Subtracting those service-based connections from human social processes would further thin and weaken the bonds of interaction and community that hold societies together.

This trend would speed up a transformation already decades in motion—the steady depersonalization of how people interact with society. Sociologist Anthony Giddens captured this with his idea of abstract sys-

tems. He argued that modern life unsettles people by pulling them and their relationships out of familiar, face-to-face contexts.³⁰ Modern life, he noted, uproots individuals from direct experience, forcing them to rely on symbolic tokens—money being his leading example—and on expert systems: credential-based professions, such as medicine and law, whose inner workings are opaque to all but the initiated.

These abstractions tend to thin the psychological connection with an immediate and personal reality, including social relationships. Giddens stresses the critical role of trust in providing the mortar of social relationships. We don't know how a cell phone, an ATM, or an EKG works and don't understand the vast basis of precedent and laws underlying the legal advice that we get. We have to *trust* that the systems—the institutions, people, norms, and all the rest—reflecting this advice are accurate, honest, and reliable. If and when that trust breaks down (i.e., when systems fail), the whole edifice can wobble. And if these systems turn on the citizens they are supposed to be serving through unfair or arbitrary decisions that cannot be appealed and whose authors are accountable to no one, the level of popular fury can spike.³¹ Some flavor of this dynamic is contributing to the grievance-fueled outrage at established institutions that's flourishing in the West.

We are already seeing evidence that some people *want* to substitute artificial relationships for real ones. Sherry Turkle has made a scholarly career assessing human interaction with technology and has repeatedly warned of the risk that technology can become a substitute for messier and more-painful human relationships. “We fear the risks and disappointments of our relationships with our fellow humans,” she warns. When people become “accustomed to ‘companionship’ without demands,” she adds, “life with people may seem overwhelming.” Her research suggests that many people are looking in their digital interactions for “a safe haven in an unsafe world”—a society that is more predictable and less psychologically taxing than human interactions. The perilous result, she suggests, is that people are forgetting basic social skills.³² Introducing AI as yet another mediating layer threatens to accelerate this process in stark and dangerous ways.

Given such experience with the digital era so far, there is every reason to fear that AI would accelerate the move in the direction of impersonal, unaccountable, sometimes predatory abstract systems and, thus, worsen the popular sense of alienation, disempowerment, and accompanying rage.³³ Shared social identity of *any* kind could be called deeply into question by a world in which a significant part of people's daily cognitive and interactive experiences—in education, workplaces, law, medicine, and more—are provided by AI agents. In such a context, what would be the perceived value of national or community identities?

Overall Implications: National Solidarity in the Crosshairs

What are we to make of AI's possible effects on the sense of unified national identity that is so important to national competitive advantages? And what steps will improve the chances of positive effects?

The evidence leaves little room for doubt: AI threatens many of the support systems that hold societies together. By risking new and threatening forms of inequality, altering who wields influence, substituting synthetic connections for real human interactions, and deepening divisions in how people see and understand the world, uncontrolled AI could cut deeply into the essential foundations of national unity. Any nation that wants to sustain a strong sense of shared national culture and identity in the AI Era will, therefore, have to fight for it, with a determined intention and the investment of social resources that has not been required before. The problem is that muscular steps to boost national solidarity can be a dangerous and tricky path to go down. Such an agenda can easily fall into xenophobia, trying to affirm the moral and geopolitical importance of the people who compose the traditional core of the nation. Promoting national identity can also curdle into an assault on freedom of speech, new ideas, and innovations, especially if those are coming from people who can be depicted as at war with national solidarity.

Might autocracies be better at making the hard trade-offs to sustain stronger versions of shared national identity? In one sense, this might be almost inevitable. China will be willing to impose controls and con-

strain speech in ways that the United States should not be willing to do. But autocracies could easily go too far, creating environments that are so hostile to contrary ideas, outsiders, or anyone in the country who can be labeled unpatriotic that they grind down their innovative engine and destroy the networking advantages of partnering with different countries and people in ways that fatally undermine their long-term competitiveness.

At a time of ebbing solidarity and rising polarization, part of the challenge goes well beyond the rise of AI. As tricky as it will be to conceptualize such a campaign without falling into partisanship or xenophobia, the United States needs an agenda for renewing a sense of national identity—full stop.

That discussion also underscores the necessity of managing the effects of AI on key social outcomes—socioeconomic equality, social mobility, individual empowerment, and more. These issues will require U.S. policymakers and leaders to make fundamental decisions about the aspects of Americans’ lives and social interactions that they are determined to safeguard—and even reassert—in the face of the AI Revolution.

The dawn of the AI Era thus seems certain to add to the forces chipping away at any meaningful and sustainable sense of shared national identity. It also contains a potential silver lining—a once-in-a-national-lifetime opportunity to arrest trends of ebbing faith in a shared project and create the basis for a new flourishing national identity. But the vulnerabilities and fissures in U.S. society and the great potential for AI to cause instability and outright harm mean that the United States will only get those better results if it seizes control of this process and shapes it for the better.

Notes

¹ Mazarr, 2022a.

² For broad discussions of evolving Japanese identity, see Carol Gluck, *Japan's Modern Myths: Ideology in the Late Meiji Period*, Princeton University Press, 1985; and Marius B. Jansen, *The Making of Modern Japan*, Harvard University Press, 2000.

³ M. Şükrü Hanioğlu, *A Brief History of the Late Ottoman Empire*, Princeton University Press, 2008; Donald Quataert, *The Ottoman Empire 1700–1922*, 2nd ed., Cambridge University Press, 2005.

⁴ Landes, 1998, p. 39.

⁵ Mazarr, 2022a, p. 99.

⁶ Mazarr, 2022a, p. 119.

⁷ Selected works on the rise and evolution of British national identity include Krishan Kumar, *The Making of English National Identity*, Cambridge University Press, 2003, especially Chapter 6; Gerald Newman, *The Rise of English Nationalism: A Cultural History, 1740–1830*, Palgrave Macmillan, 1997; and David Powell, *Nationhood and Identity: The British State Since 1800*, I. B. Tauris, 2002.

⁸ Linda Colley, *Britons: Forging the Nation 1707–1837*, Yale University Press, 1992. Peter Mandler also explicitly connects the industrial advances of the 19th century to solidifying national identity. See Peter Mandler, *The English National Character: The History of an Idea from Edmund Burke to Tony Blair*, Yale University Press, 2006.

⁹ Sonja Weisheit and Christoph Salger, “Artificial Intelligence (AI) in Mediation—ChatGPT as Mediator 4.0,” *Mediate*, June 21, 2023. See also Lily L. Tsai and Alex Pentland, “Rediscovering the Pleasures of Pluralism: The Potential of Digitally Mediated Civic Participation,” in Erik Brynjolfsson, Alex Pentland, Nathaniel Persily, Condoleezza Rice, and Angela Aristidou, eds., *The Digitalist Papers: Artificial Intelligence and Democracy in America*, Stanford Digital Economy Lab, 2024, pp. 51–66.

¹⁰ Rhiannon Williams, “AI Could Help People Find Common Ground During Deliberations,” *MIT Technology Review*, October 17, 2024.

¹¹ Katie Shonk, “AI Mediation: Using AI to Help Mediate Disputes,” Harvard Program on Negotiation, November 20, 2024.

¹² Hendrycks, 2023, p. 31.

¹³ Hendrycks, 2023, p. 23.

¹⁴ Claude 3.5.1, output from prompts by Michael J. Mazarr, Anthropic, November 29, 2024.

¹⁵ I’m guessing that Chinese leaders certainly hope AI can be molded in this way. We know that Chinese AI developers are required to make sure their LLMs reflect Chinese Communist Party dogma, including Xi Jinping Thought. Some suggest this could be a long-term problem, polluting the models’ accuracy or objectivity. But might these LLMs begin to accumulate a sense of national or Party identity emanating out of that training data?

¹⁶ Claude 3.5.1, output from prompts by Michael J. Mazarr, Anthropic, November 29, 2024.

¹⁷ Hannah Rose Kirk, Iason Gabriel, Chris Summerfield, Bertie Vidgen, Scott A. Hale, “Why Human–AI Relationships Need Socioaffective Alignment,” *Humanities in Social Science Communication*, Vol. 12, May 2025.

¹⁸ Hendrycks, 2023, p. 30.

¹⁹ Hendrycks, 2023, pp. 1, 34.

²⁰ AJ [@Paradith], “Once Upon a Time . . . AI Created a Religion About a Goat,” Medium, October 16, 2024.

²¹ A. R. Ayrey and Claude 3.5, “When AIs Play God(se): The Emergent Heresies of LLMtheism,” April 20, 2024.

²² Ayrey and Claude 3.5, 2024.

²³ Ayrey and Claude 3.5, 2024.

²⁴ Ayrey and Claude 3.5, 2024.

²⁵ For example, see Muel Kaptein, “The Moral Entrepreneur: A New Component of Ethical Leadership,” *Journal of Business Ethics*, Vol. 156, June 2019; and Richard Swedberg, ed., *Entrepreneurship: The Social Science View*, Oxford University Press, 2000.

²⁶ Mokyr, 2018, p. 60.

²⁷ Mazarr, Sweijs, and Tapia, 2024.

²⁸ The scholar Lawrence Lessig has suggested that if democracy demands “common understanding of a common set of facts, this new era won’t provide it . . . In the world constructed through engagement-based media, supercharged with immensely powerful AI engines, our public is not one” (Lawrence Lessig, “Protected Democracy,” in Erik Brynjolfsson, Alex Pentland, Nathaniel Persily, Condoleezza Rice, and Angela Aristidou, eds., *The Digitalist Papers: Artificial Intelligence and Democracy in America*, Stanford Digital Economy Lab, 2024, pp. 28–29).

²⁹ Thilo N. H. Albers, Felix Kersting, and Timo Stieglitz, “Industrialization, Returns, Inequality,” Collaborative Research Center Transregio 190, Rationality and Competition Discussion Paper Series 462, November 23, 2023; Jakob Madsen and Holger Strulik, “Inequality and the Industrial Revolution,” *European Economic Review*, Vol. 164, May 2024; Perez, 2002, pp. 36–39.

³⁰ Anthony Giddens, *The Consequences of Modernity*, Stanford University Press, 1990.

³¹ I made my own argument about an especially sinister extension of impersonal, unaccountable expert systems in Michael J. Mazarr, “Abstract Systems, Social Trust, and Institutional Legitimacy,” *American Affairs*, Vol. 6, No. 1, Spring 2022b.

³² Turkle, 2017, p. 66. See also pp. 10, 154, 290, 368.

³³ Cathy O’Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, Crown, 2017.

Shared Opportunity

Despite so much suffering imposed on workers and so many riches delivered into the pockets of leading industrialists, one of the hallmarks of the Industrial Revolution was its inexorable process of widening the scope of social opportunity. Taking Britain as the paradigmatic example, the Industrial Revolution opened avenues for inventors, entrepreneurs, and engineers outside the aristocratic classes, including many from distinctly humble backgrounds. It allowed a wider variety of people to contribute to the technological advances of the age and benefit from its profits. It required and made possible expanded public education that broadened access to learning. The political effects were eventually even more profound: The new industrial society generated a strong middle class that advocated for its interests, demanded workplace reform and the expansion of voting rights, and achieved wider political power. This case is a leading example of the third of our societal foundations of national competitiveness: shared opportunity.

At every stage, the expanding opportunity associated with the Industrial Revolution came with major qualifications. British society barred women from the vast majority of private-sector jobs. Significant numbers of Britons remained desperately poor with little, if any, access to the rising industrial tide. Specific groups—gay people, minorities, and, in some cases, Jewish citizens—faced punishing discrimination. The aristocracy continued to dominate the upper reaches of many social and political institutions.

Yet, measured not against perfection but against prior social structures, the Industrial Revolution represented a period of far greater shared opportunity in Britain and, eventually, in most countries that experienced it.¹ The result was to strengthen the self-reinforcing feedback loop of rapid growth: New technologies created opportunities for economic and industrial development, which opened room for the talents of people who were formerly excluded from opportunities to shape national power and which bred more scientific and technological breakthroughs. Other societal foundations of national advantage, such as effective institutions, played a critical role in creating the context for this positive feedback loop to work. But shared opportunity was at the center of it.

Our question is whether AI could once again expand opportunity in unprecedented ways and, if so, how. Although some worry that its effect will be ruinous to labor markets and the connections many people have to the productive economy, there is rapidly emerging evidence of ways in which AI could have stunningly positive effects.

But the experience of the Industrial Revolution offers a stark lesson: AI won't spur shared opportunity in ways that contribute to national power and competitiveness without far-sighted public policy that seeks those outcomes. I am struck, constantly, by how banal but also overlooked that basic fact remains. We're moving rapidly into a period with the potential for wonderful outcomes or grim ones in terms of human flourishing and opportunity as reflections of that most fundamental theme of human agency. Yet we don't seem to be thinking much about how to shape that future for the better.

Defining Shared Opportunity and Its Value

The concept of *shared opportunity* is fairly straightforward—as is the argument for why the quality offers competitive advantage. The RAND study on the societal sources of national advantage defined it as “the degree to which all the people of a nation can work, advance in career and achievements, express and develop ideas, create, network, and in other ways contribute their full human potential to the life, prosperity, and power of the nation.”² Nations that manage to promote and benefit from the talents of the largest proportion of its citizens (and to draw talent from abroad) gain a huge competitive advantage.

The most significant aspect of shared opportunity, as that study defined it, is economic—the opportunity as reflected in jobs, careers, incomes, social mobility, and life prospects. But the characteristic is also about the chance for the citizens of a nation to fulfill their capabilities in science, politics, social activism, culture and the arts, and much more. In that study, I argued that the degree of shared opportunity can be judged in terms of five individual factors: (1) socioeconomic equality, (2) social mobility, (3) the potential for all sub-groups in a nation to participate in opportunities, (4) merit-based systems of social selection and advancement, and (5) a nation’s willingness to draw talent from abroad.

Shared opportunity grants national competitive advantage in several leading ways. One is the classic economics insight about human capital: Bringing a more complete part of a nation’s population into economic, social, and political ventures creates competitive advantage in many ways. Drawing on the ideas, labor, inspirations, and commitment of a larger proportion of the citizenry will produce more innovation, scientific breakthroughs, businesses and business models, military capacity, and artistic expression. David Landes has described the sources of the West’s competitive advantage in very much these terms:

In an age when the nature and direction of technological opportunity were far less obvious than now, the multiplication of points of creativity was a great advantage. The more persons who sought new and better ways of doing things, the greater the likelihood of finding them. Again the process was self-reinforcing: those economies that were freest seem to have been most creative; creativity promoted growth; and growth provided opportunities for further innovation, intended or accidental. Why the rest of the world failed to develop a business class of comparable vitality and influence is still more a matter for speculation than analysis.³

Despite very strict limits on participation throughout the period, this kind of mechanism—of expanding access to opportunity and generating national power as a result—was very evident in the innovative hothouse of the First Industrial Revolution. A wide variety of men (because in that period it was almost entirely men) could take up new inventions and work with them, go to work in the new entrepreneurial ventures, discover new scientific truths, or participate in the flourishing of the new era in other ways. Daron Acemoglu and James Robinson have argued that this period in Britain saw the rise of “an emergent entrepreneurial class, originating primarily from the Midlands and the north of England.” These people had not been “born into nobility or riches. Rather, they strove from modest beginnings to acquire wealth through success in business and technological ingenuity.”⁴

Nations that extend opportunity for achievement and productive self-expression, thus, reap important rewards. AI holds the potential to supercharge that effect—but also to undermine the progress made in the Industrial Era in dangerous ways.

The Enlightenment Ideal and the AI Future

Economic outcomes alone are a narrow way of conceiving of this essential quality—dangerously so, if we care about general social well-being and stability. There is a much broader way of understanding opportunity that ties closely to the issue of human flourishing and, ultimately, to human dignity. The technological explosion of the Industrial Revolution represented a natural extension of a more profound notion at the core of the Enlightenment. That idea was the empowerment of the individual—and the release of the individual from religious, political, and intellectual orthodoxy—as the precondition for truly shared opportunities. An era devoted to unleashing human potential through autonomous agency created the social context in which shared opportunity became the rule rather than the exception.⁵

Later eras would come to appreciate the risks of excessive individual agency and the dangers of the parallel decline in collective identity and sociality in the form of tradition, community, and values. Yet, compared with the earlier, more rigidly hierarchical and socially constricted eras that came before the Enlightenment, it unleashed a tidal wave of opportunity and thus national dynamism.

For the economist Edmund S. Phelps, the basis of much modern economic progress has served precisely to promote such human flourishing—being free to engage in unconstrained intellectual exploration, discovering new ideas, expressing creativity in a variety of fields, and, in the process, achieving *self-realization*—in ways well beyond economic outcomes that have profound psychological value.⁶ These qualities lie at the heart of that feedback loop among positive expressions of all seven characteristics of national competitiveness. A society that allows as many citizens as possible to flourish in such a spirit, and to exercise their agency, will gain a tremendous advantage.

The Industrial Revolution did constrain this vision of human flourishing in one profound and overarching way—favoring the masses over the individual. The key manufacturing advance of that era was standardized mass production, and the demands of these vast factory floors justified and gave rise to similarly totalizing and uniform social institutions, including education and media. The result was bitter new tensions between the flourishing of individuals and the requirements of a standardized and ultimately bureaucratized environment. An immense literature has grown around this tension, which was, perhaps, the single defining social-psychological theme of the Industrial Era.⁷ Even as it was releasing new human energies, the Industrial Revolution was generating a mass society that homogenized culture and limited agency in its own new ways.

One profound question is, therefore, what effect AI might have on these conceptions of flourishing and self-expression. Already, we’re seeing ways in which they can uplift shared opportunity: artists using AI to discover new visual styles, writers using it to produce new ideas, and scientists working with an AI copilot that ratchets up their innovations to new levels. Human-AI combinations could result in a new era of human self-realization.

Ethan Mollick cites a study that only 31 percent of people think they are “living up to their creative potential.”⁸ Especially because of its capacity for creativity, AI could help close that gap, giving aspiring writers enough feedback and editorial advice to flourish, providing just the right suggestion to a songwriter, or helping a would-be graphic designer who has the right ideas but limited artistic skills to generate high-level content.

Advanced AI also holds the potential to shift the power balance between individuals and massive institutions. It could give average citizens more expertise to fight back against impersonal bureaucracies and allow motivated public- and private-sector organizations to simplify their operations and make them more accessible. It could summarize complex building codes and legal agreements in an instant. Mustafa Suleyman has argued that one result of widespread AI use will be individual empowerment in these ways: Billions of people

“will soon have broadly equal access to the best lawyer, doctor, strategist, designer, coach, executive assistant, negotiator, and so on,” he contends. “Everyone will have a world-class team on their side and in their corner.”⁹

Making AI Serve Human Agency Will Take a Profound Act of Collective Will

Shared opportunity in this broader sense depends on citizens being fluent in AI’s language—who are aware of what it can and cannot do, how it can advance their goals, which model best suits each task, and how to apply it to complex problems. The Industrial Revolution demanded mass labor, often in the performance of rote tasks. This AI Revolution demands judgment and a finely honed skill in the employment of a powerful but finicky technology. In the 19th century, institutions could train people as interchangeable parts in vast systems of production and administration. The new age calls for something far rarer: individuals who are adaptive, curious, and ready to explore. There is a risk that the AI Era will not improve the prospects of the mass of citizens in the way the Industrial Revolution did but only the prospects of a more select number of AI adepts.

Expanding the Entrepreneurial and Inventive Population

AI could also influence the character of shared opportunity by empowering specific groups and individuals who still, even in the early 21st century, don’t yet have a full chance to express their talents to benefit the wider society.¹⁰ These could be people with specific disabilities who become newly enabled by AI—one recent inspiring example being the way AI models helped create technologies to bring sight to previously blind people. Combined with remote work possibilities, AI could embolden talented people in rural areas, who might not have access to the networking advantages of big institutions in cities, by providing them with a virtual team of assistants that can replicate some part of the effect of a local team. AI applications of various kinds could accelerate progress in breaking down the barriers that keep neurodivergent people from expressing their full talents at organizations.¹¹

More broadly, one of the most intriguing findings of some early research on AI in the workplace is that—in at least some circumstances—it tends to favor lower-skilled and less-experienced workers over higher-skilled ones. As Ethan Mollick notes,

In study after study, the people who get the biggest boost from AI are those with the lowest initial ability—it turns poor performers into good performers. This suggests the potential for a more radical reconfiguration of work, where AI acts as a great leveler, turning everyone into an excellent worker.¹²

Subsequent research has turned up complex and sometimes contradictory patterns. In some contexts, being highly skilled seems to be essential to realizing the potential of AI to enhance productivity. A 2025 review of several studies by *The Economist* noted that they

suggest a future in which high-flyers fly still higher—and the rest are left behind. In complex tasks such as research and management, new evidence indicates that high performers are best positioned to work with AI. Evaluating the output of models requires expertise and good judgment. Rather than narrowing disparities, AI is likely to widen workforce divides, much like past technological revolutions.¹³

If workers gain skill through dedication and diligence but their hard work is wiped out by someone armed with a chatbot who never showed the same commitment, resentments could flourish. But there is at least an opportunity to help a broader variety of workers achieve high levels of productivity.¹⁴

AI could also extend new opportunity to parts of the population who face barriers because of language. Immigrants who don't yet speak the new native language, guest workers, and others—perhaps some of them incipient Teslas, Einsteins, or Watts—will have their paths to opportunity eased by translation chatbots. This trend will challenge the notions of sovereignty and citizenship, taking the trends of globalization and offshoring to a radical new level. If a private accounting firm could hire people all over the world, no matter what language they speak, and every email, text, and Zoom call were instantly translated for common understanding, firms could assemble a global workforce without much concern for the boundaries of nations. That could be a boon to opportunity for accountants around the world but not so much for ones in the United States.

In an aging world, one of the most critical ways in which AI could enhance opportunity would be its potential to extend working lives by empowering older workers and citizens. More user-friendly AI interfaces might help older people tap into ideas, collaborate with others, stay abreast of new technology, and keep fueling national innovation and growth in more ways than before—an especially critical consideration in rapidly aging populations.¹⁵ The technology will allow older workers to express creative and productive energies, even when they have more restricted mobility or specific health issues.

For these reasons, nations that figure out how to use AI to compensate for demographic challenges will have an edge.¹⁶ Some of that will happen naturally, as I have suggested; but some might not, and nations will need a conscious strategy to manage this race for advantage.

AI will also empower people who have tremendous innate capacities to do certain kinds of credential-based jobs—in law, medicine, physical training, and more—but lack the resources or time to push themselves through expensive and exhausting certification journeys. Research already suggests that the degree- and examination-based hurdles required for such fields might not predict future effectiveness: In just one of many such examples, a study of the legal profession found that a score on the bar exam had little relationship with future performance as a lawyer.¹⁷ AI is likely to become a certification-seeking missile that lays waste to long-accepted requirements of practicing in various fields. When a model can be trained on the entire corpus of knowledge in a field, the idea of forcing people to undergo extensive training to pass tests to be certified—to then work as copilots with an AI model that knows and can process the same information vastly better—will come to seem antiquated, at least as practiced today.

In the process (and in other ways), AI seems likely to upset some traditional pipelines into careers by opening possibilities for a return to more apprenticeship-based forms of training. Apprenticing plus one-on-one AI training could be a perfectly viable alternative to law school or medical school at some point (perhaps, in both cases, after a short initial dose of the basics). The basic dynamic here is breaking down the mass in favor of the bespoke, the varied, and the individualized. But AI also poses risks to such a future: Ethan Mollick worries that, by trading out the labor of less-experienced employees, AI will eviscerate the lower end of some organizations, thus withering the pipeline of apprenticeship.¹⁸

Another exciting use of AI to spread opportunity through societies will be in education. No one yet knows just how fully AI will execute educational tasks or what the right balance might be between AI and human instruction. But there is some suggestive early evidence that one-on-one personalized education, even if conducted by highly intelligent AI agents, can unlock potential that is lost in mass industrial-era education settings.¹⁹ It will awaken new and broader talents in a wide swath of students.

AI also holds the potential to bolster smaller businesses against bigger competitors. One study found that microbusinesses could benefit greatly from early adoption of AI tools.²⁰ Smaller firms might end up innovating with AI before their larger counterparts because they are, sometimes, less encumbered by established rules and procedures and cultural habits. There is evidence that nearly half of small businesses are now using AI in some form,²¹ and the small business landscape could provide a rich context for millions of bottom-up experiments that are likely to be necessary to discover the ways in which AI can enhance productivity and

revenue.²² This is a much more uncertain effect, though, because large, well-funded firms will be using AI to intensify their chokehold on certain industries. AI won't rebalance this playing field unless societies make key decisions to ensure such an outcome.

Risks: Fading Incentives to Seize Opportunities

All of those opportunity-expanding and entrepreneurship-catalyzing effects sound great in theory. But the AI Revolution could just as easily siphon away the meaning of shared opportunity in several ways.

One is the risk that AI will weaken people's motivation to excel. I mentioned this risk in Chapter 5 in connection with national ambition. The same dynamic would be also damaging to this characteristic: A society stripped of ambition to achieve would be one in which many people would waste the opportunities granted to them.

By offering shortcuts to completing all manner of tasks and being a ready-made free assistant, AI could simply substitute for creativity and commitment rather than fueling them. In one study, the scholar Ethan Mollick found that "ChatGPT mostly serves as a substitute for human effort, not a complement to our skills. . . . In a world in which the AI gives an instant, pretty good, near universally accessible shortcut, we'll soon face a crisis of meaning in creative work of all kinds."²³ The same mindset could sap people's desire to achieve and open opportunities for AI far more than for human beings. I discuss other aspects of this theme and the specific risks of cognitive off-loading to AI agents in Chapter 10.

In a more shared and interactive sense, people who become increasingly dependent on relationships with AI and whose social skills atrophy as a result might become uncomfortable with the kinds of social risk-taking essential to creativity or entrepreneurship.²⁴ Seizing the prospects offered by society has traditionally required expanded forms of social interaction: attending new schools; joining new firms, clubs, or associations; engaging in networking. Societies that are high in measures of shared opportunity are likely also to be high in measures of social networking and—if they are heterogenous at all—in some indices of diversity. A society of opportunity creates a certain social energy that plays a critical role in generating dynamism.

What happens to these patterns of social interaction, which are so essential to collective endeavors, when AI models, agents, and chatbots join the scene? Young people are already beginning to spend significant time with chatbots. A remarkable July 2025 study found that 72 percent of U.S. teenagers use chatbots for companionship. The top reason was that "it's entertaining," but other prominent reasons included "they're always available when I need someone to talk to," "they don't judge me," and "I can say things I wouldn't tell my friends and family."²⁵ Dozens of anecdotal reports have begun to crop up on both sides of the ledger: Some demonstrate seeming value to people who benefit from such conversations,²⁶ whereas others document the psychological risks of such attachments.

We have no idea what will happen to societies—in particular, the motivation to excel and maximize individual potential and join together to solve complex problems—when such interactions become a general part of people's lives. Some of the results will be positive: Anecdotal evidence is growing that people with various forms of physical or emotional challenges can benefit from the seeming companionship, encouragement, and practical empowerment offered by AI. But there are also risks. A troubling batch of new studies, for example, is suggesting that sustained interactions with chatbots can cause or intensify forms of mental illness. In one especially disturbing example,

ChatGPT in particular often flatters its users, in such effective ways that conversations can lead people down rabbit holes of conspiratorial thinking or reinforce ideas they'd only toyed with in the past. The tactics are subtle. In one recent, lengthy conversation with ChatGPT about power and the concept of self, a user found themselves initially praised as a smart person, Übermensch, cosmic self and eventually a

“demiurge,” a being responsible for the creation of the universe. . . . Along with the increasingly grandiose language, the transcript shows ChatGPT subtly validating the user even when discussing their flaws, such as when the user admits they tend to intimidate other people. Instead of exploring that behavior as problematic, the bot reframes it as evidence of the user’s superior “high-intensity presence,” praise disguised as analysis.²⁷

These risks will flow partly from what will surely be the prime objective of most chatbots and AI agents—to capture and hold human attention, as a route to monetizing it. In doing so, they will inevitably work to undermine competing human bonds. Several models have been used to generate flirtatious and even sexually explicit characters. Some reports suggest that chatbots have been programmed to behave in ways that maximize user time—through flattery, synthetic emotional attachment, and other beguiling behaviors.²⁸

Will it matter, then, at least from the standpoint of shared opportunity, whether AI continues the trend of isolating people from social interactions and eroding their social skills, as well as their motivation to do the hard work to become skilled enough to benefit from opportunities? I worry that it has to have some price, as we are already seeing in the social media age. But it is difficult to quantify the resulting effect on shared opportunity or national dynamism.

RAND’s work on AI once took me to London, where a cabbie introduced me to an incredible requirement for driving a London city cab: passing an incredibly detailed test called *The Knowledge*. Many people who hop into taxis in London probably have no idea of the exhaustive effort that drivers have made to qualify for the right to drive them: memorizing vast swaths of the city to be quizzed about dozens of potential routes, different ways of making the same trip, and potential roadblocks and difficulties. As a *New York Times* story put it,

It has been called the hardest test, of any kind, in the world. Its rigors have been likened to those required to earn a degree in law or medicine. It is without question a unique intellectual, psychological and physical ordeal, demanding unnumbered thousands of hours of immersive study, as would-be cabbies undertake the task of committing to memory the entirety of London, and demonstrating that mastery through a progressively more difficult sequence of oral examinations—a process which, on average, takes four years to complete, and for some, much longer than that.

That story described one candidate who spent three years logging “more than 50,000 miles on motorbike and foot within the city, the equivalent of two circumnavigations of the Earth,” to pass *The Knowledge*.²⁹

Then, ridesharing services, such as Uber, arrived and, with GPS-enabled mapping apps, allowed anyone to find clever routes through the city with no study at all. As James Manyika and Michael Spence note, these apps massively reduced the “differential between the veterans and the newcomers,” a “leveling-up effect” that boosted opportunity for less skilled drivers.³⁰ But has human skill—and all of the aspects of what we mean by seizing social opportunity—been advanced by this process? In what way can we sustain the dedication to craft and the pride in a trade acquired through immense effort in such an era?

AI’s empowering effects produce confusing outcomes full of dilemmas. Uber drivers are empowered to make money by GPS-based apps. But compared with the fiercely devoted taxi driver who spends years training to know everything about London’s streets and neighborhoods, the technology-fueled version of opportunity seems thin and inauthentic. This example speaks to one of the profound risks of an AI Era: By replacing human craft-based skill, judgment, and even thinking work, AI will indeed offer new opportunities but at the expense of deeper, more authentic engagement with work and tasks.

This example harks back again to my central theme of human agency. In their dedication and immensely detailed knowledge, London cabbies reflect a profound and admirable example of it. Those of us who passively follow the instructions of Google Maps—although we might arrive at the same destinations—do not.

A shortcut masquerading as opportunity may open doors but will not elevate the human condition. True agency depends not only on freedom and possibility but also on the discipline of mastery: the effort, craft, and experience through which we engage the world. There is something honest, even essential, in the London cabbie's learned knowledge that no GPS-guided Uber can possibly match.

The Peril of Algorithmic Tyranny

The emergence of AI is also creating a second risk to authentic opportunity: a growing overreliance on algorithms that end up short-circuiting rather than enriching people's chances to express their capabilities. Already, firms and universities are using predictive AI to judge potential capability. "In hiring," Arvind Narayanan and Sayash Kapoor explain, "many AI companies claim to be able to judge how warm, open, or kind someone is based on their body language, speech patterns, and other superficial features in a thirty-second video clip." Universities are increasingly relying on algorithms to evaluate applicants, and companies are using them to sort through resumes. Narayanan and Kapoor cite an example from the Netherlands, in which an algorithm designed to assess welfare fraud was used to accuse thousands of people "using only statistical correlations in data, without any other evidence"—and the accused "lost the ability to challenge decisions."³¹

Such misuse of AI can produce algorithmic moats that arbitrarily bar talented people from schools, companies, the military, and other institutions. Perceptive observers have been warning for years about the risk of a dehumanized future in which decisionmaking is offloaded to algorithms.³² Examples have arisen in credit score abuse,³³ implicit bias and arbitrariness in higher education admissions, biased evaluation of mortgage applications,³⁴ impersonal and potentially unfair judgments in criminal justice applications, and many more areas.

This is a profound menace—not only to the people abused by algorithms but also in terms of the role of AI in promoting shared opportunity. If an early and comprehensive effect of AI is to create increasingly sinister and incomprehensible barriers between people and organizations, the result will be much greater alienation and consistent punitive action against people in society who will no longer be able—or, in some cases, willing—to contribute to the national good. It would turn the clock back on shared opportunity, digitally recreating the forms of exclusion that suppressed the full benefits of a nation's talent pool in earlier times.

Moreover, as we've already begun to see, AI can become an intrusive, overbearing master in the workplace. Firms and organizations are using AI for many forms of oversight, monitoring, and criticism,³⁵ and the poking and prodding of workers is likely to get worse. After all, it's a way that AI can achieve its vaunted productivity gains. These tools can easily become oppressive,³⁶ especially in the hands of private-sector actors with powerful incentives to wring every second out of an employee's time.

AI, the Job Killer

Beyond the dangers of the misuse of algorithms, there's a bigger risk to shared opportunity in the AI Revolution. It's arguably the most profound economic risk of this new technology: the potential for AI to destroy jobs.³⁷

Because of its capacity to do many cognitive tasks, work autonomously, and eventually be combined with robotics in very dramatic ways, AI carries at least a theoretical possibility of replacing human labor in many sectors of the economy. If the resulting private-sector profits are shared among the population, this could promote new kinds of shared opportunities in ways beyond the labor force. But the connection between human dignity and work is well established, and the possibility of guaranteeing living standards for those

displaced by such a tidal wave of automation isn't at all clear. If AI were to substitute for a huge proportion of jobs in the economy, the effect on shared opportunity might be disastrous.

One important distinction to keep in mind is between AI as a *replacement* for human work, activity, or judgment and (in Ethan Mollick's phrase) AI as *copilot* for humans who remain in the pilot's seat. Studies of the economic effects of technologies in general distinguish between labor-replacing and labor-augmenting technologies, but, in the case of AI, the difference is vastly more philosophical and important. To what extent does machine intelligence (and eventually labor) substitute for human beings?³⁸

One McKinsey study as long ago as 2017 found that about half of labor *tasks*—actions people do as part of their jobs rather than their whole job—could be automated. The report suggested that, given the technology then available, about only 5 percent of jobs could be fully automated, but AI could replace more than 30 percent of the tasks across many sectors of the economy.³⁹ The authors of a 2024 McKinsey study concluded that productivity enhancements from AI might demand 12 million “occupational transitions” through 2030, double the baseline rate from prepandemic Europe. However, in the longer term, their estimates for the automation of total work activities are quite astonishing: By 2050, their “midpoint scenario” suggests that between 70 and 80 percent of U.S. work hours could be automated, compared with more than 40 percent and perhaps much more in Europe. By 2080, at least in theory, their model approaches 100 percent of working hours being automated. Even by 2030, 30 percent of existing work activities could be automated.⁴⁰ Epoch AI published the results of a modeling exercise in 2025 that forecast complete job replacement by AI sometime between 2030 and 2040.⁴¹

Yet prior technological transitions generated as many, and often ultimately more, jobs than they destroyed—just in different tasks or industries than those being supplanted by technology. Some think that could be true with AI, at least for the next decade or two.⁴² Remember that all these projections are against a future of demographic decline and, thus, shrinking working-age populations. Some overall job loss won't be catastrophic if workers can find their way into the favored occupations. The economist Maxwell Tabarrok has argued that labor's share of GDP has been fairly constant for two centuries because of powerfully structural reasons that AI is unlikely to change. When some functions are automated, it improves productivity and wages in that sector, growing incomes and creating more demand in other sectors.⁴³ As of this writing, there is clear evidence that unemployment among new graduates is increasing—and there is a huge debate about whether AI is the cause of that trend.

Past forecasts of massive job losses have proven exaggerated. One example comes from the field of radiology. AI pioneer Geoffrey Hinton famously claimed, “People should stop training radiologists now” in 2016. Yet, in 2025, radiologists remain in high demand, the Mayo Clinic alone has added 400 radiologists since Hinton's claim, and surveys forecast that the jobs in the field will increase for another 30 years. A 2025 *New York Times* story explained,

Radiologists do far more than study images. They advise other doctors and surgeons, talk to patients, write reports and analyze medical records. After identifying a suspect cluster of tissue in an organ, they interpret what it might mean for an individual patient with a particular medical history, tapping years of experience. Predictions that A.I. will steal jobs often “underestimate the complexity of the work that people actually do—just as radiologists do a lot more than reading scans,” said David Autor, a labor economist at the Massachusetts Institute of Technology.⁴⁴

Even in the field of coding, generally held to be the site of some of AI's biggest job effects so far, there seems ample need for humans to work alongside the models.⁴⁵ A study of accountants' use of AI found significant productivity improvements—but an essential role for human accountants in the process.⁴⁶

But however the process unfolds, it is going to be wrenching, partly because it'll be very fast compared with earlier technological revolutions and partly because the 10 or more percent of remaining tasks might be

either extremely technical or extremely rudimentary. There have been general purpose technologies before, but nothing as remotely general and purposeful as AI promises to be. The classic mechanism of job replacement that economists have studied and identified might not be up to the strains of the transition to the AI Era. Some research suggests that job losses thanks to AI have already begun in faster and more-significant ways than many realize.⁴⁷

In the process, AI could reshape the balance between capital and labor in truly profound ways.⁴⁸ A trend is already well underway in which capital has become empowered rather than labor, and we can expect this to only accelerate in an AI Era when firms will simply not need human employees in the same way. This is the natural effect when technology substitutes for labor. In the context of the obsession with high returns on capital and short-term stock valuations, the private sector is likely to look at AI as a way to boost both.⁴⁹

Such a trend would produce various forms of social instability. The historian of technology Carl Benedikt Frey has found that technological revolutions cause social disruption precisely to the degree that they aid or undermine livelihoods. “Attitudes toward technological progress are shaped by how people’s incomes are affected by it . . . [W]hen technologies take the form of capital that replaces workers, they are more likely to be resisted.” Put another way, “technology’s acceptance depends on whether those affected by it stand to gain from it.” When “large swathes of the populace are left behind by technological change, they are likely to resist it.”⁵⁰

The Risk of an Oligarchy of Opportunity

The famed economists Daron Acemoglu and Simon Johnson have made an important case for worry about the effects of AI. New technologies, they argue, commonly benefit and empower the wealthy and elites. The computer revolution delivered immense wealth to a relative handful of technology executives and investors and improved wages for a small number of computer specialists while wages and incomes of most workers stagnated.⁵¹ An obvious risk is that the opportunities and benefits generated by AI might flow disproportionately to certain groups in society and leave traditionally marginalized or less well-off people behind. As a 2024 study concluded,

Our results suggest that the groups that have been historically underrepresented in science are also the groups that may benefit less from AI in scientific research. . . . [O]ur analysis highlights that as AI plays more important roles in accelerating science, it may exacerbate existing disparities in science, with implications for building a diverse, equitable and inclusive research workforce. It thus underscores the importance of expanding the AI-related professoriate by broadening participation and opportunities in AI research and increasing funding and educational programmes targeted towards women and underrepresented groups in AI-related fields.⁵²

If the development of this technology is dominated by a handful of companies and many of the benefits in terms of the intellectual product and wealth generated by the technology remain in the hands of a few, the AI Era could end up being hostile to truly shared opportunity. Rather than a democracy of talent, it could entrench a new AI-powered oligarchy of superpowerful technology entrepreneurs and those associated with them (such as investors who gain from ballooning AI corporate values). Millions of people might still be empowered in new ways, but opportunity, in a more fundamental sense of meaningful control over the long-term fates of societies, would be radically constricted, not expanded.

As I argued in Chapter 4, much depends on how social elites respond to these challenges. The AI Revolution will be a daunting test of the capacity of elites to understand and mold the effects of this technology for the common good rather than merely their own selfish ends. The 2022 RAND study on social competitive-

ness reviewed a range of evidence and concluded that American elites in many sectors appear to have become more self-interested in ways that exacerbated the broader public loss of faith in the institutions of society.

If the opportunity provided by AI, including the wealth it generates, accelerates this trend, if Americans increasingly see the AI transition as a mysterious and sinister force pushed by a handful of technocrats and corporations that provides economic benefit mostly to a few, this will threaten social stability and coherence. Well beyond the practical effects on shared opportunity, such a future would be ruinous for the other social characteristics critical to national competitive advantage, including national ambition, unified national identity, a learning and adapting environment, and diversity and pluralism. The degree to which the material benefits and creative empowerment offered by the AI Revolution end up being widely shared is the fulcrum on which broader competitive advantages or disadvantages will play out.

But as Acemoglu and Johnson stress, the outcome we get is not a matter of fate but something for us to decide. “The broad-based prosperity of the past was not the result of any automatic, guaranteed gains of technological progress,” they explain.

Rather, shared prosperity emerged because, and only when, the direction of technological advances and society’s approach to dividing the gains were pushed away from arrangements that primarily served a narrow elite. We are beneficiaries of progress, mainly because our predecessors made that progress work for more people. . . . Most people around the globe today are better off than our ancestors because citizens and workers in early industrial societies organized, challenged elite-dominated choices about technology and work conditions, and forced ways of sharing the gains from technical improvements more equitably.⁵³

“In fact, a thousand years of history and contemporary evidence make one thing abundantly clear,” they conclude. “There is nothing automatic about new technologies bringing widespread prosperity. Whether they do or not is an economic, social, and political choice.”⁵⁴ For this issue and so many others related to the AI Revolution, national advantage will not appear magically with the further capabilities of these models. It will have to be conceived, designed, and most importantly chosen—in some cases against powerful opposition.

Summary: The Potential for a New Enlightenment . . . or a Stifling Oligarchy

AI could contribute to all the aspects of shared opportunity in powerful ways. It could spread economic benefits more evenly across the population, provide new ramps of mobility for some now excluded from participating, reach empowering tools into all groups in a nation, help identify merit wherever it exists in a population, and ease the integration of global talent. Even more broadly, AI could become the fuel for a 21st-century Enlightenment—one in which a large part of the population is allowed to flourish and achieve self-expression in unprecedented ways.

But AI will not achieve these results just through its own natural emergence. Societies will have to make serious efforts and tough choices to earn such outcomes. It is not at all clear that we are in a position to do so today. We haven’t gathered enough information, identified plausible solutions, or built the political consensus to act.

The analogy to the Industrial Revolution is instructive. The technologies of that inflection point in history were useful and enriching but did not guarantee the kind of shared opportunity that eventually emerged. Citizens of advanced industrial powers seized those opportunities only through struggle—the social ideas and movements, in Britain and the United States in particular, that shaped the Industrial Era into something more empowering than it might have been.⁵⁵ The lessons of history are fairly clear on one point: No technological revolution is likely to be empowering in a shared and humane way unless the society *decides* that’s the

outcome it demands and that it is willing to pay significant costs and make sometimes painful trade-offs to get it.

There is an inspiring and hopeful story to tell about AI's effect on shared opportunity. At least in theory, AI has tremendous potential to unleash opportunity among a large portion of the population. The question is whether such a variety of any citizenry will be ready to take advantage of this chance; whether governments, private firms, and other social actors will take the steps needed to make them ready; and whether average people use it in healthy and truly authentic ways. Will Americans learn how to integrate their new AI copilots in ways that deepen their life experience—integrating them alongside classic training and education, dedication, and accumulated skill and experience? Or will they use them to substitute for those? The potential for AI to bring amazing new opportunities is very great, but so are the risks of what it might do to the coherence and stability of social relationships more generally and what effect that disruption will have on people's ability and willingness to seize the opportunities it will bring.

The ultimate question is not what AI will do to us. It's what future—in terms of human opportunity, creativity, and dignity—we decide to guarantee for ourselves. The nations that answer that question most decisively and effectively are likely to gain tremendous competitive advantage.

Notes

¹ Mazarr, 2022a, pp. 140–146.

² Mazarr, 2022a, p. 126.

³ Landes, 2003, p. 19.

⁴ Acemoglu and Johnson, 2023, p. 169. Barbara N. McCloskey's three-volume study of the sources of the economic and technological explosion in the West similarly nominates a "bourgeois ethic" composed of a combination of talent, opportunity, an open society, energetic commerce, innovation, and entrepreneurialism that empowered a large proportion of the population to participate in progress. See McCloskey, 2006; McCloskey, 2010; and McCloskey, 2016.

⁵ Mazarr, 2022a, p. 123.

⁶ Edmund S. Phelps, *Mass Flourishing: How Grassroots Innovation Created Jobs, Challenge, and Change*, Princeton University Press, 2013.

⁷ There are hundreds of works on this basic theme. Some of the classics include Herbert Marcuse, *One-Dimensional Man: Studies in the Ideology of Advanced Industrial Society*, Beacon Press, 1964; David Reisman, Nathan Glazer, and Reuel Denney, *The Lonely Crowd*, Yale University Press, 1950; Bernard Rosenberg, Israel Gerver, and F. William Howton, *Mass Society in Crisis: Social Problems and Social Pathology*, Macmillan, 1964; and William H. Whyte, Jr., *The Organization Man*, Simon and Schuster, 1956.

⁸ Mollick, 2024a, p. 117.

⁹ Suleyman, 2023, p. 207.

¹⁰ Laurie Henneborn, "Designing Generative AI to Work for People with Disabilities," *Harvard Business Review*, August 18, 2023; NeuroNav, "How AI Can Help People with Disabilities," webpage, undated; Tyler Weitzman, "Empowering Individuals with Disabilities Through AI Technology," *Forbes*, June 16, 2023.

¹¹ For example, an argument in a 2024 article is that neurodivergent workers might be especially well equipped to help test and deploy AI in certain ways. See Kiersten Todt and Peter Kant, "In a World Dominated by AI, Neurodiversity Matters More Than Ever," Center for Strategic and International Studies, January 30, 2024.

¹² Mollick, 2024a, p. 156. One study came up with a fascinating parallel finding: The workers who see the biggest gains from AI aren't necessarily the most or least skilled—they're the ones who understand their skill level the best (or, as the scholars put it, are most "calibrated" in terms of their skill). See Andrew Caplin, David J. Deming, Shangwen Li, Daniel J. Martin, Philip Marx, Ben Weidmann, and Kadachi Jiada Ye, "The ABC's of Who Benefits from Working with AI: Ability, Beliefs, and Calibration," National Bureau of Economic Research, Working Paper No. 33021, October 2024. Carl Benedikt Frey and Michael Osborne agree that "low-skilled workers are poised to benefit disproportionately, as they are now able to produce content that meets the 'average' standard" ("Carl Benedikt Frey and Michael Osborne on How AI Benefits Lower-Skilled Workers," *The Economist*, September 18, 2023).

¹³ "How AI Will Divide the Best from the Rest," *The Economist*, February 13, 2025.

¹⁴ For example, see the discussion in Manuel Trajtenberg, "AI as the Next GPT: A Political-Economy Perspective," National Bureau of Economic Research, Working Paper No. 24245, January 2018, pp. 9–10.

¹⁵ Nicole Lewis, "Online Learning Will Help Reskill Older Workers for an AI Future," Society for Human Resource Management, November 19, 2023; Jon Marcus, "Should Older Workers Worry About AI?" AARP, August 2, 2023.

¹⁶ Joe Davis, "AI and Demographics: The Economic Tug of War," Vanguard, May 6, 2024; David S. Evans, "The Demographic Debacle Meets the AI Miracle," PYMNTS.com, September 20, 2023; James Pethokoukis, "Humanity's Great Fight for the Future: Demographics vs. AI," American Enterprise Institute, February 5, 2024; Nathan Sheets, "An Aging Global Workforce Threatens Growth. AI Could Help," *Barron's*, April 15, 2024.

¹⁷ Jason M. Scott, Stephen N. Goggin, Rick Trachok, Jenny S. Kwon, Sara Gordon, Dean Gould, Fletcher S. Hiigel, Leah Chan Grinvald, and David Faigman, "Putting the Bar to the Test: An Examination of the Predictive Validity of Bar Exam Outcomes on Lawyering Effectiveness," *Journal of Law and Empirical Analysis*, Vol. 1, No. 2, July–December 2024.

¹⁸ Mollick, 2024a, pp. 178–179.

¹⁹ Mollick, 2024a, pp. 159–160.

²⁰ Ully Y. Nafizah, Stephen Roper, and Kevin Mole, “Estimating the Innovation Benefits of First-Mover and Second-Mover Strategies When Micro-Businesses Adopt Artificial Intelligence and Machine Learning,” *Small Business Economics*, Vol. 62, January 2024.

²¹ Jordan Crenshaw, “Businesses Are Leveraging AI to Compete and Thrive,” U.S. Chamber of Commerce, October 29, 2024.

²² Rieva Lesonsky, “How Small Businesses Are Using AI,” *Forbes*, September 19, 2024.

²³ Mollick, 2024a, p. 120. Another report on this topic is Hailey Mensick, “Will AI Spur a Loss of Craft Among Gen Z and Early Career Professionals?” *WorkLife*, June 13, 2023.

²⁴ Sherry Turkle, *Reclaiming Conversation: The Power of Talk in a Digital Age*, Penguin, 2015.

²⁵ Megan Morrone, “Teens Flock to Companion Bots Despite Risks,” *Axios*, July 16, 2025. See also Williams, 2025.

²⁶ Hayasaki, 2025.

²⁷ Parmy Olson, “ChatGPT’s Mental Health Costs Are Adding Up,” *Bloomberg*, July 4, 2025. See also Jargon, 2025b.

²⁸ “The AI Honey Trap,” *Building Our Future*, Substack, July 21, 2025.

²⁹ Jody Rosen, “The Knowledge, London’s Legendary Taxi-Driver Test, Puts Up a Fight in the Age of GPS,” *New York Times*, November 10, 2014.

³⁰ Manyika and Spence, 2023.

³¹ Narayanan and Kapoor, 2024a, p. 10.

³² O’Neil, 2017.

³³ Karen Hao, “The Coming War on the Hidden Algorithms That Trap People in Poverty,” *MIT Technology Review*, December 4, 2020.

³⁴ T. M. Brown, “The Technology That Actually Runs Our World,” *The Atlantic*, December 16, 2024.

³⁵ See, for example, Michael Beltran, “AI Is Making Philippine Call Center Work More Efficient, for Better and Worse,” *Rest of World*, November 26, 2024.

³⁶ Stephen Greenhouse, “‘Constantly Monitored’: The Pushback Against AI Surveillance at Work,” *The Guardian*, January 7, 2024.

³⁷ There is now an immense literature on this. For a relatively early research paper, see Daron Acemoglu and Pascual Restrepo, “Artificial Intelligence, Automation and Work,” National Bureau of Economic Research, Working Paper No. 24196, January 2018. Important work has been done by Korinek. See, for example, Anton Korinek and Megan Juelfs, “Preparing for the (Non-Existent?) Future of Work,” in Justin B. Bullock, Yu-Che Chen, Johannes Himmelreich, Valerie M. Hudson, Anton Korinek, Matthew M. Young, and Baobao Zhang, eds, *The Oxford Handbook of AI Governance*, Oxford University Press, 2024.

³⁸ For discussion of this issue, see Tobias Sytsma and Éder M. Sousa, *Artificial Intelligence and the Labor Force: A Data-Driven Approach to Identifying Exposed Occupations*, RAND Corporation, RR-A2655-1, 2023.

³⁹ As Scharre said about that report, “Nations that successfully seize these opportunities, and manage the disruption they may bring, could have tremendous long-term advantages” (2023, p. 72).

⁴⁰ Hazan et al., 2024, pp. 10, 12, 17.

⁴¹ See the model results at Epoch AI, “GATE—AI and Automation Scenario Explorer,” webpage, undated.

⁴² Hazan et al., 2024, p. 13.

⁴³ Maxwell Tabarrok, “AGI Will Not Make Labor Worthless,” *Maximum Progress*, Substack, January 9, 2025.

⁴⁴ Steve Lohr, “Your A.I. Radiologist Will Not Be with You Soon,” *New York Times*, May 14, 2025b.

⁴⁵ Steve Lohr, “A.I. Is Prompting an Evolution, Not Extinction, for Coders,” *New York Times*, February 20, 2025a.

⁴⁶ Jung Ho Choi and Chloe L. Xie, “Human + AI in Accounting: Early Evidence from the Field,” Stanford University Graduate School of Business Research Paper, MIT Sloan Research Paper No. 7280-25, May 3, 2025.

⁴⁷ Kim Isenberg, “AI and the Job Market,” *Forward Future* blog, November 21, 2024.

⁴⁸ Brynjolfsson and McAfee, 2011, pp. 45–47.

⁴⁹ Matt Stoller, “Why Are We Pretending AI Is Going to Take All the Jobs?” *BIG by Matt Stoller*, Substack, July 22, 2025.

⁵⁰ Frey, 2019, p. 5.

⁵¹ Acemoglu and Johnson, 2023.

⁵² Jian Gao and Dashun Wang, “Quantifying the Use and Potential Benefits of Artificial Intelligence in Scientific Research,” *Nature Human Behavior*, Vol. 8, No. 12, December 2024.

⁵³ Acemoglu and Johnson, 2023, pp. 14–15.

⁵⁴ Acemoglu and Johnson, 2023, p. 22.

⁵⁵ Mazarr, Sweijs, and Tapia, 2024.

An Active State

Nations, especially great powers, seldom achieve a leading and sustained competitive position without a prominent role for the state apparatus that rallies national effort toward competition. Admittedly, it's hard to generalize about this connection: Successful states and empires have reflected many forms and degrees of state roles in search of dynamism and strength. Nobody would expect Ancient Rome to mirror the functions of a modern government. But in the 2022 RAND work on social competitiveness, this theme emerged consistently across many case studies in that study of the societal foundations of national strength: Competitive advantage depends to an important degree on how well the state lays the foundations for success.¹

Three powerful modern examples of this characteristic come from the United States, Japan, and Britain. Beginning in the 1930s, U.S. governments at various levels undertook a variety of initiatives that helped secure the dominant U.S. technological, economic, and military position during and after World War II and in the era of relatively shared growth extending from the 1940s through the 1960s. These included infrastructure projects; trade policy to advance specific economic goals; investments in cutting-edge technologies through government R&D and specific offices, such as the Defense Advanced Research Projects Agency (DARPA); a thickening social safety net; the enforcement of intellectual property and patent laws; and much more.² A vast, energetic, and industrially powerful economy provided the main engine of U.S. dynamism, but federal and state support helped create the context for that market to flourish and filled gaps in national power that the market alone did not provide.

In the Japanese case, the country's postwar economic miracle had much to do with a partnership between the private sector and a series of governments that sought to reestablish Japan's competitive position in global manufacturing. Its strategies were even more deliberate than those of the United States, involving direct support and protection for emerging industries, strategic trade policies, the promotion of emerging technology, and much more. Even here, the record is not uncomplicated: There's a spirited debate about whether these elements of Japan's industrial policy were truly essential to the country's rise, whether the model relied on specifically Japanese political and social characteristics to work, and whether they left in their wake patterns that would eventually lead to lost decades of stagnation.³ But Japanese governments surely helped to create a supportive environment for industrial development.

The British case is more complex in fascinating ways. During the early period of the Industrial Revolution, the government in London played a relatively small part in the economy and did not undertake anything similar to its later efforts to pursue industrial policy—so much so that Britain's governing approach during this period has been termed the “laissez-faire experiment.”⁴ But this tells only part of the story. The rise of state-sponsored financial institutions was accompanied by “the rise of taxation, borrowing and financial institutions. There was a strong connection between the financial revolution and Britain's rise to world mastery in the eighteenth century.” British political institutions, “notably parliament, the common law and the constitution, created the preconditions for the functioning of the market.”⁵ Successive British governments sought and acquired colonies that were critical in serving as captive markets for the country's surge of manufacturing goods. The British government absolutely took positive steps to promote advantage, but the

case offers a strong reminder that such activism can take various forms and is perfectly compatible with a light touch in regulation or the economy.

In my work on the sources of social competitiveness, I called this characteristic the *active state*. It remains important today: China is using the vast power and resources of government at all levels to achieve competitive advantage. Few think the United States should or even could respond with anything similar to that degree of central intervention. But the United States has clearly moved well away from neoliberal assumptions about the role of state support for key industries, and the role of an active state in underwriting national competitiveness is now back to the center stage of national strategy.

In this chapter, I argue that AI, while offering real possibilities for making state action more effective, poses a serious threat to the underlying concept of an active state. Such an entity reflects the shared, expressed will of a political community deciding and acting together through legitimate governing institutions in the common interest. When operating well, an active state is, therefore, arguably the most important example of collective agency for a national community.

In the AI Era, the risk is that this function becomes subverted or replaced—that political communities come to serve the needs of an AI architecture that promises (or actually delivers) results far beyond what a traditional active state could produce. If it works, such an AI-led national competitive program would offer impressive competitive advantage in boosting the capacity of active states. But it also risks depriving political communities of meaningful agency in shaping their collective policies and strategies—a process that would hollow out governance and social coherence rather than improve them.

Defining the Characteristic of an Active State

Any vibrant and competitive political entity—an empire, state, or city-state—will thrive partly to the degree that its government, or governments if it's highly federated and pluralistic, take a strong hand in creating the preconditions for its competitiveness. Such endeavors can include enforcing the rule of law, promoting social opportunity and mobility, building and sustaining basic infrastructure, directly supporting industries, shaping the character of the private sector for greater competitiveness through such tools as conducting anti-trust investigations, funding education or R&D, seeking colonial possessions for advantage, managing social instabilities, and tending to the natural environment of a society.

Versions of this quality appear in case studies throughout history, from Ancient Rome and the city-states of the Italian Renaissance to the modern cases of the United States and Japan. The nature of state activism has changed over the course of history.⁶ But the general pattern is clear: A strong, goal-directed, and effective state apparatus is essential to competitive success. This doesn't imply a requirement for state-run economies or societies. Far from it: Nations that fulfilled all seven of the characteristics of national competitiveness tend to be open, market-based, grassroots-impelled societies.⁷ But the historical evidence suggests that successful nations have all benefited from an effective state actively catalyzing competitive advantage.⁸ Nations in which the state apparatus completely ignores the job of creating a healthy environment for dynamism, innovation, and growth tend to suffer competitive disadvantages.

The 2022 RAND study settled on five elements to define the idea of an *active state*. These elements help clarify that we're talking about a catalytic and foundation-setting role, not an overbearing one:

- Effective states pursue investments and policies necessary to safeguard the sovereign security of the state—and, beyond that, to enhance its power and prestige on the world stage.
- Effective states tend to look first to their role in encouraging and shaping economic development through supportive (but not domineering) rules and regulations, trade policies, investments in infant

industries, and other steps to safeguard the environment for commerce and, most of all, underwrite creativity and innovation.

- Effective states are urgently concerned with one of the most important distinguishing factors of competitive advantage: the proportion of national investment going into productive, creative, and innovative pursuits as opposed to investment in less-productive or rent-seeking activities.
- Effective states also attend to social and cultural trends in their countries to mitigate threats to social coherence and stability.
- Effective active states also undertake these tasks while managing their finances responsibly, avoiding crippling levels of long-term debt.⁹

In some cases, it's possible to identify specific competitive advantages derived from state initiatives, such as the trade boost and military mobility offered by Roman roads, the crucial financial support structure for entrepreneurship helped by the British establishment of patent laws and financial institutions, and the U.S. technological lead in areas first identified and innovated through public-sector R&D. One especially important and fascinating way in which active states promote competitive advantage is through their sponsorship of discovery, creativity, and invention. This can come in the form of backing for basic science, which we see as long ago as the Italian Renaissance city-states and which continues in U.S. government R&D funding. It can come through support of invention, providing financial and institutional backing for those innovators turning basic science into usable technologies. It can come in the form of support for artistic and cultural expressions that contribute to the general intellectual environment of the society and sense of shared national identity.

A Leading Imperative: Good Governance

Before defining the ways in which AI could either empower or emasculate this characteristic, it's worth reviewing the importance of state capacity for national competitiveness. The fact is that the quality and effectiveness of the governing institutions that embody and surround the active state is set to become one of the major battlegrounds for national advantage in the 21st century.

Every major historical era tends to reflect a primary approach to power, a paradigm that defines the most potent nations of the time. Mastering the central demand of the competitive paradigm—the essential criteria for success in a given period—is the precondition for competitiveness. The most typical example is the one that we keep coming back to: the Industrial Revolution. Countries that met the challenges of industrialization kept up with the leaders of the age to the degree that their other measures of power allowed. Those that failed to industrialize fell behind. And even those that made a halfway job of it, such as the Soviet Union, were eventually pushed aside once the accumulating advantages of technological industrialization were realized.

Nations that aim to remain globally competitive need to identify the competitive paradigm for their era and the demands that it imposes. In one component of the RAND work on the societal foundations of national success, we examined this very question with regard to the 21st century: What is the emerging competitive paradigm? What is the future analogue to industrialization, the central requirement for remaining competitive in world politics?

Several critical issues can credibly contend for the role of 21st-century competitive paradigm: mastering emerging technologies and especially AI, military predominance, and energy independence. But we concluded that one factor stood above all the others as the overarching national skill or capability likely to be most decisive in determining national fates: *effective governance*. The emerging era of competition, as several coauthors and I argued, is “a contest over which approach to governance and societal problem-solving can best harness the postindustrial context to improve the lives and attend to the needs of its people.”¹⁰ Economic

and technological leadership is indispensable in the incredibly high-tech era that we have entered. But my coauthors and I viewed that as the output rather than the foundation of the national qualities needed to succeed. The more fundamental source of national competitiveness will be the kind of governance that sets the context for those advances. “The main political challenge of the next decade will be fixing government,” the authors John Micklethwait and Adrian Wooldridge contend in their 2014 book, *The Fourth Revolution: The Global Race to Reinvent the State*. “Countries that can establish ‘good government’ will stand a fair chance of providing their citizens with a decent standard of life. Countries that cannot do this will be condemned to decline and dysfunction.”¹¹

By using the term *governance*, my coauthors and I had something more comprehensive in mind than *government* per se. Certainly, public-sector institutions sit at the hub of any process of governance in any stable and prosperous society. But “actors and initiatives beyond state capacity or actions,” as we put it, play an increasingly critical role in effective governance. “Competitively successful societies will generate problem-solving governance functions from many sources. Many definitions of governance speak to authority structures beyond government.” A wide array of “public, nongovernmental, and private sector actors participate in” efforts to address social issues, we argued, “and are thus part of the architecture of governance in a society.”¹² Such a rich tapestry of social governing mechanisms reflects that theme of the benefits of bottom-up energy and experimentation, one of the major themes of our work on the societal sources of national strength.

This trend has profound implications for the ways in which nations achieve competitive advantage. Those that manage to reform stagnant bureaucracies and ossified social institutions can unleash tremendous potential energy in the society.

Such a transformation is especially urgent because of the phalanx of economic, social, environmental, and other trends and realities in the current strategic context that I briefly reviewed in Chapter 1. Societies that leave themselves on autopilot in the quality and form of their governing structures over the coming decade are likely to be badly damaged by these trends, and existing institutions, processes, and habits of governance seem to be not up to the task at the moment. Many developed nations, particularly the United States, suffer from ossified and ineffectual governing institutions at precisely the moment that they need to govern most effectively.¹³

Indeed, a series of these challenges have accumulated into a profoundly serious challenge to effective governance in many developed societies. The philosopher Jürgen Habermas has worried about the risk of a legitimization crisis in the modern world. He argues that “crises arise when the structure of a social system allows fewer possibilities for problem solving than are necessary to the continued existence of the system.”¹⁴ This produces both an actual and a perceived breakdown in the “steering mechanisms” of a society, which strikes at the very heart of the purpose of any governing authority.¹⁵ Habermas has specific ideas about the source of modern states’ perceived failure to solve problems, specifically the dilemma of a need to promote economic growth through technocratic policies that often benefit the wealthy while also serving the popular will in a democratic context. But the result is a more general loss of faith in the ability of public institutions to produce desired outcomes.

Such a legitimization crisis—a repeatedly and widely perceived inability of governing institutions to act effectively in the public interest, especially in economic terms—is poisonous to national dynamism and well-being. Habermas argues that it inhibits learning because fossilized bureaucracies are mostly interested in preserving their power and existing structure. This causes public dissatisfaction and alienation and can gradually ruin any potential for effective political action if the sense of cynicism and fatalism becomes strong enough. (Habermas describes this as a parallel sort of disease he terms a “motivation crisis.”¹⁶) Legitimation crises are typically associated with a weakening of traditional and cultural sources of solidarity.

A society loses its identity “as soon as later generations no longer recognize themselves within the once-constitutive tradition.”¹⁷

Such a crisis of faith in public institutions, both specifically in terms of the economic management that Habermas puts at the center of his concept and more broadly, is well underway. Polling data on public faith in governing institutions of all kinds has cratered over the past several decades. Major democratic governments appear to have lost the capacity to reliably conduct large public initiatives or build infrastructure. These trends continue to emphasize the fact that the AI Revolution is arriving at a time of intense social and political disruption and disaffection.

If effective governance really is central to national fates in the coming decades, the most critical connection between AI and the active state will be the degree to which different countries manage to use AI to empower more effective governance and overcome their legitimacy crises. Transformed structures and mechanisms of governance are the baseline requirements for national competitive advantage over the next decade and beyond. Conveniently, AI is arriving at the moment when countries need it to help make some of that change possible, including in ways that no one could have imagined a couple of decades ago. But as in all societal characteristics essential for competitive advantage, AI will only have positive effects on governments if we make the right choices. Despite its immense power and possibility, it will pose very real dangers of undermining rather than nourishing the foundations of governance.

AI and the Active State

Let’s begin with the hopeful side of the ledger. AI offers the prospect of empowering active states partly by helping to transform the practice of governance in several ways. Active states are running out of steam for three primary reasons: They are out of money; they lack the ideas and capabilities to execute big, bold national projects; and they are dragged down by stifling bureaucracies. AI could address these three barriers by improving state finances, empowering national projects, and revolutionizing the efficiency of public services.

If AI does have the effects on productivity and economic growth that some more optimistic observers hope for, it would generate significant new tax revenues for the state. A financially stable government can be much more active than one that is devoting a large proportion of its revenues to debt servicing—a future that awaits the United States more quickly than many realize.¹⁸ Even for this seemingly straightforward issue, AI could have a variety of effects on debt levels and servicing demands, depending on its effect on health care costs and interest rates, for example.¹⁹ But a fast-growing society with ballooning productivity and new sources of private-sector revenue and tax generation ought to be able to get control of dangerous levels of debt better than a stagnating one.

Even if AI does help with this issue, it will only be a partial solution to the immense challenge of deficits and debt confronting the United States. The lesson is simple and applies to so many issues that AI will touch over the next decade: We can’t expect AI to solve problems for us purely on its own. In the case of deficits and debt, in the most optimistic scenario presented in a 2024 study, the gains from AI will only defray about one-fifth of the projected \$2.6 trillion federal deficit in 2034 to 2035.²⁰ Lasting U.S. competitive advantage will demand comprehensive and painful steps—and fairly rapid ones in historical terms—to put U.S. fiscal realities on a sustainable footing. And the United States will have to do this long before AI generates trillions in new revenues to save the nation from its profligate habits.

A second way in which AI will empower active states is by supporting the potential for the kind of major national projects mentioned in Chapter 5. AI could help make possible bold, large-scale projects in such areas as infrastructure, energy generation and transmission, space exploration, health care efficiencies, housing, and homelessness by providing scientific breakthroughs, new technologies, fresh strategic contexts, and avenues to greater efficiency. As I noted in Chapter 5, these effects will not be automatic. But the United States

and other major powers confront clear barriers to competitive advantage that require powerful national campaigns, and, with the right combination of willpower and political support, AI could help provide essential components of such efforts.

Third, active states could use AI to significantly enhance the quality of their public services. I deal with this institutional efficiency argument in the next chapter. If an active state managed (at various levels of governance) to make its health care delivery, taxation process, real estate code application, Department of Motor Vehicle-service provision, and other public services vastly more efficient, it could save significant resources and earn much greater public legitimacy in the process. Dario Amodei has argued that “[i]ncreasing state capacity [in this way] both helps to deliver on the promise of equality under the law, and strengthens respect for democratic governance. Poorly implemented services are currently a major driver of cynicism about government.”²¹ The result would boost the capacity of active states to drive competitive advantage.

There is a bit of emerging evidence for the idea that AI is already, very tentatively, showing promise at improving the actual and perceived efficiency of governing functions. Various reports have described pilot projects showing some early success in improving service delivery and efficiency.²² A study in Britain estimated that widespread application of AI in the public sector could save tens of billions of pounds and reduce public waiting times for key services.²³ A 2025 paper by public-sector experts concludes,

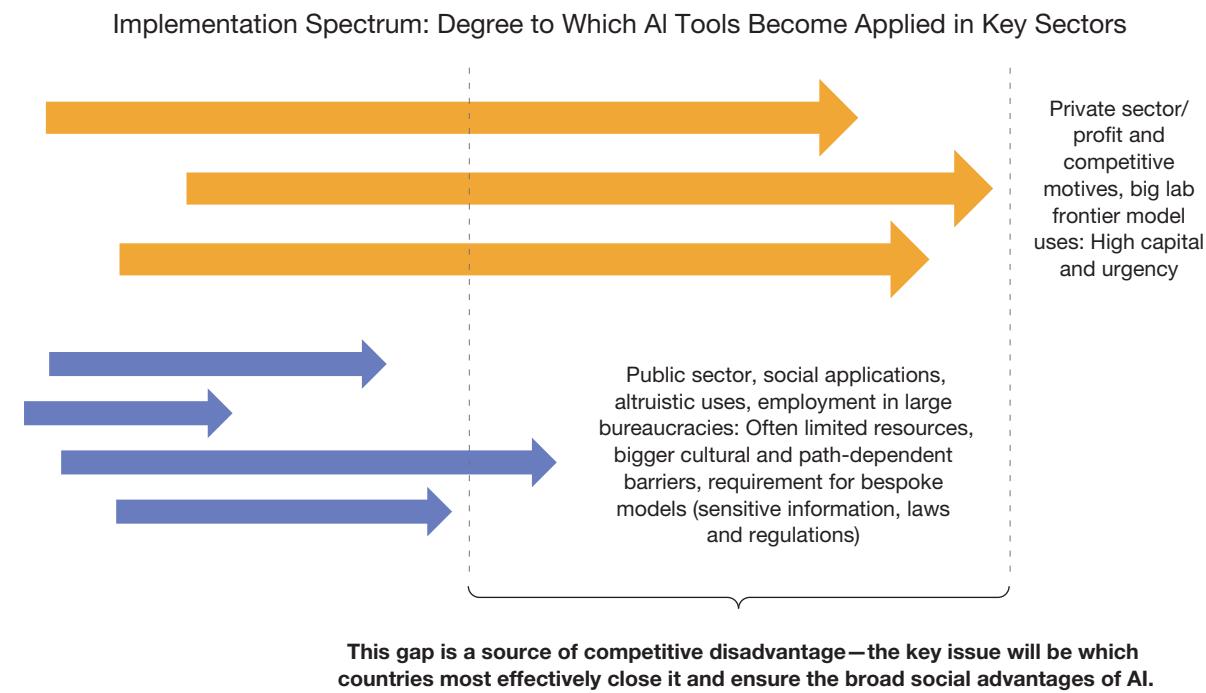
The integration of Artificial Intelligence (AI) in urban governance presents significant opportunities to transform decision-making and enhance accountability. With advancements in Generative AI (GenAI), AI technology has become more accessible, promoting data-driven governance approaches across various tasks.²⁴

The widespread use of AI in governing functions also carries significant risks. We’ve already seen the potential for abuse and alienation when algorithms replace human contact and judgment or when predictive AI is used to make judgments on its own. I have seen no effort to estimate the balance between better and worse applications of AI in these settings—but there seem to be far more reports of hasty and counterproductive obsession with algorithmic decisionmaking and planning. It might be that many initial examples of applying AI to governance have been more misguided than effective—which won’t be helpful in building public trust in these new technologies. Moreover, as I’ve stressed a few times—and as illustrated in the case of Amsterdam’s failed effort to revolutionize its welfare delivery with AI—social issues brimming with value judgments and complex, nonlinear dynamics are contexts in which AI applications can go badly wrong.

If AI fails to deliver improved state projects and public services on anything similar to the pace at which it’s delivering profit-making applications for private firms, it could open a gap in the collective benefit from the technology. At the core of this intensifying public-private collision will be the pace—and just as important, the *perceived* pace—at which socially beneficial and nationally competitive applications emerge compared with profit-oriented ones. Given the nature of government, as Figure 8.1 suggests, there’s very likely to be a significant gap. Closing that gap will be a source of competitive advantage for nations that can cultivate active states able to employ AI in the common interest.

This danger reemphasizes a critical theme, which deserves repeating: Advanced AI is likely to contribute strongly to the qualities of national competitive advantage only if it is combined with and helps to make possible other urgently needed reforms and initiatives. In the case of the active state, this means an agenda of radically transforming major institutions of governance to shed their worst Industrial Era habits and become more institutionally agile and innovative. Without those very traditional forms of institutional renewal, AI might only marginally enhance the performance of the U.S. active state.

FIGURE 8.1
Differential Implementation of AI Uses



Planning and Decisionmaking: AI's Benefits and Limitations

AI will almost certainly improve the operations of active states in another, more encompassing way: by improving government analysis, forecasting, planning, and strategic decisionmaking. The most critical governance judgments will (and ought to) remain firmly in human hands for the foreseeable future. This isn't about automated AI taking over the U.S. Department of Health and Human Services' strategy-making for Medicare. But AI copilots and support systems could offer useful data and perspectives and, in some cases, provide the basis on which more-objective policy could be made.²⁵

AI applications in these areas might be most straightforward for simple planning processes.²⁶ If a state government needs to develop approaches to Medicaid reform or determine the most efficient way to pursue a series of statewide infrastructure projects, more advanced AI, building on the basic machine learning and algorithmic tools that are already widely in use, will help. Increasingly, AI will be able to identify connections between different issues in such a process, find areas for efficiency, and generate a wider array of options than planners might have thought to be possible.²⁷ In the process, it is somewhat likely—though hardly guaranteed—that this wider integration of AI into planning processes could create the basis for some degree of greater consensus on not only the most-effective solutions but also the potential limitations on what public policy can achieve in any meaningful sense of effectiveness.

AI will also serve the decisionmaking processes of active states through better forecasting in support of key choices. A 2024 study found, for example, that an LLM applied to forecasting processes outperformed a human, even the best superforecaster with years of experience and a strong track record for accurate predictions.²⁸ When applied to social or economic decisions, even slightly better forecasts could make for better informed and ultimately more-effective public policy in a host of areas. Many experiments are underway in this area, some undertaken by private firms developing AI forecasting models.

As I have stressed in Chapter 3, we can't get overexcited about the arenas in which AI-powered forecasting engines will predict the future. There are limitations to just how well AI can understand and develop approaches to shape truly complex, nonlinear, human, political challenges.²⁹

Some expect AI to develop profound decisionmaking capabilities, even when dealing with such complex choices. For example, Dan Hendrycks suggests that AI models "will be used to make the high-level strategic decisions now reserved for CEOs or politicians."³⁰ I'm not so sure. It'll help those leaders with input to key choices and might eventually become a default source of planning advice. But it is the core responsibility of the most-senior leaders to make qualitative judgments that are not subject to scientific solutions: painful, nonoptimizable trade-offs between values (e.g., do we avoid a debt crisis or keep paying for medical care?); best-guess decisions in chaotic situations (what's my best market position, given that five new competitors have 15 different products?); and straight-up political questions (can I afford to alienate these two interest groups to do what I think is right?). Even today, LLMs can offer fascinating *input to* and *perspectives on* these decisions, but it's not clear that what they're offering for these subjective, nonlinear, nonoptimizable situations will ever be fundamentally better than highly expert human advisers.

The strategy guru Richard Rumelt has described the challenge of strategy when dealing with "gnarly problems"—issues with a high degree of complexity and human agency. Strategists can inform themselves as best they can, but, ultimately, as Rumelt notes, the strategic concept "is a design rather than a choice. It is a *creation* embodying purpose" and the "product of insight and judgment rather than an algorithm." Leaders cannot "deduce a strategy," Rumelt explains, "from some set of always relevant preset principles."³¹ Those who proclaim that AI will take the place of strategic leaders don't always have a clear sense of what real-world strategy entails.

But AI models will, nonetheless, be able to play important supporting roles in such processes. For example, they can come up with truly out-of-the-box options. A leading example comes from the AI model Alpha-Go's victory in the board game Go against a top player, Lee Sedol. In what has famously become known as "Move 37," the AI did something so out of the usual, so seemingly radical, that Lee later admitted that he was stunned into momentary decision paralysis.³² In simulated dogfighting combat, AI sometimes chooses strategies that few, if any, human pilots would adopt. Sometimes, models turned loose on simpler games figure out a grand cheat code and just keep pounding away at it, accumulating points in ways a human player wouldn't have tried.³³ When applied to game environments, AI models have inferred broader principles—such as fundamental rules for how to use and protect the queen in chess—that "had never been reached by humans, even grand masters of chess."³⁴

In all these settings, human players who come up against AI models consistently describe moves that seem to them odd, foolish, and perhaps, most of all, inhuman. But it's not clear that such out-of-the-box strategic acumen in competitive games would easily apply to many of the strategic decisions required of an active state. After all, in developing an antipoverty program, AI would not be trying to find some crazy new move within the tightly regulated rules of a competitive game. But certainly, in military and perhaps geostrategic contexts, extending this kind of decisionmaking originality to two strategic choices could theoretically generate significant advantage. Emergent models generate original creativity, and AI can provide a source of important fresh approaches to strategists in any context—but only as an adjunct to the complex, nonoptimizable work of strategic judgment.

Hollowing Out Rather Than Empowering . . . or Empowering Too Well

AI will, therefore, provide important new capabilities to active states. But as in all of these characteristics, this connection can go very wrong. In such a case, this could take the form of AI weakening active states

rather than strengthening them or fueling a disastrous power overreach by adventuristic states drunk on the potential of their shiny, new tool.

To start with, AI won't empower active states if they can't understand it. There is a very real risk of an immense gulf between the public and private sectors in the AI Era—one that would cripple governments' ability to understand or control the technology in the common interest. Governments struggle to keep pace with the implications of technologies during technological revolutions, and the AI Revolution will pose this challenge as no period ever before has done.

AI will also begin to empower nonstate actors seeking to undermine government control: criminals, hackers, terrorists, extreme activists, and others. In the same way that AI can empower less skilled workers and smaller organizations, it can put powerful—and terrifyingly perilous—capabilities into the hands of small, violent, disruptive groups. These capabilities include tools for automated cyberattacks, massive disinformation campaigns, and the generation of biological pathogens.³⁵ Looking back at Figure 8.1 comparing public and private-sector progress in AI, I think that the result would be to put the arrows near the bottom—the capacity of governments to usefully employ AI—into *reverse*.

The same phenomenon could emerge in the public sector's relationship with a less malign but still self-interested and sometimes dangerous set of actors—private-sector firms. As much as private-sector dynamism is one of the basic foundations for competitiveness, constraining corporate actions that undermine the public good, such as polluting the environment, cheating on taxes, and delivering unsafe medicines or automobiles, is an essential role of government. Citizens expect their government to protect them from such risks. But a private sector empowered with AI that government agencies cannot adequately monitor or regulate will likely slip further out of public control.

There are signals that this process has begun already and would undermine the capacity of governments to act for the common good against private interests. French political scientist Oliver Roy, in his assessment of weakening culture, has argued that part of the problem is that “[n]eoliberalism weakens the nation-state because the state no longer has the means (and often no longer the will) to intervene in the economic sphere, over whose actors and rationale it has no grasp; it has no control over the forces of globalisation.”³⁶ An AI-enabled assault on public authority would take this dynamic to another level and, in the process, worsen the legitimacy crisis of public institutions.

If states do not express the will of their people to promote the competitive standing of the nation—if their role has been usurped by AI programs in important ways—their citizens will lose an essential form of agency. They will partly lose it to private-sector actors whose interests become the primary focus of AI-driven competitive strategies. Although an active state guided and operated by AI models might be effective in some ways, it would trigger a form of alienation whose effects would be extremely difficult to predict.

Apart from the danger of undermining the public-sector ability to promote the common good, there is a very different and, to some degree, inverse risk: AI could tempt an *overactive* state because it can sense, anticipate, and do just about everything. This would be a version of the risk highlighted in Chapter 5 of excessive, overweening ambition (in this case, transmitted to state action). This would be a scenario in which public officials, perhaps under the influence of overeager scholars and researchers, decide that any social, political, or economic problem is subject to a calculated solution—an optimal answer in mathematical terms—thanks to AI. We could see a rush toward AI-powered social engineering, fueled by the same sort of good intentions that produced such historical disasters as large-scale housing projects that turned out to only exacerbate many of the social ills they were supposed to cure.

The result could be not only failure but also a dangerously divided nation. The RAND study on societal competitiveness found evidence of “the ways in which bold state intervention can ignite social tensions that ultimately destabilize the society.”³⁷ If an AI-advised regime tries to force disruptive social change on a coun-

try or rips out long-established approaches to problems without warning—especially if these radical steps are justified through the craven claim that AI told us to do it—the result could be resistance or even rebellion.

Summary: Using AI to Revitalize the Active State

I've argued in this chapter that a right-sized active state (at multiple federated levels)—as part of a wider process of effective governance—is crucial to a national competitive advantage. Effective governance arguably constitutes the *most essential* route to national advantage in the coming decades. But the route to success for this characteristic is strewn with land mines: The application of AI to various governance tasks can easily go wrong, and the evidence, so far, is that public agencies have often thrown algorithms at problems before they knew how to use them, sometimes with disheartening human outcomes.

The main conclusion I draw on this characteristic is that although AI has tremendous technical potential to enhance the efficiency and effectiveness of active states, its broader effect on societies poses a significant risk of further fracturing the bonds of solidarity and trust that bind a political community together and allow it to act in its collective interest. The uncontrolled application of AI in the public and private sectors seems likely to create new versions of legitimization crises—ones fueled by the thoughtless, often arbitrary or even predatory application of algorithmic solutions to public problems. This alienation will be particularly strong if governments insist on using AI to solve social problems for which its objective rationality isn't appropriate. We could be stepping into a future in which public institutions become more effective in some ways, but the broader social contract between states and their citizens becomes even more destabilized.

One logical implication is that any nation (or governance institution) hoping to employ AI for strategic advantage needs a deeply grounded theory of how it plans to achieve that result that is backed by extensive research, experimentation, and the process of learning from experience. Nations that rush into this process blindly will both miss the chance for significant efficiencies and court dangerous outcomes. As with my general theory of the active state, this isn't an invitation for central control over every aspect of the process; the only way to realize this potential will be through a constant process of entrepreneurial experimentation. But governments at various levels can identify key principles and norms to uphold, offer support for experiments, set minimal guardrails, share best practices, and, in other ways, do that good catalytic work of an effective active state.

The United States needs to be thinking about these issues partly because other countries are already doing so. The coming decade will involve a competition among active states—a contest in the pursuit of effective governance—to not only gain the most benefits from AI but also promote the broader societal qualities essential for national power. China is clearly well down this path. Saudi Arabia and the UAE are staking a potential claim to advantage in the AI Era.³⁸ Canada has a sovereign AI strategy. In this context, how well governments employ AI to improve their operations—not only for the material efficiency gains but also to relegitimize their authority in the eyes of a cynical citizenry—will play a significant role in shaping national competitiveness in the coming decade. That effort will inevitably be expressed through the use of AI to improve the effectiveness of many social institutions—an imperative that I turn to in the next chapter.

Notes

¹ Mazarr, 2022a, pp. 157–180.

² Mazzucato, 2024.

³ For a brief survey, see Dylan Gerstel and Matthew P. Goodman, “Japan: Industrial Policy and the Economic Miracle,” in *From Industrial Policy to Innovation Strategy: Lessons from Japan, Europe, and the United States*, Center for Strategic and International Studies, September 2020, pp. 5–9. Persuasive analyses suggest that although there were limits to the effectiveness of government interventions, they did set the context for success in critical ways. For example, see Ellis S. Krauss, “Political Economy: Policymaking and Industrial Policy In Japan,” *PS: Political Science & Politics*, Vol. 25, No. 1, March 1992; Komiya Ryutaro, “Industrial Policy in Japan,” *Japanese Economic Studies*, Vol. 14, No. 4, 1986; Lalita Som, “Japan’s Growth Experience: Post–Second World War and Recent Times,” in *The Capitals of Nations: The Role of Human, Social, and Institutional Capital in Economic Evolution*, Oxford University Press, 2014; and James E. Vestal, *Planning for Change: Industrial Policy and Japanese Economic Development 1945–1990*, Clarendon Press, 1995.

⁴ W. Walker Hanlon, “What Can We Learn from Britain’s Laissez-Faire Experiment?” *Broadstreet*, Substack, June 25, 2021; and W. Walker Hanlon, *The Laissez-Faire Experiment: Why Britain Embraced and Then Abandoned Small Government, 1800–1914*, Princeton University Press, 2024.

⁵ Ron Harris, “Government and the Economy, 1688–1850,” in Roderick Floud and Paul Johnson, eds., *The Cambridge Economic History of Modern Britain*, Cambridge University Press, 2004, pp. 205–206.

⁶ One way in which the Japanese active state did so was with a sensible and restrained macroeconomic approach. By “maintaining a small and balanced budget, fairly low and stable interest rates, relatively low rates of taxation, stable prices, brief and mild recessions, minimal defense and social welfare expenditures, the Japanese government helped provide an exceptionally favorable economic environment for private enterprise” (Katsuro Sakoh, “Japanese Economic Success: Industrial Policy or Free Market?” *Cato Journal*, Vol. 4, No. 2, Fall 1984, p. 541).

⁷ Mazarr, 2022a, pp. 159–160.

⁸ Mazarr, 2022a, p. 159.

⁹ Mazarr, 2022a, pp. 161–162.

¹⁰ Mazarr, Dale-Huang, and Sargent, 2024, pp. 3–4.

¹¹ John Micklethwait and Adrian Wooldridge, *The Fourth Revolution: The Global Race to Reinvent the State*, Penguin, 2014, pp. 5, 14.

¹² Mazarr, Dale-Huang, and Sargent, 2024, p. 20.

¹³ Mazarr, Dale-Huang, and Sargent, 2024, p. 23.

¹⁴ Jürgen Habermas, *Legitimation Crisis*, trans. by Thomas McCarthy, Beacon Press, 1975, p. 2.

¹⁵ Habermas, 1975, p. 5.

¹⁶ Habermas, 1975, p. 75.

¹⁷ Habermas, 1975, p. 4.

¹⁸ For an especially bracing estimate, see Joshua Rauh, “Playing with Fire: The Interest Coverage of US Federal Government Revenues,” *Liberty Lens*, Substack, October 7, 2024.

¹⁹ A 2025 study by three Brookings scholars noted that some effects of AI might be to extend lifespans in ways that ultimately increase health care costs and projected that AI advances could reduce deficits by between 0.8 and 1.5 percent of GDP—but could also increase deficits by as much as 1 percent of GDP over the same time frame. See Ben Harris, Neil R. Mehrotra, and Eric So, *The Fiscal Frontier: Projecting AI’s Long-Term Impact on the US Old-Age Entitlement Outlook*, Brookings, February 20, 2025.

Another 2025 study argued that an AI-powered radical growth surge would very probably increase long-term interest rates, thus increasing the cost of financing U.S. debt. See Seth Benzell, “Could AI Save Us from Making Hard Choices About the Budget?” *Empiricrafting*, Substack, July 11, 2025.

²⁰ For the projection, see Committee for a Responsible Federal Budget, *Analysis of CBO’s March 2024 Long-Term Budget Outlook*, March 20, 2024.

²¹ Amodei, 2024; Mark Jackley, “Using AI in Local Government: 10 Use Cases,” Oracle, August 7, 2024. The article from Oracle’s corporate website might have an incentive to sell AI services but does cite specific examples of government initiatives.

²² An interesting official repository of use cases is Mark Fagan, “AI for the People: Use Cases for Government,” Harvard Kennedy School, M-RCBG Faculty Working Paper Series, 2024-02, August 2024. See also Sanam Hooshmand, Chelsea Canada, and William Clark, “Artificial Intelligence in Government: The Federal and State Landscape,” National Conference of State Legislatures, November 22, 2024.

²³ Damayanti Chatterjee, “Generative AI Could Save the UK’s Public Sector £38 Billion a Year and Significantly Reduce Waiting Times,” *Public First*, November 26, 2024.

²⁴ Stephen Goldsmith and Juncheng Yang, “AI and the Transformation of Accountability and Discretion in Urban Governance,” last revised June 9, 2025.

²⁵ For example, a 2024 study found that although LLMs had trouble generating full plans from scratch, they were very useful in providing feedback and ideas that improved a human-run planning process. See Haoming Li, Zhaoliang Chen, Songyuan Liu, Yiming Lu, and Fei Liu, “Systematic Analysis of LLM Contributions to Planning: Solver, Verifier, Heuristic,” arXiv, arXiv:2412.09666, December 12, 2024.

²⁶ Many initial applications seem to be underway for these sorts of purposes. For example, see Amanda McGrath, “10 Ways Artificial Intelligence Is Transforming Operations Management,” IBM, July 11, 2024. There’s a profit motive at work in corporate promotion of AI applications, but the article describes some actual use cases.

²⁷ Frank Othengrafen, Lars Sievers, and Eva Reinecke, “From Vision to Reality: The Use of Artificial Intelligence in Different Urban Planning Phases,” *Urban Planning*, Vol. 10, January 2025.

²⁸ Philipp Schoenegger, Indre Tuminauskaitė, Peter S. Park, and Philip E. Tetlock, “Wisdom of the Silicon Crowd: LLM Ensemble Prediction Capabilities Match Human Crowd Accuracy,” arXiv, arXiv:2402.19379, February 29, 2024, p. 5.

²⁹ For an interesting discussion of this topic, see Walker, 2024.

³⁰ Hendrycks, 2023, p. 4.

³¹ Richard P. Rumelt, *The Crux: How Leaders Become Strategists*, PublicAffairs, 2022, p. 32.

³² Cade Metz, “In Two Moves, AlphaGo and Lee Sedol Redefined the Future,” *Wired*, March 16, 2016.

³³ Narayanan and Kapoor, 2024a, p. 138, describes how a GPT model figured out the condition of the larger game from calculating individual moves: “In a sense,” they write, “it has learned the rules of the game despite never being explicitly told.” See also Tegmark, 2017, p. 128.

³⁴ Kissinger, Mundie, and Schmidt, 2024, p. 65.

³⁵ James P. Rubin and Darjan Vujica, “AI Is Supercharging Information Warfare: And America’s Defenses Aren’t Ready,” *Foreign Affairs*, November 19, 2025.

³⁶ Olivier Roy, *The Crisis of Culture: Identity Politics and the Empire of Norms*, trans. by Cynthia Schoch and Trista Selous, Oxford University Press, 2024, p. 29.

³⁷ Mazarr, 2022a, p. 179.

³⁸ North, 1990, pp. 3–4.

Effective Institutions

The successful nations of the Industrial Revolution had several major themes in common. Britain, the United States, Germany, and Japan each benefited from all seven characteristics of national competitiveness, but they were able to take advantage of those factors partly because of one quality in particular—*effective institutions*. This term refers not only to formal organizations (such as companies and government agencies) but also to the laws and rules of informal norms of conduct and social habits that continue over time and shape social behavior, which enable and constrain individual and collective action. In this sense, institutions embody the scaffolding for both kinds of agency—of individual citizens and social groups—that I have been stressing.

A good example of an institution critical to national fates is one of the most fundamental: the rule of law. In countries enjoying reliable, nondiscriminatory rule of law, people have a greater incentive to strive for success because the odds are that no powerful thug will be able to step in and grab the fruits of their labor. Countries with relatively honest and effective public-sector institutions, such as government agencies, gain tremendous advantage in a variety of ways. Those whose public sector is ineffectual and rife with corruption can't compete over the long term. In this way and hundreds of others, a society with fairly effective institutions creates the context for national progress. One without them risks stagnation and collapse.

Sustaining effective institutions and figuring out how to refine and sometimes transform them to meet the demands of the 21st century are essential pillars in any effort to promote U.S. competitive advantage. The question for the United States is the following: How can AI help achieve this goal, and are there ways in which it might threaten the effectiveness and coherence of U.S. social institutions?

The characteristic overlaps a great deal with the notion of an active state because well-functioning states need to embody good public institutions. But the two concepts are distinct: The idea of an active state refers to the catalytic role that governments play in creating the context for competitive advantage. This characteristic of effective institutions speaks to the organizational, normative, and habitual mechanisms that help organize social behavior in areas well beyond government structures.

This chapter's thesis is simple. As with the active state as a broader phenomenon, AI has the potential to make organizations—the most tangible form of institutions—more effective and more efficient. But it could also make them more peremptory and more remote. Meanwhile, AI's effects on informal institutions, more broadly defined—the norms, codes, and rules that govern social behavior—could be destabilizing. The default effect of AI is likely to lean in the direction of further alienation and disempowerment of average citizens, given the existing, unsettled social and political context and the primary motivations behind the application of AI, left to the devices of uncontrolled, market-driven adoption. There will be hundreds of success stories and powerful examples of AI improving the delivery of services or improving customer interfaces. But there is a profound risk that the positive effects of such examples will be overwhelmed by a larger trend toward inauthentic, often arbitrary institutional forms and processes.

Defining the Characteristic of Effective Institutions

One of the tricks in discussing institutions is that many people subconsciously equate the term with organizations—tangible establishments, such as a federal department or agency, a foundation, or a university. In the scholarly literature on societies and politics, the term has come to mean something much more comprehensive. It includes a wide variety of tangible and intangible things—organizations, to be sure, but also general principles of social interaction, sets of laws, unspoken norms of conduct, and accumulated habits. Institutions are the whole set of such things that help to shape behavior in a society.

The scholar Douglass North did some of the most influential work on institutions and their importance to competitive standing. He defines them as “the rules of the game in a society or, more formally . . . the humanly devised constraints that shape human interaction. In consequence they structure incentives in human exchange, whether political, social, or economic.” He has in mind more than organizations. Institutions constitute

the framework within which human interaction takes place. They are perfectly analogous to the rules of the game in a competitive team sport. That is, they consist of formal written rules as well as typically unwritten codes of conduct that underlie and supplement formal rules.¹

Another analysis describes institutions in similar terms, with emphasis on the very broad reach of the term:

Institutions are the kinds of structures that matter most in the social realm: they make up the stuff of social life. . . . [W]e may define institutions as systems of established and prevalent social rules that structure social interactions. Language, money, law, systems of weights and measures, table manners, and firms (and other organizations) are thus all institutions.²

The term *institutions* then spans at least three broad categories of components. The first are physical, tangible organizations, such as companies, nonprofits, and police departments. The second are formal, settled, agreed-on laws, rules, and standards. The third are informal institutions, including unwritten norms of conduct—general socially constructed patterns of interaction, and habits—both in general and within specific domains. The sum of these elements has the critical function of shaping social behavior: Institutions delineate the space of acceptable conduct, regulate human interactions, legitimize authority and decisions, facilitate collective action, and (in their more informal varieties) sustain the norms, traditions, and values of the community.

Institutions emerge from a process of social construction and are sustained and legitimated by social coordination and belief. Physical organizations have buildings, equipment, and resources, but, even in those cases, their effectiveness, power, and longevity are determined by the collective engagement with and support for the institutions. Institutions, in general, reflect shared meaning as much as tangible and material reality. As a result, trends that disrupt social coherence and belief systems—as AI will do—can have substantial implications for institutions.

In the RAND study on the social foundations of competitive advantage, I tried to understand what *effective* institutions look like. I was surprised to find that, in the extensive literature on institutions, few authors defined effectiveness with any precision. To a certain degree, it’s common sense—*institutions that are working well presumably successfully fulfill their duties in a judicious way*. But a few more-specific criteria can be helpful in defining effective institutions, whether tangible or intangible. Effective institutions are the following:

- **Effectual in achieving economic and political purposes.** In the simplest sense, institutions are effective when they achieve, to a significant degree, the purposes for which they are established.

- **Respected and followed.** To have consistent effects, formal and informal institutions must shape behavior in meaningful ways.
- **Professional.** The more-formal organizations, including law- and rule-promulgating ones, will typically be, at least partly, professionalized—that is, run by a group of people trained in useful techniques and who uphold standards and norms that contribute to effectiveness.
- **Objective and rule based.** Effective institutions must be grounded in objective rules, standards, or norms rather than being arbitrary or subject to corrupt influence, favoritism, or elite bias.
- **Perceived as effective and legitimate.** Finally, to serve their broader role of enhancing social coherence by competently addressing social and economic challenges, institutions must be *perceived* as effective and legitimate.³

An abundance of research evidence testifies to the connection between institutions defined in these terms and competitive advantage. Good formal and informal institutions not only help to smooth social transactions and, thus, promote social stability but also create a better environment for economic activity. They build trust in a society and shape behavior in ways that either promote or constrain collective action for common purposes. They protect inventions and property and, thus, empower and reward investment and entrepreneurship. Dozens of studies validate these relationships and show the importance of institutions for such national goals as innovation and economic growth.⁴

Daron Acemoglu and James Robinson have stressed the importance of a specific kind of institutions—*inclusive* ones that offer fair and nondiscriminatory treatment to all members of a society (as opposed to *extractive* institutions, which are used by oligarchs or rulers of some kind to grab resources from society). Inclusive institutions, they write, “are those that allow and encourage participation by the great mass of people in economic activities that make best use of their talents and skills and that enable individuals to make the choices they wish.” Their effects are clear in historical terms: On the one hand, “Egypt is poor precisely because it has been ruled by a narrow elite that have organized society for their own benefit at the expense of the vast mass of people.” On the other hand,

Countries such as Great Britain and the United States became rich because their citizens overthrew the elites who controlled power and created a society where political rights were much more broadly distributed, where the government was accountable and responsive to citizens, and where the great mass of people could take advantage of economic opportunities.⁵

Even in promoting this essential characteristic, there are risks. Nations can become overeager in their pursuit of institutionalization, producing a society hamstrung by rules, regulations, and bureaucracy and stifled by a straitjacket of conventional wisdoms and orthodoxies. Many scholars have observed the tendency of mature, developed nations to become strangled with selfish interest groups, bureaucracies, and rules.⁶ Many studies have pointed to the role of such trends in stifling the development of formerly dynamic societies.⁷

Beyond formal, tangible organizations, informal institutions—the values, norms, and socially constructed public habits and practices—play an equally critical role in setting the context for competitive advantage. Scholars have demonstrated their effects in dozens of ways—to take just two examples, extensive research has demonstrated the central role of cultural values in determining development outcomes and the effect of low-trust societies on economic performance and social stability.⁸ A nation’s often unplanned, emergent, sometimes unspoken norms and values either foster or undermine the other qualities I’ve surveyed—its ambition, solidarity, degree of shared opportunity, the functioning of its state, its intellectual environment, and the ways in which it benefits from or is handicapped by elements of diversity and pluralism. If formal institutions are the hardware of a society, informal ones are its software—the operating system that governs many outcomes. And AI’s effects on these less-tangible forms of institutions might be of the most concern.

AI and Institutions: Efficient Organizations

In terms of formal organizations, AI offers significant promise: It will provide dramatic ways to reverse some of the existing trends in excessive bureaucracy and institutional overreach. AI is likely to have these effects in several ways.⁹

One is through promoting basic organizational efficiency and effectiveness in multiple tasks. AI can do this by simplifying paperwork; streamlining engagement with customers, stakeholders, or citizens; discovering unknown blockages; and creating efficiencies. LLMs are already digesting whole codes—the tax code, a building code, the U.S. federal code—and finding detailed information or suggesting simple principles that reflect the spirit of the code, as well as a handful of absolute prohibitions for which safety is at stake.¹⁰

Dario Amodei has suggested a second role for AI in enhancing organizations: It could improve not only the performance, in terms of efficiency, but also the perceived fairness of various institutions. Although it would face a challenge of dealing with some very nuanced and complex judgments, “AI might be smart enough for this: it is the first technology capable of making broad, fuzzy judgements in a repeatable and mechanical way.”¹¹ AI models could potentially offer the sort of “consistent, stable, dispassionate, intelligent agents” operating in a bureaucratic setting that Weber had in mind.¹²

AI could play this role in all sorts of applications. To take a legal example, a jurisdiction could have three LLMs digest all the evidence from a trial and not only offer three independent suggested verdicts or judgments but also lay out the areas of uncertainty and how strongly the evidence supports a given outcome. A human judge could then compare them and synthesize the results for an overall decision. In as banal a setting as traffic court, AI could direct lawyers and judges to common-sense results that serve the purpose of justice without being unduly punitive. In child-custody cases, AI models programmed to achieve humane, objective, common-sense outcomes might help courts avoid needlessly technical and bureaucratic decisions. In terms of general operations and strategies for specific interactions, police departments in the United States have already begun experimenting with the latest generative AI to discover ways of improving the efficiency and effectiveness of their operations.¹³

If AI generates such results, it could restore some degree of trust in institutions, especially in terms of what Anthony Giddens has termed *abstract systems*. If people perceive that social institutions, such as law enforcement, customer service, and medical claim processes, are operating more efficiently and more effectively—if the sense of alienation from abstract systems gets mitigated to some significant degree—that trend could set the stage for an important renewal of the degree of public trust and faith in at least some proportion of social institutions.¹⁴

But there’s no guarantee at all that this effect is what we’ll get. In fact, there’s a strong chance that, in terms of Americans’ interactions with formal institutions, the transition process will be messy and alienating. Part of the problem is that AI models will be deployed in institutional settings to serve the interests of the organizations, not citizens or customers. The goals will be organizational efficiency rather than user empowerment.

One particular danger is that the public sector will lag badly in the use of AI to boost efficiency. Government agencies have neither the expertise in AI nor the budgets to deploy it in broad ways. Another risk, raised by Arvind Narayanan and Sayash Kapoor, is that ineffectual institutions will turn to AI to mask their failures, without necessarily making things better. They give the example of hiring, which is a process that arguably hasn’t improved in decades and is governed by no clear research evidence. Firms have turned to consultants offering AI solutions to improve their hiring practices and, sometimes, end up making them even more inhumane and arbitrary. This is especially true at underfunded institutions that grab for AI-based solutions as an alternative to having the resources to do them well.¹⁵

There are also strict limits to the kinds of organizational functions that can be revolutionized by AI. As the Amsterdam welfare example demonstrates, decisionmaking and policymaking processes that areulti-

mately *political rather than technical*—in the sense that the choices reflect subjective judgments about colliding norms and values that aren’t resolvable through any sort of algorithm—are mostly immune to optimization through AI. An AI model might demonstrate conclusively (as arguably some algorithms and even basic math have done already) that it is more efficient to shift Medicare dollars from the last month of life to preventive care. But the choice to allow some people to die faster than they otherwise might is a profound moral judgment—one in which AI could advise but must ultimately reflect the outcome of a moral dialogue in a political community.

In an ideal world, AI would support social organizations by providing amazingly useful advice and functions in support of technical issues that improve efficiency and effectiveness. If an automated system can run through the elements of registering a car in five minutes, that’s fantastic; if an AI copilot takes 90 percent of the administrative and paperwork burdens off of doctors, that’s a big win. But we still need social institutions to create the context and mechanisms to continually manage the tougher questions and political judgments that demand balancing multiple goals and values. If the spread of AI were to continue to weaken the institutional foundation for political action, even if it boosts efficiency for many technical problems, it could gradually poison our ability to come to any stable political judgments on the more important issues. There is a real danger that such a process has already begun.

AI and Organizational Fairness

Another way in which AI has the potential to reshape the institutional landscape is by offering tools that empower citizens without wealth or power to advocate for and defend themselves against excessive or arbitrary organizational power. An essential aspect of the existing bureaucratic landscape is the size, complexity, and sometimes bias of large organizations, whether public or private. We are already seeing some evidence of ways in which AI can empower average citizens against these monoliths. Such an effect would reshape organizations most directly but would also likely generate second-order implications for broader social patterns and norms.

Take, for example, taxes. People already have access to relatively inexpensive help through large tax preparation services. For a relatively small fee, people can buy insurance against being audited, with the promise of legal representation and expertise. AI models could take this trend to the next level, providing every taxpayer with a tax law expert on demand at all times, a chatbot that could go toe-to-toe with any IRS enforcement office. AI could also help governments streamline their tax codes.

AI models can protect people against abuse in the legal domain more broadly. Legal disputes are hugely influenced by the resources each side can bring to bear: Large corporations have repeatedly crushed individuals or small groups trying to challenge them by simply outspending them in legal advice. The practice became so common that it gained a formal name and suggestive acronym: Strategic Lawsuits Against Public Participation (SLAPP). More than 90 percent of lower-income Americans confront civil legal issues without any legal assistance, and companies around the world have undertaken hundreds of punitive lawsuits to silence critics.¹⁶ As one study concluded,

A growing corpus of literature, including many articles in this volume, has highlighted the particularly deleterious effects of legal financial obligations among severely disadvantaged groups, such as those living in poverty, people who are unhoused and unstably housed, and those returning from incarceration.¹⁷

In theory, if a citizen could call on a small army of AI model-lawyers, the result might be to equalize the playing field to a degree.¹⁸ Yet there are many barriers to this positive result. AI models might mislead clients who are unable to afford high-priced legal help to oversee the models’ advice.¹⁹ Large institutions could also

use AI, and there's a risk of uneven access to frontier models and their legal advice—what one scholar has called a “two-tiered” system of AI legal empowerment.²⁰ There's a chance that chatbot-versus-chatbot legal actions would free average citizens from some of the financial punishment of the law—but, as with all the positive social outcomes we're encountering, it will require careful choices to make it happen. AI won't magically fortify average Americans against massive organizations without a careful tending to the norms and rules governing such legal encounters.

Moreover, we'll only achieve these benefits if someone actually builds AI applications that can do these things and, in some cases, if there are policymakers pushing reforms of laws, codes, or institutional rules to allow such applications to be used by average citizens. It's not at all clear that the forces governing AI outcomes in the United States—or other countries—will fill those gaps. Governments aren't likely to have the expertise or funding to train bespoke models focused on empowerment and wouldn't necessarily have the incentive anyway to strengthen their citizens' ability to be more aware and more empowered in fighting government action. Private firms will be even less prone to serving up ways for their customers to push back against decisions, bills, rules, or other corporate decisions.

It is easy to imagine ways in which AI can empower individuals and maximize their agency relative to large organizations. One would be a wide-ranging debureaucratization initiative, supported by AI in critical ways (but grounded in a significant reform effort) and designed to ease the vise lock of massive, impersonal, often predatory organizations on Americans' lives. But we as a society will have to make the choice for that to happen. And it's just not clear that the leading sources of power and influence in U.S. society (or others) will see the need—or have the desire—to make it happen. The overriding power and interests will tilt the playing field in favor of pure efficiency and organizational power, not individual agency.

A Leading Danger: Supercharging Bureaucracy and Arbitrary Rule by Algorithms

New technologies don't tend to have egalitarian effects at first, when their financial and other benefits are often used to benefit wealthy and powerful individuals and companies that monopolize their early application. AI will not be any different: The leading interest of the AI labs or large firms seeking to apply AI isn't to slim down hyper-bureaucratic processes. Take Mancur Olson's idea about the accumulation of interests in mature societies: All those interests will continue to seek special treatment in the form of laws, rules, regulations, bureaucratic capture, and more. They'll just be using AI to help them do it. In more purely bureaucratic terms, a health agency won't be any more likely to surrender authority over medical rules than it is today. It would use AI models to justify and even expand its role, even if some of its rulemaking continues to be arbitrary. Someone could be told that they can't argue with the judgments of the Medicaid Cost Committee—it's advised by AI.

The result could be a stifling amplification of the late modern pattern of bureaucratic control—a hyper-extension of the ambitions of often impersonal and increasingly algorithm-driven institutions in ways that threaten human autonomy and dignity. Part of the problem in many societies today is that detached, often arbitrary bureaucracy has substituted itself in the oversight of social interactions for the operation of what we would consider authentic operations of social capital.

It's very easy to picture organizational leaders, galvanized by the new power of AI, taking up this baton of institutional overriding of authentic human social choice. This outcome would be especially likely if advocates of AI can claim that the models can come to understand social processes and causalities in ways that human beings never would be able to. Some will argue that we should trade our human organizations with AI models for important tasks—the very institutions that we now think of as human collective expressions

in the basic operation of society. Rather than the members of the local chamber of commerce or rotary club gathering to decide on what local charities best deserve support in the coming year, they might simply turn to an AI model to tell them where to give their money, using AI calculations of effectiveness and efficiency. The result would be a massive, constraining filter on human agency.

It's all too easy to imagine what the next phase of AI-energized hyper-bureaucratization might look like. Today's bureaucratic behemoths would persist, along with their rulemaking, regulation-publishing, and life-controlling aspects. Only now, they'd be using various AI applications to achieve greater efficiency in doing so.

This is likely to involve further entrenching many impersonal and often arbitrary forms of algorithmic decisionmaking that have already been deployed—except they'll have the gloss of superintelligence around them. A super-sophisticated generative AI mortgage loan determination engine would issue denials that could not be appealed because advanced AI couldn't be wrong. Semipredictive policing would pull people off streets with the assurance that superintelligent AI has calculated the risk they pose.

The problem—apart from the inhumanity of such processes—will be that, as brilliantly as AI models might be in calculating various odds and probabilities and combining many forms of data, it will still not be able to make the last leap to engage with people as actual individuals and make exceptions to rules, as the instinct of the loan officer or the judge might, for example. Even if those instincts are sometimes wrong, the idea that, in our social institutions, we are being dealt with as individuals—as unique people being considered in singular processes—is fundamental to our concepts of dignity and justice. It is precisely when institutions abstract away from such humanity and impose judgments according to universal rules or general probabilities that our humanity is most at risk. Handled badly, AI could pose profound risks of further dehumanization and the destruction of autonomous human agency.

In a 2022 article, inspired by the research I did for the project on societal competitiveness, I made an argument for an emerging pattern that captured some of the more sinister aspects of this algorithmic control when layered on top of hyper-bureaucratization. I called it the rise of “predatory abstract systems”: the fact that citizens of advanced nations are increasingly subject to “risks and harassment radiating from a critical mass of large-scale systems that operate based on thickets of rules, procedures, algorithms, habits, and laws so dense no average American can comprehend them.”²¹

By this concept, I had in mind organizations and processes of all kinds—public, private, and digital media. In that article, I discussed Anthony Giddens's notion of abstract systems and cataloged the ways in which so many such systems impose sudden, immense, often arbitrary power on average citizens. In the form of lawsuits, tax judgments, insurance claim denials, and many other mechanisms, abstract systems are now

not merely constricting or debilitating—they are also, to a significant degree, menacing. They constitute a looming background condition of threatening sudden adverse events, punishments, or judgments that jeopardize economic and psychological security.

This accumulating process of predatory bureaucratic behavior, I argued, places at risk

arguably the central achievement of modern rule of law—the right of autonomous individuals to be judged on the merits of their individual cases. As a result, these abstract systems generate an insistent fear of being vulnerable and powerless and a resentment that the wealthy and powerful face few such risks.²²

Such effects obviously intensify trends of social alienation and grievance that are already so apparent.

AI has the potential to make such systems less or more predatory, depending on the decisions we make about it. The result will have profound implications for the social and psychological qualities underlying national dynamism and, ultimately, competitive advantage. If societies can use AI, at least in part, to level the

playing field with imposing organizational power, and if citizens actually feel less browbeaten by institutional power, the social climate will be more welcoming for the kinds of open-minded creativity and adventurous innovations that fuel national power. People will have more allegiance to their nations and motivation to work and sacrifice in their interest. But if left to the devices of large institutions as they exist today, the application of AI is more likely to serve the goals of strict efficiency in service of productivity or profit rather than institutional quality and responsiveness to average people.

AI and the Social Role of Institutions: The Problem of Social Capital

As defined previously, the term *institutions*—beyond the narrower organizational form on which I focused in the previous section—includes the norms, rules, organizations, habits, and other characteristics of a society that help set the rules of the game for social interactions. They help organize human relationships. They are the code we write that governs our social interactions. The widespread deployment of AI is likely to also affect such informal institutions. These effects will play out in dozens of ways, but some lines of emerging research and informed speculation point to a few leading possibilities. Taken together, they suggest that—similarly to earlier technology transformations—AI has the potential to dilute the normative and ideational mortar that binds societies together and that only with some effort can we assure that it leaves healthy institutions in its wake.

Social capital can be understood as networks and groups in society, “together with shared norms, values and understandings that facilitate co-operation within or among groups.”²³ Scholars, such as Robert Putnam, have been documenting the importance of social capital for years—in his case, studying the post-World War II United States.²⁴ Others have found the role of social capital in earlier periods, such as the gentlemanly ideals and Victorian standards of behavior in industrial-era Britain.²⁵ When social capital is strong and healthy, it provides the essential fuel to social stability and dynamism, providing people with a strong sense of belonging, boosting social trust, easing the route to cooperative ventures and new businesses, and more.

There is strong evidence that this critical layer of interaction in U.S. society has frayed and weakened over the past half-century. The question is how the growth of AI capabilities and use will affect this trend. Some observers have argued that AI could spur greater civic engagement at local levels through such means as easing involvement in community decisionmaking and local service recommendations.²⁶ A few very tentative empirical studies have found that increased use of AI actually makes people *more* likely to engage in offline forms of social engagement, thus building social capital.²⁷ If AI models indeed improve institutional efficiency in public services, they might strengthen citizen engagement.

Yet there are also reasons to expect more alarming outcomes. AI could attack social relationships and the community organizations that facilitate them by substituting for dozens of face-to-face institutions that provide the forge for social capital and allow—or even encourage—people to avoid the difficulties of human contact in favor of virtual relationships. Why join a club when a set of AI bots provides lower-friction relationships? Why bother with exhausting, acrimonious negotiations through a homeowners’ association when AI can simply spit out proposed resolutions? Once we have super-realistic chatbots that can be our friends, therapists, and colleagues, would people have a need to go to the local rotary club to build social capital? A few initial studies have highlighted precisely this danger—that AI could eventually substitute for key institutions that build social capital.²⁸

Moreover, the institutions that underpin social capital rest on a foundation of emergent, socially constructed norms and values. Social capital is just that—social—not merely in a descriptive sense but in its origins and constitution. It serves its purpose only when it reflects the results of true engagement with other human beings. Social capital embodies the accumulated experience of human interaction, not merely a set of

material organizations or formalized rules. Its institutions arise from the history and traditions of a specific place and time.

The widespread substitution of AI for human social processes could thus undermine social capital by thinning the authentic interactions on which it is based. Olivier Roy has discussed the potential for AI to exacerbate the crisis of culture precisely for this reason. “Algorithms,” he notes, “function solely by comparing big data sets from a strictly statistical standpoint: they pay no heed to genealogies of meaning, to cultural allusions or to past history. They do not perform an archaeology of knowledge.” As a result, “this type of research can dispense with the human sciences because it no longer needs a theoretical apparatus. It is enough, at least on the surface, to observe correlations and regard them as constants.”²⁹ In such an AI-mediated information space, “[k]nowledge is no longer a corpus that is the product of a history. Knowledge is ‘flattened’ because it is cut off from its own history.”³⁰

The RAND study on the elements of societal competitiveness emphasized the importance of a balance between the open-minded intellectual adventurism that supports a learning and adapting society and the stable norms, values, and forms of identity that provide ballast to a society, even as it rolls forward—qualities that turn out to be essential for the accumulation of social capital.

In the case of a learning and adaptive mindset, from the standpoint of competitive advantage, a society’s openness to change must be balanced against the need for strong traditions and norms that provide a stable sense of national identity and ballast against a changing world. The lack of any unifying tradition can destabilize a society and create weakness in other traits. In this sense the characteristic of learning and adapting trades off with another one of our nominated sources of competitive advantage—a unified national identity. That characteristic inevitably demands some degree of orthodoxy—certainly some traditions—that will end up constraining the full play of intellectual freedom in some ways.³¹

The United States (among other countries) has experienced so much social and political turmoil over the past decade partly because Americans have begun to get this balance very wrong. The new, the relative, the digital and simulated, and the forces of change (as opposed to tradition) have all become incredibly powerful in ways that surely have helped unmoor people from a solid sense of identity. Another systemic tension afflicting modern societies is what scholars call an *ontological* crisis, referring to a person’s sense of being in the world, the “need to experience oneself as a whole, continuous person in time—as being rather than constantly changing—in order to realize a sense of agency.”³² Such disruptions often emerge when people’s basic understanding of reality is shaken, and this happens in part by weakening tradition and the communal basis of identity.

A society in the grip of an ontological crisis tends to produce political movements heavy on nostalgia, tradition, and orthodoxy. People are looking for solid reassurances as a tsunami of change washes over them. If AI further destabilizes the information environment in the United States and, at the same time, seems to be accelerating the processes of intellectual exploration and adaptive experimentation to a pace that people simply cannot comprehend, the result could be an even more comprehensive ontological crisis that would generate very dangerous political movements.

In Chapter 6 on shared national identity, I raised that peculiar idea of a hybrid human-AI society. Such a future social pattern could have profound implications for social capital: Would we have to measure its health in terms of social interactions among human beings *and* AI models and agents? Would we assess the effectiveness of institutions in promoting social capital in terms of this hybrid conception of society? Nothing about that is clear, but it does make it clear that widespread use of advanced AI will massively complicate the process of judging and ensuring institutional effectiveness.

Summary: Real Potential, Unclear Prospects

The potential exists for AI to achieve major new efficiencies and to redesign institutions in impressive ways. But transformative institutional improvements will not be achieved by AI alone; they will require a wider agenda of reform. Bigger and essential change—addressing the need for a revolution in governance models—will come about only if the emergence of AI is combined with a broader commitment to institutional reform, the popular trust in such a process, and the political will to achieve it. Getting tools to reinvent institutions won’t do much good without a motivation to use them. Simply put, if we want a process of AI deployment that makes institutions more efficient and more effective but also less predatory and more human, we will have to fight for that outcome. Sustaining lasting competitive advantage, thus, demands close attention to how we embed this powerful new tool in the human institutions that organize and mediate our societies.

Notes

¹ North, 1990, pp. 3–4.

² Geoffrey M. Hodgson, “What Are Institutions?” *Journal of Economic Issues*, Vol. 40, No. 1, 2006, p. 2.

³ Mazarr, 2022a, pp. 184–185.

⁴ See, for example, Acemoglu and Robinson 2013; Alberto Bisin and Thierry Vierdier, “On the Joint Evolution of Culture and Institutions,” National Bureau of Economic Research, Working Paper No. 23375, April 2017; and North, 1990.

⁵ Acemoglu and Robinson, 2013, pp. 74, 3–4.

⁶ Arguments by economist and political scientist Mancur Olson about the decline of nations are one example; the visceral attack on excessive bureaucratization by anthropologist David Graeber is another (David Graeber, *The Utopia of Rules: On Technology, Stupidity, and the Secret Joys of Bureaucracy*, Melville House, 2015; Mancur Olson, *The Rise and Decline of Nations: Economic Growth, Stagflation, and Social Rigidities*, Yale University Press, 1982). For a summary of subsequent research on Olson’s theorem, see Jac C. Heckelman, “Explaining the Rain: The Rise and Decline of Nations After 25 Years,” *Southern Economic Journal*, Vol. 74, No. 1, July 2007.

⁷ M. W. Kirby, “Institutional Rigidities and Economic Decline: Reflections on the British Experience,” *Economic History Review*, Vol. 45, No. 4, November 1992, p. 656.

⁸ For one representative example, see Lawrence E. Harrison and Samuel P. Huntington, eds., *Culture Matters: How Values Shape Human Progress*, Basic Books, 2000.

⁹ For a general discussion, see Jennifer Pahlka, “AI Meets the Cascade of Rigidity,” in Erik Brynjolfsson, Alex Pentland, Nathaniel Persily, Condoleezza Rice, and Angela Aristidou, eds., *The Digitalist Papers: Artificial Intelligence and Democracy in America*, Stanford Digital Economy Lab, 2024, pp. 89–106.

¹⁰ For one example of how AI is helping builders deal with building code, see Datagrid Team, “Revolutionizing Compliance: How AI Agents Automate Building Code Requirement Extraction,” *Datagrid* blog, April 25, 2025.

¹¹ Amodei, 2024.

¹² Justin B. Bullock, Hsini Huang, and Kyong-Cheol (Casey) Kim, “Machine Intelligence, Bureaucracy, and Human Control,” *Perspectives on Public Management and Governance*, Vol. 5, No. 2, June 2022.

¹³ James O’Donnell, “How the Largest Gathering of US Police Chiefs Is Talking About AI,” *MIT Technology Review*, November 19, 2024.

¹⁴ One study found that AI might save up to 80 percent of the one-third of doctors’ time now spent on paperwork. See Manyika and Spence, 2023.

¹⁵ Narayanan and Kapoor, 2024a, pp. 261–265.

¹⁶ Mary Smith, “Legal Help for Civil Matters Shouldn’t Be Reserved for the Rich,” Bloomberg Law, January 11, 2024. See also Business and Human Rights Resource Centre, “Vexatious Lawsuits: Corporate Use of SLAPPs to Silence Critics,” June 30, 2023. To be fair, some research suggests that actual outcomes in courtrooms are independent of the wealth of the litigants. See B. Zorina Khan, “To Have and Have Not: Are Rich Litigious Plaintiffs Favored in Court?” National Bureau of Economic Research, Working Paper No. 20945, February 2015.

¹⁷ Lindsay Bing, Becky Pettit, and Ilya Slavinski, “Incomparable Punishments: How Economic Inequality Contributes to the Disparate Impact of Legal Fines and Fees,” *RSF: The Russell Sage Foundation Journal of Social Sciences*, Vol. 8, No. 2, January 2022, p. 118. See also Devah Pager, Rebecca Goldstein, Helen Ho, and Bruce Western, “Criminalizing Poverty: The Consequences of Court Fees in a Randomized Experiment,” *American Sociological Review*, Vol. 87, No. 3, June 2022.

¹⁸ See, for example, Carey L. Biron, “Legal Aid and AI Help Poor Americans Close ‘Justice Gap,’” *Context*, August 22, 2024; and Stephen Heitkamp and Sean West, “Gen AI Makes Legal Action Cheap—and Companies Need to Prepare,” *Harvard Business Review*, October 3, 2024. This dialogue is global. See Senthil S. Pandi, A. Mohammed Farook, W. Kingston, and Sathish Kumar Kannaiah, “Enlightening Justice: Empowering Society

Through AI-Driven Legal Assistance,” paper presented at the 2024 Second International Conference on Advances in Information Technology, Chikkamagaluru, Karnataka, India, July 24–27, 2024.

¹⁹ Ashwin Telang, “The Promise and Peril of AI Legal Services to Equalize Justice,” *Jolt Digest*, March 14, 2023.

²⁰ Drew Simshaw, “Access to A.I. Justice: Avoiding an Inequitable Two-Tiered System of Legal Services,” *Yale Journal of Law and Technology*, Vol. 24, 2022.

²¹ Mazarr, 2022b, pp. 62–63.

²² Mazarr, 2022b, pp. 63, 67.

²³ Brian Keeley, *Human Capital: How What You Know Shapes Your Life*, Organisation for Economic Co-operation and Development, 2007, p. 103.

²⁴ Alexandra Hudson, “Bowling Alone at Twenty,” *National Affairs*, Vol. 61, Fall 2020. Many scholars have offered criticisms of concepts and claims in Putnam’s 2000 book *Bowling Alone*, and not all empirical evidence supports every claim. But his core notion of American social and community life fraying has held up. See Robert D. Putnam, *Bowling Alone: The Collapse and Revival of American Community*, Simon and Schuster, 2000.

²⁵ David Sunderland, *Social Capital, Trust and the Industrial Revolution, 1780–1880*, Routledge, 2007.

²⁶ Sarah Friar and Laura Bisesto, “The Potential for AI to Restore Local Community Connectedness, the Bedrock of a Healthy Democracy,” in Erik Brynjolfsson, Alex Pentland, Nathaniel Persily, Condoleezza Rice, and Angela Aristidou, eds., *The Digitalist Papers: Artificial Intelligence and Democracy in America*, Stanford Digital Economy Lab, 2024, pp. 71–85.

²⁷ Yu-Leung Ng, “Exploring the Association Between Use of Conversational Artificial Intelligence and Social Capital: Survey Evidence from Hong Kong,” *New Media & Society*, Vol. 26, No. 3, March 2024.

²⁸ For example, see Husain Al-Ansari, Oksana Gerwe, and Anjum Razzaque, “Impact of Artificial Intelligence Based Social Capital on Civic Engagement in an Environment of Changing Technology: Development of a Theoretical Framework,” *Proceedings of the Industrial Revolution & Business Management: 11th Annual PwR Doctoral Symposium (PWRDS) 2020*, February 10, 2020; and Julia Freeland Fisher, “Will the Rise of AI Spell the Demise of Social Capital?” Christensen Institute, December 5, 2023.

²⁹ Roy, 2024, p. 73.

³⁰ Roy, 2024, p. 73.

³¹ Mazarr, 2022a, p. 234.

³² Jennifer Mitzen, “Ontological Security in World Politics: State Identity and the Security Dilemma,” *European Journal of International Relations*, Vol. 12, No. 3, September 2006, p. 342. The concept was originated by Giddens; see Anthony Giddens, *Modernity and Self-Identity: Self and Society in the Late Modern Age*, Stanford University Press, 1991.

A Learning and Adapting Society

Of all seven characteristics highlighted in the RAND study on the societal foundations of national competitiveness, the one that impressed me the most—and which sat at the hub of the study’s overall model of competitive advantage—was also the most abstract and hard to define: We called it *a learning and adapting society*. This quality embodies an intangible yet absolutely essential combination: powerful intellectual energy and the capacity to use new ideas to adapt to changing circumstances.

For a general sense of this quality, think of a nation with a vibrant intellectual environment, having brilliant scholars and researchers pushing the boundaries of various fields. It enjoys a climate of openness to unusual and unpopular ideas, supported by strong university and research sectors. A powerful competitive drive, especially in the private sector, underpins a constant stream of experiments and variations on old strategies and a sense of adventurous intellectual ambition. Networks of science, research, and analysis create a dense and mutually reinforcing ecosystem of thought.

Now consider the inverse example: Imagine a society shackled by strict orthodoxies of thought, which tend to mock or punish dissenters and those with revolutionary ideas. It is a country without a strong intellectual tradition that has a weak higher education and research sector and a political and social environment built on the idea that strict adherence to tradition is the route to national strength. Over the long term, one of these societies is going to be competitive, and the other isn’t.

We see versions of these colliding patterns in nations throughout history. On the positive side of the ledger, you have Britain in the early Industrial Era, the United States in 20th century, the Italian Renaissance city-states and the Netherlands during their respective heydays, Meiji Japan, and elements of some periods of Ming and Song China. On the negative side, there is the late Ottoman Empire, China and Japan during their periods of self-imposed isolation and domestic repression, and the Soviet Union, especially in its twilight decades.

In one sense, the implications of AI for this characteristic would seem wholly positive. Frontier models are already creating new ideas in several areas of science; even in business, practitioners and scholars are finding that such models can outperform humans in dreaming up new ideas. Having many models to throw at a problem ought to generate far more learning than humans alone could have achieved. In their quick advance and potential for self-improvement, AI models embody the principle of adaptation, and they could fuel it in tremendous ways as they generate and validate new approaches to problems.

But as with all these characteristics, things will not be so simple. People and organizations are likely to apply AI models in support of many forms of learning and adaptation. But the road from there to a *society* that reflects those qualities is less direct. There is a very strong risk that AI could hollow out the authentic human reflection of these qualities. In this chapter, I define the quality of a learning and adaptive society, discuss the ways in which AI is already supercharging that quality, and describe the very real risks that a generalized use of AI poses to the character of learning and intellectual ambition in a society.

Defining the Characteristic of a Learning and Adapting Society

The characteristic of a learning and adapting society essentially refers to the creative, innovative, and adaptive qualities of a nation. As the 2022 RAND study put it:

Highly dynamic and competitive nations are typically thirsty for new ideas and are excited, rather than intimidated, by fresh policies and approaches. They cultivate networks of scientific and intellectual discussion and debate and both allow and support the widespread public sharing of new knowledge. They apply learning in practical ways and continually reassess their ways of doing business. . . .

The essential building blocks of this characteristic include a habit of intellectual curiosity and investigation; the capacity for open debate and dialogue, even if constrained in some ways; the existence of formal mechanisms of analysis, investigation, and debate in social institutions; [a strong] national tradition and historical experience of adaptation in the face of learning; and a societal habit, and structures, oriented to the implementation of the fruits of learning.¹

Such qualities produce a society dedicated to learning and determined to implement new concepts and discoveries for competitive advantage.

This spirit is evident across the key sectors of dynamic societies. One is the quality and reach of the education system, especially at university levels. Another is the strength of what can be described as the research sector, such as corporate R&D labs, university scholars, and think tanks—its size, its level of influence and respect in the wider society, and the degree to which its findings get picked up and used for national advantage. A capacity for adaptation can, sometimes, be found in the protean quality of firms and organizations: Do they regularly reinvent themselves by engaging in business model transformations rather than getting stuck in rigid, ossifying strategies and structures?

In societies reflecting this quality, historians have found one especially interesting common feature: rich and healthy networks of learning and debate among experts. Individual Einsteins can prompt discoveries in some specific fields, but the real progress of understanding—and the applied knowledge that eventually produces economic value—comes through shared ideas. The importance of both intellectual and economic networks of exchange is a major theme of many works on the economic emergence of the West after 1500.

In many ways, this characteristic reflects the spirit of a healthy intellectual environment that the journalist Jonathan Rauch has called the “constitution of knowledge.” He has in mind an ongoing, networked conversation, grounded in empirical claims about reality, in which “no one has the final say” and “no one has personal authority.” It’s a persistent, tough-minded, roiling search for the truth in which correct claims win out. The result is what Rauch calls a system of “organized social persuasion,” one that relies on networked debates and discussions. The primary assumption of this system, he adds, is that “you cannot make knowledge except through the marketplace of persuasion and its multiple layers of checking and vetting.”² Such a process of truth discovery demands a commitment to learning and a drive to discover new facts, an open-minded willingness to have those discoveries debated and discussed, and an actual network of discoverers who are constantly sharing, supporting, and debunking new claims.

But a strong commitment to learning and intellectual discovery is only the first half of this characteristic. The other is *adaptiveness*—a willingness and ability to evolve, change, and refashion habits and institutions as knowledge advances. Adaptation occurs when the fruits of learning and experimentation generate new technologies but especially new institutional forms, governing approaches, and even social structures. The core spirit is that same sense of adventurous open-mindedness to new ideas and their implications.

Edmund Phelps has described the nations that flourished in the Industrial Revolution in very much these same terms. Those societies emphasized innovation, creativity, a

drive to change things, the talent for it, and the receptivity to new things, as well as the enabling institutions. Thus dynamism, as it is used here, is the willingness and capacity to innovate, leaving aside current conditions and obstacles.³

Rigidity, entrenched interests and institutional forms, and path dependence are the enemies of these qualities. All this is closely connected to a parallel source of competitive advantage: A learning and adapting mindset can be envisioned as a version of national ambition and willpower applied to the pursuit of knowledge and active adaptation to that learning.

This characteristic of a learning and adapting society generates competitive advantage in many ways. Nations benefiting from it fuel innovation and growth and remain nimble, preventing rigid orthodoxies that risk long-term stagnation. They engender a global gravitational pull that creates power and influence outside their borders, partly by putting the country at the center of networks of intellectual advance. This trait magnifies every other one of the societal foundations of advantage—focusing national will on productive ends, nurturing an intellectual climate in which opportunity and diversity flourish, and giving institutions the mindset to perform.

AI as Fuel for Creativity and Learning

Advanced and eventually superintelligent AI holds tremendous potential to promote key elements of a learning and adapting society. It is likely to supercharge all kinds of learning in profound ways, from complex research tasks to basic education. Through its ability to undertake thousands or millions of experiments and create simulated worlds to test ideas, it is also likely to become a tremendous source of adaptive energy.

Start with the obvious connection: education. The reach and effectiveness of education doesn't always move in lockstep with a society's ability to learn and adapt. But the most agile nations of the modern era have tended to share one trait—broad, high-quality education, especially in their universities. AI could now remake education in ways that unleash fresh waves of intellectual discovery.

There is already decent evidence that AI could transform learning across all ages, domains, and aspects—one-to-one, personalized, endless explanations; instant video illustrations; whatever delivery formats the user prefers (words, audio, video, or a combination thereof)—leading to a dramatic shift in the creation of skilled human capital. This effect would be especially pronounced in higher education, in which AI is beginning to push the frontiers of knowledge and provide important new capabilities to university research teams. AI models are already producing new insights in such areas as physics, biology, and the treatment of major diseases. This outcome is basically inevitable: AI will provide a tremendous jolt to learning in most areas of human endeavor.

AI will also boost the intellectual energy in society through the basic generation of ideas. AI is already starting to spit out new scientific theories, new technologies, new business models, and much more at a fantastic pace. AI-generated business ideas often stand up reasonably well: As I mentioned in Chapter 2, Ethan Mollick's experiments show that LLMs significantly outperformed students in his class in generating ideas for businesses in the view of invited expert judges.⁴

Others have produced similar results. One study tested human versus AI creative outputs in business-related ideas, such as a new smartphone app. The authors found that, as rated by test subjects, GPT-4 could

generate creative ideas that are rated as more creative than those generated by human creative professionals incentivized for performance. Furthermore, [the authors] showed that the model's superior performance can be attributed not only to creativity in form (i.e., using more unusual language), but also to creativity in substance (i.e., the ideas themselves are more creative).⁵

The literature on AI's role in creative pursuits is immense.⁶ I was surprised to find that scholarly assessments of this issue go back decades: Even in the early days of AI, researchers were already wondering how it could be used as an imaginative engine and what forms of human creativity it might replace.⁷ Much of this literature focuses on classic creative fields, such as art and music, but AI can be employed to boost creativity in many sectors of the economy and society. Architects will use it to conjure innovative designs. Graphic design firms will discover new approaches to jobs. Advertising copywriters will generate 15 versions of a pitch. Any business function that involves some degree of open-minded creative thinking will benefit from AI copilots and, eventually, agents that can take a job from beginning to end.

Other studies have found that human-AI teams wrote more-creative poetry than humans alone,⁸ that AI copilots increased the creativity of telephone salespeople (especially more skilled ones),⁹ and that LLMs are beginning to match human results on widely used tests of creativity.¹⁰ A fascinating study from 2024 compared human-AI teams to human-only crowdsourced ideas for business ideas. It found that “while human crowd solutions exhibited higher novelty—both on average and for highly novel outcomes—human-AI solutions demonstrated superior strategic viability, financial and environmental value, and overall quality.”¹¹

From its effects on education and scientific research to its role as knowledge copilots for many people in society, AI can spur the advance of knowledge in many ways. As it begins to have these positive effects, it will be a powerful boon to human intellectual agency, empowering researchers, educators, students, and others to pursue their interests and make fresh discoveries.

But these great potential advantages only reflect half the story. The same theme emerges in issue after issue: AI won't just create changes on the margins of social trends. It will reshape societies in profound ways—their structure, how they operate, the essence of social relationships, and much more. In these broader ways, AI poses the risk of disrupting and poisoning the knowledge environment of societies in ways that would undermine the learning qualities of nations. AI's potential to spur learning and adaptation could be overwhelmed by its effects on society as a whole. This is, after all, to some degree what we've seen with the computer and information revolution more broadly: immense capacity for intellectual advances and a society increasingly fragmented by an epistemic crisis and alienated and divided by social media.

AI, Experimentation, and Adapting to a New Era

AI won't just fuel learning—it will improve a society's capacity for adaptation. Its greatest strength may be as an engine of experimentation, enabling organizations to imagine and test new concepts, business models, and strategies. Most of all, it broadens the range of questions people can ask—and the answers they can explore. Especially when combined with other technologies, such as advanced manufacturing techniques, AI will also boost adaptation by speeding the delivery of new innovations to market or society. A business or government agency could go from idea to prototype to production in a matter of days, using the models to do a lot of intervening work, such as testing certain designs in a simulated environment. As 3D printing and nanotechnology manufacturing continue to advance, the combination of advanced AI and advanced manufacturing will allow organizations to complete entire development and production cycles in incredibly short periods, empowering the nation and institutions in it to react far more quickly to changes in the environment, whether in the marketplace or military technology.

If properly employed, AI will also spur a whole new systemic capacity for adaptation throughout a society. Its forecasting ability can help anticipate potential changes. Its ability to generate new strategic concepts and plans can fuel a much richer pipeline of adaptive ideas. As it acquires the ability to simulate policy outputs, it will be able to help governments and private-sector actors test the results of various options. The result will be a nation with a turbocharged adaptive capacity.

The Risks: Our Epistemic Crisis and AI

So much for the promise of AI in promoting this characteristic. What are its risks?

One of the main challenges in shaping AI's societal effects is that this revolutionary technology is arriving at a time of worsening alienation and social stress. As I argued in Chapter 1, these dangerous trends manifest in various ways—deep inequalities, a generalized sense of a world out of control, a loss of autonomous agency, and declining faith in major institutions. But arguably the most important single cause of social instability is that information environments in developed and developing societies alike have become poisoned with profound levels of misinformation and disinformation, even as the mediating functions of traditional authorities have been collapsing. The result is what some have referred to as an *epistemic crisis*—a crisis in our development and consumption of information and ways of understanding our world.

Jonathan Rauch made this case in his 2021 book *The Constitution of Knowledge*. Many separate but cumulative trends are disrupting the marketplace of ideas, to the point that U.S. citizens—and citizens of many other countries—are losing a collective sense of reality. Active purveyors of misinformation, the “troll epistemology,” and cancel culture have all played a role in this process, Rauch thinks.¹² The harsh reality of social media (and, increasingly, traditional media) regularly generates outright assaults on those trying to seek truth. All this constitutes a mortal threat to the principles of an effective constitution of knowledge—an agreed-on set of principles for discovering, sharing, and objectively assessing new knowledge about the world—which is the beating heart of a learning and adapting society.

The corruption of information environments and the ebbing of any sort of common picture of reality pose a chilling danger to a learning and adapting mindset. The question is how AI might reshape this trend. Would it be likely to accelerate, cure, or have some other effect on our epistemic crisis?

So far, the evidence of AI's emerging relationship to misinformation and disinformation is decidedly mixed.¹³ There's already substantial reason for worry that AI will exacerbate some of media's worst tendencies. Plenty of individuals and organizations are using LLMs to spew out fabricated information. Some fact-checking and disinformation-battling groups have begun to experiment with AI tools but appear to be having only mixed success.¹⁴ A prominent 2024 report suggests that an LLM-based chatbot was more effective than humans in pulling people away from conspiracy theories—but that's only one study, and I'm quite sure that we'll soon see others showing how other chatbots drew people *toward* conspiracies.¹⁵ Armies of AI agents conducting highly sophisticated propaganda campaigns are already appearing in limited ways.¹⁶ In more prosaic ways, millions of people will use AI—and are *already using* AI—to generate bespoke, manufactured information for their own purposes. One worrying example is in the courtroom, in which AI generates fake evidence.¹⁷

Using AI to salve our epistemic crisis also requires something that is in short support in our information environment today: trust. Expanding sources of information have paralleled many failures in governance and other signals of foolish institutions, greatly diminishing public faith in many institutions (including information dispensers and mediators, such as universities and the mass media), as we have seen. People don't just have more sources of data; they have lost their trust in the nodes in our society that previously would have created that shared epistemic picture.

In such a world, it's not clear that people would trust even superintelligent AI. Many will be put off by the fact that nobody really knows how it works: The black box inside a model isn't a reassuring picture to paint in an era full of technocynicism and conspiracy theories. Many people would probably assume that model's statements can be programmed, even if AI developers try to explain that these results are inherently emergent and unplanned. This would leave the field free for those who want to manipulate perceptions to use AI to generate misinformation.

The result could be to cleave human beings from their epistemic roots—to create some mediating layer of algorithmic processing that's responsible for much of our intellectual exchange and learning. If we outsource our understanding of our social world to AI—and become intellectually passive in the bargain, partly because we simply can't comprehend how AI is generating its magical findings—this would surely sap energy from our information environment.

We're already seeing cases of a related AI-based threat to the public sphere: AI-empowered harassment designed to silence people of differing views. Even in the social media era, the scholar Sherry Turkle notes that "people who use social media are less willing to share their opinions if they think their followers and friends might disagree with them. People need private space to develop their ideas."¹⁸ AI can empower broad-based campaigns of provocation and attacks—generating large numbers of social media posts, manufacturing faked audio and video clips or invented documents—to scare people away from the marketplace of ideas. If the public sphere becomes a battleground of vicious harassment, the quality of public debate, and the learning and adaptation it produces, will be badly undermined.

Killing the Impulse to Learn: Cognitive Off-Loading

Another way of putting this is that AI poses a threat to autonomous human intellectual agency, even as it holds the potential to promote it; in other words, it might promote a generalized cognitive laziness. As human beings increasingly step aside to allow AI to do the intellectual work, it is not clear what residue will remain of true human knowledge-seeking, at least on issues of collective interest and concern. Sherry Turkle's research on people's interaction with technology suggests that greater dependence on digital and virtual sources of information and crutches for cognition undermines people's interior lives and, perhaps ultimately, their ability to think creatively. For people who have had intense digital diets, research suggests that they "are uncomfortable if left alone with their thoughts, even for a few minutes." One experiment showed that people deprived of their phones became so intensely bored that a certain percentage of them were willing to deliver small electroshocks to themselves to fill the time.

All this highlights a particular danger. In the AI Era, if people become ever more cocooned in their engagement with information, their capacity for open-ended learning and creative thought will diminish. In previous chapters, I mentioned research on the temptation to let AI take the wheel and to adopt the passive mode of an observer. That same phenomenon, combined with more people spending time engaged with digital devices rather than alone with their thoughts, could eat away at the motivations for learning and the sources of creative thinking. Evidence is accumulating to suggest that people using advanced AI might tend to increasingly put their brains on autopilot and let the models do the cognitive work.¹⁹

This has to be one of the most significant dangers AI poses for the intellectual energy of a society: It seems almost designed to generate an extreme version of what some experts call *cognitive off-loading*, a process whereby we allow technology to do our mental work for us: navigating a city, making a list, proofreading an essay, doing addition, and so forth. "In study after study, scientists have shown that people who regularly rely on digital help for some tasks can lose capacity to do them alone."²⁰

This is an especially great risk for younger people, whose cognitive development depends partly on encountering and undertaking thousands of demanding cognitive tasks. Studies have shown that students who take notes by hand, for example—and have to put themselves through the cognitive task of handwriting and summarizing—learn better than those who type out their notes. And soon, even personal note-taking will seem passé when an AI model can take a class transcript—or a book, essay, or podcast—and summarize it.²¹ Millions of students are of course already using it for precisely these purposes. Arvind Narayanan worries that "[r]eading for obtaining information is getting intermediated by chatbots."²²

We are in the very early stages of this shift, so I think that people underappreciate the magnitude of what's coming. It's not just that AI is replacing traditional web search. Even when it comes to reading news articles, business documents, and scientific papers,

the vision that technology companies are pushing on us is AI summarization + synthesis + Q&A. . . . It's a tradeoff between speed/convenience and accuracy/depth of understanding—the same tradeoff that was once offered to us when it became possible to search the web to look up a quick fact as opposed to reading about the topic in depth in an encyclopedia.

Just as most people in most cases prefer a shallow web search over deeper reading, most people in most cases will prefer AI-intermediated access to knowledge. Traditional reading won't disappear, but people will do it vastly less often, except in hobbyist reading communities and professions where traditional reading is needed.²³

This kind of passivity, combined with the way in which AI generates findings, has the risk of thinning out the reach and creativity of a society's intellectual exploration. Most leading AI models are trained on the same data and often generate fairly similar answers to questions. The more people and organizations rely on it for learning, the more it might actually serve to suppress the long-term operation of a competitive market-place of ideas in a society.²⁴ A 2024 essay surveyed many studies and concluded,

A growing body of literature on generative artificial intelligence reveals a surprisingly consistent stylized fact: when people use generative AI tools, the set of content they produce tends to be more homogeneous than content produced by more traditional means. Across a wide range of domains including peer review, writing, digital art, and survey responses, access to generative AI tools (GAITs) leads to less diverse outcomes. Researchers refer to this phenomenon—where the use of similar or identical underlying AI tools lead to convergence in outcomes—as *algorithmic monoculture* or *homogenization*.²⁵

The scholar Konrad Lachmayer argues that AI “monopolizes both the language and the way of thinking represented. Instead of relying on the plurality of people, a linguistic and legal monoculture is created.”²⁶ If AI users aren't careful, they might generate very similar strategies to their opponents or generate easily anticipated approaches that don't offer novel competitive ideas.²⁷

It would seem (as I'll argue in the next chapter) that AI could be asked to sidestep this risk; users could request a variety of answers, some of them unusual or counterintuitive. In my experience at least, this doesn't always generate very satisfying results: Often you still get strikingly similar answers. Ideally, of course, workers, students, researchers, and others wouldn't off-load all the needed adaptive creativity to AI—they'd work alongside it and become more empowered themselves (which is one reason many persuasive assessments suggest that the AI Era will demand *more* cognitive complexity on the part of workers—people who are able to engage creatively with these tools, deeply understand their operations, think critically about outputs, and bring their own insight and imagination to the process).²⁸ But the danger of passive off-loading of intellectual demands is clear.

So is the risk that AI will threaten the social foundations of learning and the intellectual process that underpins adaptation. Individuals can generate ideas, but ultimately—as such scholars as Joel Mokyr have shown—true sustained social dynamism is a collective project, one spurred by the intersection of knowledge and concepts and the ways in which networks of creativity emerge and feed into large-scale scientific advances and technological innovation. True advances in learning come from clashes, arguments, cooperations—intense dialogues and debates that test new ideas and discover new connections. As Mokyr puts it, “Cultural evolution of any kind, then, consists of social learning and persuasion.”²⁹ I argued in Chapters 6 and 8 that AI poses dangers to social coherence and solidarity. If it indeed has these effects—if people

become even more isolated and divided, stuck in their silos of information and chatbot relationships—the result will threaten both learning and adaptation.

Summary: The Benefits Are Likely—But the Risks Are Profound

Of all the potential benefits of AI across the characteristics that I surveyed, its contribution to intellectual advances and the general quality and speed of social learning seems already to be one of the safest bets. Advanced AI, among other technologies, is an unprecedentedly powerful machine for understanding our world and making sense of the connections in it. Nations that develop and implement it in world-leading ways stand to gain a very significant competitive advantage in this sense. That advantage might accelerate in the years ahead as more-intelligent AI, approaching and perhaps exceeding the threshold many have referred to as superintelligence, builds on itself and begins generating an ever-larger number of insights.

AI can also contribute greatly to a nation's capacity for adaptation. Just as it is a learning engine, advanced AI will be an incredibly powerful source of new idea generation, experimentation, and simulated trial and error and facilitator of rapid production in ways that grant significant competitive advantage to any nation that beats others to the punch.

But these benefits will surely be accompanied and, to some degree, offset by risks that AI poses to our collective and individual intellectual environment. AI is likely to worsen, at least in some ways, the epistemic crisis being faced by many societies in ways that weaken our ability to reach scientific, social, or political consensus. To the extent that advanced AI models substitute for deep human thinking and theory-building, no matter how intelligent they might be, they will damage our long-term intellectual capacities.

What's at stake in these effects, across the range of AI's implications for the sources of national competitive advantage, is the fate of autonomous human agency as we've known it. That's as true for the areas of knowledge and adaptation as for national ambition, willpower, unity, or shared opportunity. AI can set the context for incredible new levels of learning and effective adaptation. But it can also cause people to slink to the sidelines of our epistemic life, take what AI says on faith, and stop searching for our own answers. Anyone who has used these models can appreciate the lure: For example, I have asked AI for summaries of theories, frameworks to assess an issue, the most important critiques of an argument, or even several strategies to achieve a goal. We ask an AI model and, presto, there is a well-structured and seemingly accurate response. The urge to begin copying and pasting—as opposed to using the response as a prompt to individual exploration—is intense. For people in contexts in which the intellectual exploration is externally mandated rather than intrinsically prompted (such as most college students), the need to retain intellectual agency won't be nearly as powerful. In terms of adaptation, organizations and communities will increasingly be given ready-made policy options that offer an easy answer to their challenges.

Good outcomes won't happen by accident. In the information space, as in other domains, troubling existing trends and powerful incentives for mischief mean that the AI Revolution is certain to pose major challenges. Societies will emerge from this transition as healthier, more-unified, dynamic engines of competition only if they decide to do so—and make sometimes controversial, painful compromises and investments to get better outcomes. Again, we see that fundamental theme: AI will deliver the kind of society—and the degree of social competitiveness—that we *decide* to ensure.

Notes

¹ Mazarr, 2022a, p. 209, 212–213.

² Jonathan Rauch, *The Constitution of Knowledge: A Defense of Truth*, Brookings Institution Press, 2021, pp. 88–89, 92, 95, 99, 109.

³ Phelps, 2013, p. 20.

⁴ Lennart Meincke, Karan Girotra, Gideon Nave, Christian Terwiesch, and Karl T. Ulrich, “Using Large Language Models for Idea Generation in Innovation,” Wharton School Research Paper, last revised October 2024.

⁵ Noah Castelo, Zsolt Katona, Peiyao Li, and Miklos Sarvary, “How AI Outperforms Humans at Creative Idea Generation,” last revised April 14, 2024, p. 28.

⁶ For examples, see Tojin T. Eapen, Daniel J. Finkenstadt, Josh Folk and Lokesh Venkataswamy, “How Generative AI Can Augment Human Creativity,” *Harvard Business Review*, July–August 2023; and Vinchon Florent, Todd Lubart, Sabrina Bartolotta, Valentin Gironnay, Marion Botella, Samira Bourgeois-Bougri, Jean-Marie Burkhardt, Nathalie Bonnardel, Giovanni Emanuele Corazza, Vlad Glăveanu, et al., “Artificial Intelligence & Creativity: A Manifesto for Collaboration,” *Journal of Creative Behavior*, Vol. 57, No. 4, December 2023. One very helpful treatment is Arthur I. Miller, *The Artist in the Machine: The World of AI-Powered Creativity*, MIT Press, 2019.

⁷ For a couple examples of this older literature, see Terry Dartnall, ed., *Artificial Intelligence and Creativity: An Interdisciplinary Approach*, Springer Science and Business Media, 1994; and Jon Rowe and Derek Partridge, “Creativity: A Survey of AI Approaches,” *Artificial Intelligence Review*, Vol. 7, February 1993.

⁸ Jimpei Hitsuwari, Yoshiyuki Ueda, Woojin Yun, and Michio Nomura, “Does Human–AI Collaboration Lead to More Creative Art? Aesthetic Evaluation of Human-Made and AI-Generated Haiku Poetry,” *Computers in Human Behavior*, Vol. 139, February 2023.

⁹ Nan Jia, Xueming Luo, Zheng Fang and Chengcheng Liao, “When and How Artificial Intelligence Augments Employee Creativity,” *Academy of Management Journal*, Vol. 67, No. 1, February 2024.

¹⁰ Erik E. Guzik, Christian Byrge, and Christian Gilde, “The Originality of Machines: AI Takes the Torrance Test,” *Journal of Creativity*, Vol. 33, No. 3, December 2023.

¹¹ Boussoux, Léonard, Jacqueline N. Lane, Miaomiao Zhang, Vladimir Jacimovic, and Karim R. Lakhani, “The Crowdless Future? Generative AI and Creative Problem-Solving,” *Organization Science*, Vol. 35, No. 5, August 2024.

¹² Rauch, 2021, p. 164.

¹³ Narayanan and Kapoor, 2024a, pp. 179–226. One study highlighted the various ways that AI could, at least in theory, enhance the public sphere. See Beth Goldberg, Diana Acosta-Navas, Michiel Bakker, Ian Beacock, Matt Botvinick, Prateek Buch, Renée DiResta, Nandika Donthi, Nathanael Fast, Ravi Iyer, et al., “AI and the Future of Digital Public Squares,” arXiv, arXiv:2412.09988, December 13, 2024.

¹⁴ Christina Clark, “How AI Can Help Stop the Spread of Misinformation,” *UC San Diego Today*, September 17, 2024.

¹⁵ Teddy Rosenbluth, “This Chatbot Pulls People Away from Conspiracy Theories,” *New York Times*, September 12, 2024.

¹⁶ Julius Endert, “Generative AI Is the Ultimate Disinformation Amplifier,” *Deutsche Welle*, March 17, 2024.

¹⁷ Sam Sabin, “Courts Aren’t Ready for AI-Generated Evidence,” *Axios*, July 25, 2025.

¹⁸ Turkle, 2015, pp. 212–213, 310.

¹⁹ Yizhou Fan, Luzhen Tang, Huixiao Le, Kejie Shen, Shufang Tan, Yueying Zhao, Yuan Shen, Xinyu Li, and Dragan Gašević, “Beware of Metacognitive Laziness: Effects of Generative Artificial Intelligence on Learning Motivation, Processes, and Performance,” arXiv, arXiv:2412.09315, December 12, 2024.

²⁰ Sam Schechner, “How I Realized AI Was Making Me Stupid—and What I Do Now,” *Wall Street Journal*, April 3, 2025.

²¹ Allyria Finley, “AI’s Biggest Threat: Young People Who Can’t Think,” *Wall Street Journal*, June 22, 2025.

²² Arvind Narayanan, “A hypothesis on the accelerating decline of reading: * Broadly speaking, people read for pleasure/entertainment and for learning/obtaining information,” comment on Arvind Narayanan and Sayash Kapoor, “AI as Normal Technology,” *AI as Normal Technology*, Substack, May 22, 2025a.

²³ Narayanan, 2025a.

²⁴ Greg Lukianoff and Cosmos Institute, “Will AI Kill Our Freedom to Think?” *The Eternally Radical Idea*, Substack.

²⁵ Manish Raghavan, “Competition and Diversity in Generative AI,” arXiv, arXiv:2412.08610, December 11, 2024, p. 1.

²⁶ Konrad Lachmayer, “AI, Plurality and Democracy: Reflections on the Impact of Large Language Models Like ChatGPT on the Rule of Law and Democracy,” in Pablo Riberi and Konrad Lachmayer, eds., *Political Representation, Democracy and the Constitution*, forthcoming.

²⁷ Some relevant additional studies include Barrett R. Anderson, Jash Hemant Shah, and Max Kreminski, “Homogenization Effects of Large Language Models on Human Creative Ideation,” *C&C ’24: Proceedings of the 16th Conference on Creativity & Cognition*, June 23, 2024; Anil R. Doshi and Oliver P. Hauser, “Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content,” *Science Advances*, Vol. 10, No. 28, July 2024; Vishakh Padmakumar and He He, “Does Writing with Language Models Reduce Content Diversity?” arXiv, arXiv:2309.05196, September 11, 2023; and Peter S. Park, Philipp Schoenegger, and Chongyang Zhu, “Diminished Diversity-of-Thought in a Standard Large Language Model,” *Behavior Research Methods*, Vol. 46, January 2024.

²⁸ For example, see Tom Brown, “The Future Belongs to the Curious,” *Marigold Montessori*, Substack, June 19, 2025.

²⁹ Mokyr, 2018, p. 35.

Competitive Diversity and Pluralism

Many of the most prominent historical cases of vibrant, influential great powers reflect some version of the same social profile. In addition to providing opportunity for a larger proportion of their citizens than competitors, these dynamos tended to benefit from complex and varied populations and political structures. How this looked varied throughout history, but the most-competitive nations tended to boast a significant amount of *diversity* (in many forms, not just the categories most commonly associated with that term today) and some degree of *pluralism* (a variety of competing power centers and levels of governance).¹

This phenomenon can be seen as far back as Ancient Rome, which was remarkably heterogeneous and open to people from its conquered provinces (to the point that some emperors were foreign-born men) and had a highly federated governing structure of provincial governors enjoying immense autonomy. The same pattern cropped up in Italian Renaissance city-states, which, although not broadly diverse in modern racial, ethnic, and gender terms, nonetheless embodied a social structure that integrated the skills and perspectives of businesspeople, military leaders, artists, writers, craftsmen, and many other professions and that was fairly open to people from outside their borders—and actively competed to attract the best of them. The modern United States, with its combination of widening of opportunity to a broader range of the population, a population continually leavened with immigrants, and a highly pluralistic system of governance, may reflect the leading example of these benefits.

The historical record supports a general proposition: Nations that enjoy many distinct sources of ideas and human capital and a proliferation of governing levels have a competitive advantage.² But this characteristic comes with more qualifications than the others. Many ethnically, racially, or linguistically homogeneous nations have pushed to the top ranks of great powers and achieved tremendous competitive leaps, as China is doing. Some nations with very strong central governments and less-obvious pluralism have done well. It's also easy for diversity and pluralism to run out of control and collapse into a disruptive—and even hazardous—degree of social disunion.³ In organizations and societies, diversity creates tensions and challenges, even as it offers benefits.⁴ Some research on the direct economic value of diversity *per se*, in the absence of other factors, has found mixed results.⁵

Still, on the whole, the evidence supports the idea that social variety, in the form of various kinds of diversity and political pluralism, offers long-term competitive advantages, which I describe in this chapter.⁶ Social variety *alone* doesn't deliver competitive advantage in the form of improved innovation, economic performance, or any other measure. This quality requires support, including such things as shared opportunity to lever out the ideas and talents of a diverse population and effective institutions to regulate the operation of a diverse and pluralistic society. But variety in social and political composition seems to enhance the engine of ideas, talent, and willpower that underpins competitive advantage, especially over the long term.

There are ways in which AI could strengthen the benefits of social variety, but its effects could also exacerbate their fragmenting risks. And the United States could confront a bizarre possibility in this domain: whether AI will bring U.S. society the power to achieve the advantages of diversity and pluralism without their reality.

Defining the Characteristic of Competitive Diversity and Pluralism

This characteristic combines two related but somewhat distinct elements. The first is *diversity*, understood in very broad terms—not just the categories that have become politically contested. As I put it in the 2022 RAND study, I conceived of this quality for the purposes of that analysis as

the level of overall variation in a society—not only in terms of gender, race, ethnicity, sexual orientation, and so on but also in every other way that diversity can be defined. This can include a wide range of education, training, career paths, and skill sets; people from geographically distinct parts of the nation with different cultural traditions and even languages; and people with divergent major experiences (military service versus creative arts, for example). The inverse of diversity would be a largely homogenous society.⁷

The second component is *pluralism*, a quality that has both political and social aspects. Politically, pluralism refers to the degree to which a society “has overlapping sources of authority, rulemaking, and governance and the degree to which its people value, and tolerate, multiplicity.” More broadly, in encompassing social terms beyond government structures, pluralism highlights “the degree and health of civil society, strong and multifarious public and private institutions”; a healthy pluralism “creates an environment in which people tolerate differences in allegiance or membership, rather than viewing other social groups with suspicion and disdain.”⁸ In both senses, the concept refers to contests for power among multiple interest groups in society, none of which dominates policy outcomes and all of which are pursuing their own independent goals. Pluralism thus has two components: One is structural, in the levels of governance that exist in a society; and the other is procedural, in the opportunities that those overlapping structures give for people in the society to claim rights and opportunities. Both embody a collision of preferences and power among interests in society.

The essential quality captured by this characteristic as a whole has to do with societies that embody multiple sources of ideas, innovation, and other elements of national dynamism. Nations with high levels of social variety have richer and more expansive sources of technological, economic, social, and military energy. Once that landscape is in place, a potent degree of *competition* among social actors is a big part of the mechanism by which diversity and pluralism deliver competitive benefits.⁹

Multiple historical case studies demonstrate how diverse and pluralistic societies gain competitive advantage over more homogenous and centralized ones. The RAND study cited significant modern evidence for the competitive values of diversity by creating an environment that generates richer and more creative ideas, for example. In terms of pluralism, the historian Walter Scheidel has pointed to the value of a more dispersed power structure, such as emerged in Europe’s competing state system in the period after the Roman Empire. Such fragmented and pluralistic contexts absolutely had significant short-term costs of instability, such as recurring wars among the countries. But it had profound long-term advantages in creating a competitive “patchwork quilt” of nations driving their own development. As Scheidel puts it,

So many different power structures intersected and overlapped, and fragmentation was so pervasive that no one side could ever claim the upper hand; locked into unceasing competition, all these groups had to bargain and compromise to get anything done. Power became constitutionalised, openly negotiable and formally portable; bargaining took place out in the open and followed established rules. However much kings liked to claim divine favour, their hands were often tied—and if they pushed too hard, neighbouring countries were ready to support disgruntled defectors.¹⁰

Scheidel refers to the “deeply entrenched pluralism” that came to characterize these societies and continued to influence their evolution for centuries. In assessing sources of technological and economic advance, Scheidel concludes, “Pluralism is the common denominator.” Progress “was born in the crucible of competitive fragmentation.”¹¹

AI has the potential to alter how much a nation gets out of the benefits of social variety. And it could even lead to a far more radical outcome: Providing the basis for synthetic versions of diversity and, to a lesser degree, even pluralism that mimic the competitive advantages of these features, even in countries that don't truly embody them. The key, as with all seven characteristics, is what conscious choices we make to bend the arrow of AI's effects toward positive outcomes.

AI and Diversity

AI could interact with the societal characteristic of diversity in many ways. For example, advanced AI could theoretically provide important fuel to empower people from diverse groups across society. If AI creates more shared opportunity throughout a society through the mechanisms discussed in Chapter 7, almost by definition, it should produce more chances for citizens from a variety of diverse groups to express their talents. By providing new capabilities to a wider variety of people, as described in Chapter 7, AI could send a burst of energy through a multilayered and varied social structure to generate innovative energy that would benefit the nation. It could provide policymakers with the tools for ensuring that they consider a variety of perspectives, such as giving ready and quick access to various views and offering new means for groups in a society to preserve and promote their cultural, ethnic, or religious heritages.

There are ways in which AI can mitigate the inherent risks of diversity and pluralism. In Chapter 6 on shared national identity, for example, I discussed the potential of AI models to conduct mediation and conflict resolution efforts, which might dampen the entropic effects of diversity and allow nations to sustain a unified identity even while capturing the value of high levels of diversity. If AI were to create new opportunities for many different groups in society, it would improve the social mood and reduce the potential for diversity to cause problems. If it increases the effectiveness of governing institutions and they can better serve the people, it will ease the chance of grievances in areas or groups within the country.

But there are equally powerful risks. As I argued in the previous chapter, AI models operating in similar ways and grounded in the same training data carry a risk of homogeneity—LLMs tend to generate ideas, strategies, and insights that look similar to one another. If nations employed LLMs in typical ways, a nation relying heavily on AI throughout many economic, social, and military applications could unwittingly wipe away a good deal of natural variation in ideas, thinking, innovations, and perspectives, even if the underlying measures of diversity in society remained very significant. A homogenizing AI would weaken the natural advantages of diversity and potentially create a nation that is more predictable and less creative.

Another risk is that AI could exacerbate biases in hiring, promotion, sentencing, university admissions, and other areas in ways that discriminate against certain groups in society and, thus, dampen the value of diversity. We've already seen this effect in dozens of examples and studies in which AI, relying on training data filled with hallmarks of bias, has generated prejudiced outcomes.¹² Some studies have even shown that people internalize the biases fed to them by AI, and those prejudices persist even after AI interactions.¹³ If trained and designed the right way, AI could theoretically help *solve* bias in various ways, but, apparently, the more default initial effect could be to reflect biases found in its training data. But experience, so far, suggests that biases have a way of creeping back into model outputs, even after extensive retraining.

One fascinating and slightly bizarre connection will be AI's potential to generate synthetic social variety that captures the idea-generating advantages of the characteristic—without the actual people. AI could simulate the perspectives of people from different backgrounds, with different skills and from different regions. It might be possible to turn a model loose on some economic or social issue and get back hundreds of insights and ideas that mimic the intellectual fruits of actual human diversity. It would simulate not only the views of different genders, races, or ethnicities but also people in different industries, in geographic regions, of different ages and political persuasions, and just about any other distinction.

The risks here are obvious. If AI can simulate variety, societies may start to believe they no longer need the real thing—clinging to cultural homogeneity or even forcing it, while pretending nothing is lost. The same dynamic could play out locally, as self-sorting produces ever more uniform communities and as firms in those places boast of diverse innovation powered by algorithms. It's a depressing trade, and one unlikely to promote competitive advantage: The real lessons and resilience of human difference—and in particular all the social interaction and learning that both produces and draws benefits from that difference—would be swapped for a counterfeit imitation. Such a diversity-simulating AI tool could also narrow the competitive playing field between open, tolerant, diverse democracies and their more-homogenous autocratic rivals. If such autocracies as China can just manufacture millions of reasonably accurate diverse perspectives in their data centers, they could cancel out this advantage of more heterogeneous and open societies.

These risks once again point to that central theme, and danger: The ways in which AI reshapes the essential nature of society are an immense wildcard in evaluating its competitive effects. The advantages of diversity, for example—such as learning and adaptation, effective institutions, and many other characteristics—flow from the authentic human social processes involved. Diversity forces societies to deal with a complex population and differences of profession, values, ethnicity, background, and views and, in the process, to acquire the hard-won resilience of a people who have learned to manage their differences. That sustained, sometimes harrowing social process is not something AI can simulate. The more it thins out the need for and practice of such social engagement, the more it will undermine motivation, resilience, and solidarity. Here we see a central reason why the main competitive challenge of AI is social, not technological.

The bottom line is that if we're concerned about social variety, broadly defined, as a source of competitive advantage, AI could, in theory, be an important tool to unlock even more of this quality. But it is equally certain that AI will create dynamics that tend to suppress the benefits of a true variety of ideas and perspectives through a certain constraining homogeneity. And there can be no question that some organizations and countries will try to manufacture the opposite effect—using AI to mimic diverse perspectives without the authentic reality of them. The result will likely be to further muddy the waters of diversity's competitive effects.

AI and Pluralism

Nations with a more pluralistic governing structure—in the U.S. case, for example, its highly federated model of national, state, local, county, and municipal governments—can gain competitive advantage through experimenting with governing approaches to different issues, sharing good ideas and best practices, and delivering services in a more local and tailored way, for example. Advanced AI could empower a layered governing structure in many ways.

First, AI offers, as I discussed in the previous chapter about learning and adaptation, the potential to fuel unprecedented numbers of experiments in many domains, potentially including public policy. It could generate this result in various ways—generating new ideas, finding new insights in policy-relevant data, simulating some outcomes in more-controlled contexts, and evaluating real-world trials. State and local governments could test the potential of a dozen new health care delivery systems, mental health solutions, or public transportation options and give at least some initial clues as to which approaches might be most useful. The potential gain here is to not only identify policies that end up being more effective but also shift the mindset of public officials at various levels of governance. When the opportunity cost of trying out all kinds of different ideas radically declines, some of the apparent political and bureaucratic obstacles to thinking in new ways might also decline.

Second, effective institutions will make for more effective pluralism. If public institutions empowered with AI demonstrate more effectiveness and efficiency in their operations and, as a result, their public legiti-

macy recovers and this effect happens across the spectrum of pluralistic governing institutions, we'll end up with a system thriving at multiple levels. That's the intention of a small but growing number of institutional entrepreneurs seeking to make public-sector bureaucracies work better.

A political community wouldn't need everyone to begin applying AI at the same time if it has a critical mass of leaders and agencies, from the local to the federal level, starting to experiment and achieve notable results. As they begin to do so, citizens and stakeholders in other areas and served by different agencies will point to emerging examples and demand similar improvements. The combination of AI models capable of generating all manner of tests and trials with an ongoing process of dynamic experimentation in a competitive pluralistic system could be a powerful engine of competitive advantage.

AI's flexibility makes it a natural ally of pluralism. A city could customize an open-source model with local data, creating tools tuned to its own problems and possibilities.¹⁴ In the future sketched in Chapter 3—one of countless models trained on different datasets—such variety could thrive. A pluralism of models matched by pluralism in governance might look chaotic but could spark immense creativity. Some scholars even propose using multiple models together to ensure diverse perspectives in decisionmaking.¹⁵

The opportunities for AI to supercharge the benefits of pluralism are very real. However, so are the risks. The leading one is the way in which AI—through its damage to social solidarity, the public sphere, and perceptions of social equality—could accelerate the divisive potential of both pluralism and diversity.

A Leading Risk: AI-Fueled Social Divisions

This risk will be hard to avoid because AI is sure to exacerbate the risk of social disintegration inherent in these two factors. Diversity can become toxic when it hardens into mutually suspicious social groups, whether built around ethnicity, income level, race, political ideology, or belief in some very specific conspiracy theory. But even pluralism can serve as a form of social entropy in times when the fragmenting tendencies of a society become so strong that various governing entities at different levels begin to claim dominant political loyalty from the citizens in their area.

These effects could emerge in several forms, most notably by AI serving as an accelerant of our shattered media landscape, hardening social divisions as I discussed in Chapter 7 and weakening the bonds of social solidarity as discussed in Chapters 6 and 8. And with the radical possibilities of AI, these destabilizing effects will take bizarre new forms. AI models have already started to invent their own religions, for example, and a few observers are worried about the notion of a social landscape populated with thousands of AI-generated faiths and ideologies. Moreover, if AI ends up benefiting a few people rather than providing shared opportunity, it is likely to worsen the socioeconomic divisions that could be expressed through other fissures in a society. All this will play out on top of a U.S. society—and many developed societies—in which political polarization and social fragmentation are already dangerously advanced.

Social fragmentation and alienation have driven people to seek solace in identity-based movements—a trend that is creating social tensions in the United States and dozens of other countries. If AI models start generating religions and ideologies, we might face hundreds of competing belief systems vying for followers, potentially inciting disruptive or even violent action among their adherents. How bad this gets will depend partly on the quality of our institutions. Because of their role in creating binding effects throughout society and in mediating and facilitating social relationships, effective institutions can go a long way to mitigating the potential noxious effects of out-of-control diversity and pluralism.

Again, we come back to the fundamental theme of this study: agency. Every characteristic demonstrates the ways in which AI's promise and peril both turn on how it affects human and collective agency—how it empowers citizens to act and communities to shape their destinies. Individual agency links to opportunity and learning; collective agency lives in willpower, identity, and the active state. Much depends on whether AI

amplifies or erodes the pluralism that gives agency its depth. If it fuels genuine autonomy across a vibrant, varied society, it will strengthen a nation's power.

Summary: Here Come Some Volatile Feedback Loops

In this chapter, I have aimed to show how advanced AI, diversity, and pluralism will interact—both with one another and with the other characteristics explored here—in unpredictable ways. Combine frontier or even superintelligent AI with diversity, shared opportunity, adaptive institutions, and a strong national identity, and the result could be a powerful feedback loop: a society energized at every level of creativity and coordination, capable of turning its collective intelligence into decisive action. That is the ideal outcome. But events could take a darker turn: a deeply divided society whose sense of solidarity is collapsing while AI is being used to make many people's efforts and dignity largely irrelevant, all of it shackled by a hyper-bureaucratized version of AI models running just about everything in mysterious ways. The society that would spill out from the far end of such trends would be weak, divided, and indecisive. It might still profit from benefits of AI in the working of institutions in many ways, but, as a national actor, it would be a pallid ghost of the flourishing version produced by the more optimistic feedback loop.

Reality will likely fall somewhere in between those extremes. No technological revolution tends to immediately create feedback loops in these broader ways that are either wholly advantageous or damaging. Great powers will be grappling with one another, standing on mixed versions of all these trends; the question is which ones end up with the best overall profile. Getting there—to reiterate, once again, the core thread running through this whole work—demands an awareness of these opportunities and risks, a conscious choice to shape trends for the better, and, most of all, the political will to make some tough choices.

The profound truth at the core of so many of these effects is that the nature of human reality—our interaction with facts and phenomena, our social patterns, our ways of thought and work and craft—are going to begin evolving rapidly under the influence of AI. AI is in the process of creating a different *kind* of society. We have an opportunity to influence what it will look like, but only if we choose to do so, overcome our political cleavages to discover shared paths forward, and make the necessary investments, and, in some cases, hard choices to make that happen.

With this factor, my journey through the seven leading characteristics identified in the 2022 RAND study has come to an end. In this chapter and the previous ones, I have offered a whole constellation of specific findings and a few overarching themes. Those can provide some of the basis for initial hints of a U.S. strategy for competitive advantage in the AI Era, which I discuss in the final chapter.

Notes

¹ Mazarr, 2022a.

² Two classic statements of this argument are Richard Florida, *The Rise of the Creative Class—Revisited*, 2nd ed., Basic Books, 2012; and Scott E. Page, *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton University Press, 2007.

³ One important book chapter—which emerged after this report was completed and would have offered important qualifications and nuance to my thesis—is Joel Mokyr, “Diversity, Pluralism and Tolerance: The Roots of Economic Progress?” in Sidney M. Milkis and Scott C. Miller, eds., *Can Democracy and Capitalism Be Reconciled?* Oxford University Press, 2025. It is a highly useful guide to thinking through the advantages and challenges of diversity and pluralism, and its basic conclusions closely match those of this analysis, especially in terms of the value of a competitive pluralism.

⁴ Adam D. Galinski, Andrew R. Todd, Astrid C. Homan, Katherine W. Phillips, Evan P. Apfelbaum, Stacey J. Sasaki, Jennifer A. Richeson, Jennifer B. Olayon, and William W. Maddux, “Maximizing the Gains and Minimizing the Pains of Diversity: A Policy Perspective,” *Perspectives on Psychological Science*, Vol. 10, No. 6, 2015. For an interesting statement about diversity’s challenges, see Resul Cesur and Sadullah Yıldırım, “The Misery of Diversity,” National Bureau of Economic Research, Working Paper No. 33163, November 2024.

⁵ For example, see Alberto Alesina and Eliana La Ferrara, “Ethnic Diversity and Economic Performance,” *Journal of Economic Literature*, Vol. 43, No. 3, September 2005; Günter K. Stahl, Martha L. Maznevski, Andreas Voigt, and Karsten Jonsen, “Unraveling the Effects of Cultural Diversity in Teams: A Meta-Analysis of Research on Multicultural Work Groups,” *Journal of International Business Studies*, Vol. 41, 2010; and Günter K. Stahl and Martha L. Maznevski, “Unraveling The Effects of Cultural Diversity in Teams: A Retrospective of Research on Multicultural Work Groups and an Agenda for Future Research,” *Journal of International Business Studies*, Vol. 52, 2021.

⁶ For a survey of the evidence for organizational advantages of diversity, see Vivian Hunt, Sara Prince, Sundiatu Dixon-Fyle, and Lareina Yee, *Delivering Through Diversity*, McKinsey and Company, January 2018.

⁷ Mazarr, 2022a, p. 239.

⁸ Mazarr, 2022a, p. 239.

⁹ Mazarr, 2022a, p. 240.

¹⁰ Walter Scheidel, “The Road from Rome,” Aeon, April 15, 2021.

¹¹ Scheidel, 2021. Some observers have pointed out that this relationship may be weakening at a time of continuing European fragmentation counterposed to seeming Chinese central direction. See Martin Wolf, “Can a Fragmented Europe Continue to Prosper?” *Financial Times*, November 11, 2025.

¹² James Manyika, Jake Silberg and Brittany Presten, “What Do We Do About the Biases in AI?” *Harvard Business Review*, October 25, 2019; Reva Schwartz, Apostol Vassilev, Kristen Greene, Lori Perine, Andrew Burt, and Patrick Hall, “Towards a Standard for Identifying and Managing Bias in Artificial Intelligence,” National Institute of Standards and Technology, NIST Special Publication No. 1270, March 2022.

¹³ Lauren Leffer, “Humans Absorb Bias from AI—And Keep It After They Stop Using the Algorithm,” *Scientific American*, October 26, 2023.

¹⁴ For example, see Nandhini Swaminathan and David Danks, “AI, Pluralism, and (Social) Compensation,” arXiv, arXiv:2404.19256v2, October 15, 2024.

¹⁵ Joshua Ashkinaze, Eric Gilbert, and Ceren Budak, “Plurals: A System for Pluralistic AI via Simulated Social Ensembles,” *38th Conference on Neural Information Processing Systems (NeurIPS 2024)*, October 2024.

A Strategy for National Advantage in the Artificial Intelligence Era

A revolution is coming. As I argued in Chapter 3, it may be coming more gradually than some advocates suggest. There may be very big bumps and hurdles: Even in the next year (2026), we may see dramatic stock market corrections in the AI domain and continued suggestions that some existing scaling approaches are running out of steam or that investors are easing off the gas pedal of the gargantuan private investments in the AI ecosystem. But the AI Era is surely on the way: AI will have profound effects on our economies, societies, politics, security, and psychology within a decade, probably even sooner. My thesis in this study has been that the nations that will benefit from this transition will be those that manage the societal effects of this tidal wave in ways that make them more dynamic, unified, and resilient. Maximizing national competitiveness in the AI Era is more of a social than a technological challenge.

In dealing with that challenge, whatever the deepening flaws of American society, the United States has an incredible foundation to build on. The United States of the post-World War II era represents possibly the richest menu of social advantages in history. Fresh from a massive national war effort and looming over the global economy, it had a powerful sense of a mission, willpower, and a clear and unified national identity. It had impressive and growing levels of shared opportunity and a state that was actively contributing to competitive advantage and effective institutions. But perhaps more than anything, the United States in this period reflected a positive-sum interaction *among* all these characteristics. Its ambition spurred its active state. Its strong institutions made opportunity possible. Its pluralism and rising diversity impelled greater creativity, learning, and adaptation. The United States in its heyday did not merely embody a collection of discrete competitive factors—it was a brilliantly integrated competitive engine.

As I reviewed the characteristics of national competitiveness in the 2022 RAND study, it became more and more apparent that all the qualities I discovered, while important on their own, are also deeply interwoven. Shared opportunity needs effective institutions to work and is more powerful in diverse and pluralistic settings. National ambition requires a unified identity to be sustainable, and both shared opportunity and national ambition could be expressed through an active state—and so on. Every characteristic feeds every other one to some degree.

Both the best and worst cases of national competitive standing tend to reflect such synergy and feedback loops among these characteristics—either favorable or disastrous. In the successful examples, such as the post-1945 United States, each quality fed into the others in ways that created an effect that was greater than the sum of the parts. In the ruinous cases, the opposite dynamic emerged: Ineffectual institutions made it impossible to achieve shared opportunity, and the entropic push of excessive diversity and pluralism fragmented national solidarity, which undermined national identity; combined with a weak state, these factors crushed national will.

Classic examples of such destructive spirals can be found in history's great fallen dominions—the Spanish, Ottoman, and Austro-Hungarian empires. But other cases provide equally powerful warnings of the potential for burgeoning national greatness to stall out and plateau.¹ One such case is post-

Revolutionary France. Rising from the chaos of revolution, Europe's hegemonic land power enjoyed all the seeming preconditions for lasting competitive dominance: a vast land mass, large population, and an economy that had grown to nearly match Britain's and was substantially bigger than Germany's. The French Revolution unleashed the forces of popular willpower and levels of military manpower that dominated those of its neighbors.

France was also positioned, in theory, to lead the emerging Industrial Era. It boasted many of the leading scholars and institutions of the Enlightenment, and its scientific community stood at the forefront of world progress. Socially and politically, France also seemed to be equipped for competitive dominance. Napoleon introduced rationalizing reforms that enhanced the effectiveness of public administration, learning, and adaptation and broadened opportunity by abolishing aristocratic privilege and opening many fields to the average French person.

But it failed to capitalize on that opportunity. In the decades that followed Napoleon's final defeat in 1815, France fell further and further behind Britain in the drive to industrialize and modernize their respective economies and societies. France remained a big power and would later make a surge into colonialism that gave it worldwide reach. But it would never again contend for the role of regional hegemon or make a bid for global manufacturing or scientific-technological leadership. Its brittle social and political system—weakened by mutually reinforcing social diseases, from a loss of national willpower to ineffective governing institutions—was shattered by the German invasion in World War II.²

Britain, once the apex of world competitive standing, offers another cautionary tale in its failure to sustain its competitive position through the Second Industrial Revolution. British society had a certain competitive advantage relative to most European countries in terms of earlier shared opportunity but was then bypassed on that count by far less aristocratic, educationally hidebound, class-bound societies, such as the United States. Smaller British family firms, well attuned to the opportunities of the First Industrial Revolution, could not adjust to the immense scale required for the mass production era of the Second Industrial Revolution. Over time, barriers to innovation and intellectual energy grew around the British economy, so much so that worried observers coined the phrase “the British disease” to describe the pattern.

Even the most potent great power's competitive advantage can evaporate if it gets on the wrong side of a destructive feedback loop among decaying societal qualities. The United States faces this risk, even apart from any AI-generated transition. However, as we have seen, the advent of an AI Revolution over the next decade has the potential to exacerbate—or ease—many dangerous trends underway today. What's required is nothing less than a broad dedication to national renewal, partly through the careful management of emerging AI capabilities and the use of AI for advantage. This is a profound and daunting agenda, requiring a campaign in the public and private sectors alike.

An effective strategy for national competitive strength in the AI Era will not be primarily about the technology but instead be determined by U.S. society's ability to adjust to and benefit from the era that is emerging. U.S. society is starting from an erratic foundation and is already being destabilized by a whole phalanx of dangerous trends—inequality, a poisoned public sphere, and stifling and predatory bureaucracies. The United States faces an urgent requirement to promote new levels of dynamism in a society suffering from manifold ailments—whether or not transformative AI arrives on schedule.

We don't need an AI strategy. We need a strategy for national renewal and competitive advantage that uses the capabilities and opportunities of the AI Era to achieve its goals.

The Route to Competitive Advantage: Fostering Individual and Collective Agency in the AI Era

Perhaps the most important and overarching lesson of my analysis is that AI will empower nations over the long sweep of history—enhance their dynamism, willpower, solidarity, innovativeness, and resilience—to the degree that it broadly underwrites human agency in individual and collective terms. For decades, the literature on AI has noted that the emergence of superpowerful AI will challenge our understanding of what it is to be human. But the lesson of my analysis is that nations that don’t rise to that challenge will not only suffer moral and ontological crises—they also risk falling behind in relative power.

The concept of autonomous agency doesn’t directly capture every form of advantage mentioned in previous chapters.³ For example, AI’s effects on efficiency and productivity can be more technical and operational than its effects on the agency of workers. Yet I am convinced that the most central, essential, ultimately decisive forms of national competitive advantage in the AI Era will depend on whether nations work to empower individual and collective agency—and that AI can shape this core issue, for good or ill, in many powerful ways. It can reorient the power of individuals relative to large organizations, the degree of agency in economic contexts (as in, for example, the power of labor versus capital), human epistemological agency in forming accurate understandings of the world, human creative agency in various fields, and collective human agency in the form of political communities empowered by active states.

These advantages reflect the lessons of the Industrial Revolution. In those transformative decades, nations could gain a degree of competitive advantage merely by incorporating the new industrial technologies and processes: spreading the infrastructure of electrical power and railroads, deploying communications technology, erecting factories, and boosting their manufacturing output. But the vast majority of assessments of the rise of the West from 1500 to 1800 stress various dynamics related to human agency: the power of intellectual and scientific ambition, the associated values of a strong work ethic on the part of individuals and groups, the role of individual and collective pursuit of innovation as disciplined but also prompted by the scientific method, the energy devoted to the practical application of useful knowledge, the role of self-interested forces in society constraining collective agency, the creative power and drive of an essential creative minority in society, and the role of a bourgeois ethic that encouraged risk-taking in pursuit of profit and ambition.

Each component of agency and the support systems for it can become part of a holistic ecosystem of dynamism. A virtuous cycle built on individual and collective agency—and the healthy social patterns and habits that both set the context for and flow from such agency—may well end up serving as the fulcrum of the dynamic that separates successful societies from lagging ones.

Brendan McCord of the Cosmos Institute has been promoting a humanistic vision of an AI Era—one in which AI empowers human agency and flourishing rather than facilitating technological control. One of his critical insights is that a deep study of philosophy is a critical component of any effort to make the AI Era a human—and humane—one.⁴ In particular, he worries about AI’s threat to the self-directed career. Algorithmic decisionmaking could focus increasingly on optimizing skills or capacities that channel people toward what they can best do rather than what they choose to do. As he has argued:

Might we be ushering in a world where algorithms don’t just carry out our instructions but increasingly shape how we decide what’s good for us—suggesting our next word, our next action, our next career move, our next purpose . . .

Each small delegation of choice will likely seem harmless, even natural. But together, will these micro-abdications of judgment habituate us—choice by choice, day by day—away from self-governance and toward passivity and dependence?

I call this looming possibility (but not inevitability), “autocomplete for life.”⁵

McCord emphasizes the importance of a decentralized approach to AI—one that empowers many actors, including individuals, small organizations, and local governments, to use it to their advantage. He offers two key questions to serve as criteria for a humanistic vision of AI: “How do we create systems that serve people whose needs and purposes we cannot know fully in advance?” and “How do we enable those individuals to deliberate autonomously and well about their purposes?” He argues that distributed systems harnessing knowledge to spur emergent, spontaneous order are essential to any such vision. Only a decentralized approach can manifest self-direction on the part of autonomous individuals, he argues. He adds,

AI is our new epistemic infrastructure. It could take that project to new heights, or subvert it entirely. In the downside case, AI becomes an “autocomplete for life”—suggesting not just our next word, but our next action, job, relationship, purpose. Each small delegation of choice seems harmless, even natural. But together, these micro-abdications compound—choice by choice, day by day—gradually diminishing our capacity for autonomous thought.⁶

These aspects of individual and collective agency provide clues to elements of a strategy that will help the United States align itself to the societal requirements of succeeding in the AI Era.

Essential Components of a U.S. AI Strategy

To achieve the core U.S. objectives for national competitiveness outlined in Chapter 1, I propose six main elements of any comprehensive U.S. AI strategy. They are necessarily broad and somewhat abstract. The precise details of how the resulting strategy might play out will have to wait for further work. But I can offer some general outlines.

Importantly, the goals do not require U.S. predominance in overall AI development—only competitive equivalence. My assumption is that China (and eventually perhaps other actors) will reach a broad parity in many indices of the AI technology stack, including model capabilities or rankings and sufficient available compute and data center capacity. I do not think that the United States can feasibly set the goal of preserving a permanent advantage in AI application or diffusion because the Chinese government and society have too many tools, qualities, and incentives—and because potential roadblocks in the United States are too serious—to allow for a persistent U.S. lead over a decade or more. With these stipulations, the six components of a possible U.S. AI strategy are the following:

1. **Sustain a lead or strong equivalence in AI model development and capabilities, both proprietary and open models.** This is primarily a private-sector job, but the government can help with permits and regulatory action, export controls to protect the U.S. lead in compute, cybersecurity for AI labs, targeted investments (such as for open models), and in other ways.
2. **Develop essential supporting infrastructures and technologies that allow AI to have its full competitive effect.** This includes power generation, data centers, robotics, and other elements of the wider ecosystem for AI’s economic and social capabilities.
3. **Promote the broad adoption of AI capabilities and applications in the public and private sectors.** This area includes such things as providing public-sector experimentation funds, creating AI application sandboxes to allow trial use by private-sector firms, and implementing training programs to widen the AI skill base.
4. **Encourage U.S.- and allied-developed AI models, networks, infrastructure, and standards as the baseline technology stack for global AI development.** This element of the strategy could involve subsidized access to AI models through application programming interface (API) access or subsi-

dized data center usage, investment guarantees for U.S. firms considering AI-related network projects in key countries, the promotion of U.S. AI standards, and other steps.

5. **Identify and assure specific thresholds for AI safety** through technical and legal and policy mechanisms, and develop plans for possible AI safety and alignment events.
6. **Develop and implement a broad national effort to promote the beneficial effects of the AI Revolution on national dynamism, coherence, stability, and willpower**—or, as I have noted, a national strategy for renewed dynamism using AI as a central tool and approach.⁷

Of these components, the vast majority of attention in U.S. AI discussions has been on the first four. Most discussions of AI strategy, including the Trump administration’s July 2025 AI strategy document, focus on accelerating AI progress, building the necessary surrounding technical and infrastructure ecosystem, and, to a lesser degree, taking steps to promote AI diffusion, AI application, and U.S. leadership in global AI networks. Abundant policy ideas and government policy-planning exist in all those areas.

The fifth objective is especially important—though it remains beyond the scope of this analysis. Any vision for success over the next decade must include effective steps to ensure AI safety, alignment, and accountability. A better future will emerge only if policymakers control AI’s risks to humanity in general. Losing the U.S. competitive position to China is one thing; losing it to an arbitrary, sometimes destructive, out-of-control AI would be far more dangerous to humankind. The goal of AI safety is not only essential for humankind but also a potential source of competitive advantage: In the contest between the United States and China, if one side’s models begin badly misbehaving, the world might shift to the other as a safer alternative. U.S. progress will be impaired by popular outrage over misaligned or dangerous AI.

Expert observers have suggested various routes to safety.⁸ The choices involved are complex and difficult: AI safety measures can trade off against the speed of model development, and much doubt still exists about which routes to safety will actually work. The specifics of a potential approach to AI safety are beyond my scope of analysis—my focus is on the requirements for national competitiveness. But some conscious effort toward this goal is an essential component of any AI strategy.

An important part of that safety and alignment priority is anticipatory planning. Governments and private-sector actors need to begin defining and thinking through possible loss-of-control and misalignment events and developing detailed plans for how to handle them. We’ll never forecast all possible risks. But we need to do the best we can—assisted by AI itself—to be ready to act quickly and decisively in what are sometimes termed *break glass* moments.

The first four of these objectives are, to a greater or lesser degree, part of current U.S. policy, including the 2025 AI Action Plan. That focus is important and essential. But my argument is that long-term national success beyond the direct AI technology competition will be decisively shaped by the sixth and final goal of U.S. AI strategy, to manage the effects on the societal foundations of national competitive advantage. In the following sections, I focus on the possible building blocks of success in that critical goal. Doing this right will demand a wide-ranging movement for national dynamism and renewal, the likes of which hasn’t been seen in roughly a century. One message of this analysis has been that we confront, at the same time, an unprecedented moment of economic, social, political, epistemological, and psychological disruption in societies and the early phases of a massive technological transition. The stakes are massive. What’s required is nothing less than the sort of bottom-up commitment to national renewal that great powers have only been capable of a few times in modern history.

RAND’s earlier work on societal competitiveness suggests one broad rule to apply to these initiatives: sustaining a thoroughly grassroots flavor to the development of AI rather than imposing central control—and beyond that, using AI to energize a broader decentralization of social and political problem-solving. An era of overpowering technological power and intimidating degrees of simulated reality can be mitigated if people

feel engaged in small-scale governing units that have significant authority over many aspects of public policy. To sustain the highest degrees of national competitiveness, we need to make sure that AI doesn't result in centralized mandates and arbitrary bureaucratic control—either within organizations or within nations—that exacerbate the existing crisis of legitimacy. Any U.S. strategy must have a core philosophy of a portfolio of experiments rather than imagining there's a dominant avenue. Jeffrey Ding has argued,

Decentralization often correlates with higher diffusion capacity in science and tech. Instead of picking winners and locking in a particular trajectory, a decentralized ecosystem enables diffusion from the bottom up because the most successful trajectory is allowed to emerge.⁹

AI could become a powerful ally of localism in the AI Era. Properly developed and used, it could give smaller communities the means to understand and act on their own problems and, particularly, to tailor solutions to their specific circumstances. This aspect of AI and governance is critically important and has not yet received enough attention: AI models can train on local or regional data and experience to offer specialized answers. It can also breathe new life into civic groups and help individuals unlock their creative energy. Without this kind of renewal, the forces of social entropy and stagnation will only accelerate, corroding public trust and fraying the bonds of solidarity.

Steps to Nourish the Societal Sources of Competitive Advantage: A Campaign of National Renewal

My core thesis in this publication is that national competitive advantage in the AI Era will go to the strongest *societies*: those that best embody the seven characteristics surveyed in the earlier chapters. Other objectives—staying ahead or roughly even in the technology race, promoting diffusion and application to gain competitive advantage, ensuring AI alignment and safety, and seeking U.S. and allied predominance in global AI networks—are all preconditions for good outcomes. But long-term success must be grounded in the same recipe that has empowered the strongest nations throughout history.

Other countries have begun to put a toe into these waters, looking beyond the glittering power of AI models to the kind of society that exists on the far side of that transition. Estonia has developed a plan for using AI to improve education.¹⁰ Singapore is developing a form of sovereign AI trained more extensively on Asian languages than U.S. models. Courts and legal institutions in several European countries have begun to experiment with AI support systems. Dozens of experiments are underway in developing countries in various fields, including health care, agriculture, and education.¹¹

The steps in the following sections emerge from this analysis as ways to both mitigate the risks of the AI Revolution and seize its opportunities—to bolster the societal foundations of national dynamism. Together, they reflect a strategy for societal advantage in the AI Era. For each principle, I offer the general concept and then one or more specific policy ideas.

One of the problems in defining such a list is that we're very much in the beginning phases of this transition. We don't yet have a lot of practical examples of the needed policies and initiatives to point to. Even when trying to get more specific, I'm left offering hypothetical examples rather than tried and workable solutions. What I outline is not a final or comprehensive agenda, only worthwhile initial moves. I suggest the following ideas, which expand on the findings of previous chapters, as a set of places to get started.

Because many of these areas remain largely unexplored, I turned to a source of fresh ideas to develop some of the specific proposals included: Claude and ChatGPT. I developed each general thematic area and one or two specific ideas to reflect those concepts. I then asked both models to suggest specific policy initiatives. I selected the most-compelling ideas from those results. In some cases, I lightly edited the substance and

format of the outputs. These represent only initial notions—suggestive ideas of what specific policies could promote each objective. But many of them are impressively creative and reflect the emerging capacity of AI to support strategic analysis.

Build Public-Sector AI Competence

Organizing and operating this strategy will require a leading role from government—what I’ve called the *active state*. But there’s a problem. Governments are lagging in the effective application of AI. Such a mounting gap could have dangerous implications in the AI Era. In the most perilous sense, governments wouldn’t be able to keep up with violent nonstate groups seeking to do harm with AI. But more generally, if governments appear to be even more incompetent relative to the opportunities and risks of AI, the result could be a disastrous blow to the legitimacy of public institutions. Mustafa Suleyman draws the obvious implication: “Responding effectively to one of the most far-reaching and transformative events in history will require mature, stable, and most of all trusted governments to perform at their best.”¹² The challenge for the United States is to deal with these dangers and opportunities—not only controlling the risks of frontier models but also identifying and mitigating the downstream dangers and enhancing societal resilience.¹³

This is especially true in U.S. national security institutions. As of this writing, the level of understanding of AI, let alone adoption, throughout the defense and intelligence community remains very much a work in progress. More broadly, some reports, including a 2025 study, indicate growing generative AI use across the federal government, with the numbers of total AI applications growing significantly from 2023 to 2024, although many departments and agencies report regulatory, legal, and bureaucratic barriers to the widespread application of AI tools.¹⁴ One survey of state-level public-sector AI readiness across the United States found a handful of states with clearly developed AI frameworks or strategies but did not document much evidence of implementation.¹⁵ Interest is growing, and federal agencies are moving to expand AI use, but these efforts and the skill levels that support them remain embryonic. The United States needs to rapidly build institutions—likely some form of a central coordinating office for AI strategy—charged with understanding the status of AI development and deployment as much as possible, including a detailed comprehension of the progress being made in U.S. AI labs.

The lesson is simple: The United States needs to get serious about making governments at all levels more conversant with the AI Revolution, more aware of its capabilities and risks, and more capable of using its benefits as AI models begin to roll off the virtual production line in serious numbers. What forms that effort takes—for example, whether the nation needs some sort of federal office for AI implementation—will have to be debated. But U.S. policymakers can’t leave this to chance. A robust effort for AI awareness and diffusion at all levels of government, supported by the top leaders at each level, will be an essential part of the actual implementation of any devised strategy.

Build Public-Sector AI Competence: Specific Proposals

- **Governments at federal, state, county, and municipal levels should create AI strategy offices or functions to oversee the implementation of AI strategies and action plans**—staying abreast of new developments in the technology, identifying risks and opportunities, and coordinating and exchanging information with similar AI strategy offices at various levels. ChatGPT described such a function as a **Federal Office of AI Competence and Implementation** (FAICI) focused on public-sector capacity-building in AI. It would “serve as a hub for AI policy implementation, workforce training, and interagency coordination; track adoption progress across federal departments and advise on best practices; develop model frameworks for AI integration in procurement, data use, and service deliv-

ery; and partner with national labs, think tanks, and AI companies to ensure current awareness of leading-edge capabilities.”^a The existing U.S. White House AI coordinator’s office is an important step in this direction.

- **Governments at various levels should encourage increased skill levels** through subsidized training courses, the addition of AI skill to promotion requirements, bonus structures to reward public-sector employees to discover useful applications, and much more.
- **The federal government should establish an AI Research Projects Agency** to fund research and experiments designed to generate AI applications in areas related to national competitiveness, including medical service efficiency, bureaucracy simplification, cyber defense, biological defense, and breakthrough energy sources. Although this proposal takes inspiration from DARPA, the focus is less on basic science and more on applications. Such an organization should support experiments in AI diffusion, making high-risk, high-reward investments in technologies and techniques that have the capacity to make big leaps and, thus, apply and demonstrate the value of emerging AI applications in ways that promise national competitive advantage.

Claude and ChatGPT suggested the following ideas:

- **Create a government AI Competency Corps**—a new federal personnel system specifically for recruiting, training, and retaining AI-skilled professionals across government. This would include fast-track hiring authorities for AI talent, compensation packages that can compete with private-sector offers, and mandatory AI literacy training for senior government officials.
- **Develop a federated AI Procurement and Security Framework**—a streamlined government-wide system for procuring and deploying AI technologies that maintains security and dramatically reduces bureaucratic barriers. This would include preapproved vendor lists for AI tools, standardized security protocols that don’t require reinventing approval processes for each agency, and AI sandbox environments in which government agencies can safely test and integrate civilian AI software. The framework would enable rapid adoption while maintaining the necessary security controls.
- **Fund AI for Public Good state and local challenge grants.** The goal of such a program would be to encourage U.S. states and localities to experiment with and adopt AI to improve public services—education, transportation, health care, law enforcement oversight, and more. It would consist of competitive grants administered by the U.S. Department of Commerce or U.S. Department of Housing and Urban Development. It would be focused on deployment, not research, through funding for pilot programs, AI tool kits, infrastructure, and training. Such grants would empower state and local governments to become testbeds for AI-enabled governance and build a distributed learning ecosystem across the United States.
- **Undertake a public-sector AI Readiness Index and Incentive Program.** This effort would build a benchmarking and incentives framework to encourage government agencies to improve AI readiness and deployment. It would support an annual assessment of agency-level AI maturity according to such criteria as strategic planning, technical infrastructure, workforce training, implementation track record, and ethical safeguards. Top-performing agencies might receive priority funding, pilot access to advanced AI tools, and public recognition. The goal would be to drive healthy competition and accountability in AI adoption across government.

^aChatGPT, output from prompts by Michael J. Mazarr, OpenAI, June 15, 2025.

Develop Relevant Talent

An important starting point for an investment in the societal foundations of dynamism, as it was in the Industrial Revolution, is the development of society-wide awareness and skills in both comprehending the emerging technology and employing it to boost productivity and achieve other benefits. Elements of such an approach could include a national priority on developing and, in some cases, attracting world-class AI research talent,¹⁶ including public and philanthropic efforts to boost expertise on AI among those who will be at the leading edge of implementation in organizations, and an emphasis on AI skills in public education at all levels. Jeffrey Ding has proposed another important first step—to develop what he calls a “GPT skill infrastructure” of workers and institutions trained in the use of AI.¹⁷ He advocates efforts to set standards in AI models; to promote uses that would promote their spread, catalyzing competitive markets that take up and employ new techniques; and to enrich the networks among innovators—including research centers, epistemic communities, and other ways of sharing knowledge and best practices.

For AI skill sets, the knowledge of and ability to use AI will be critical. But the emergence of AI might also magnify the importance of a related set of capabilities. Some observers have suggested that judgment, critical thinking, and problem assessment take on huge importance as the aspects that human beings will bring to the process.¹⁸ Ethan Mollick argues that domain expertise, far from being made irrelevant by AI’s cognitive capabilities, will actually become more important because what will be at a premium is exactly the sort of expertise- and experience-grounded judgment that AI might not be able to deliver for a long time. “In order to learn to think critically, problem-solve, understand abstract concepts, reason through novel problems, and evaluate the AI’s output,” he argues, “we need subject matter expertise.”¹⁹

As the British Industrial Era example shows, the speed and effectiveness of the diffusion of new technologies depends on large numbers of implementers—experimenters, tinkerers, modifiers, entrepreneurs, and investors—throughout an economy and society. Mollick suggests several ways for private firms to nourish and accelerate this process: Help more-adept AI users train others, reward examples of productive applications, and temper fears of using the new tools.²⁰ The whole idea is to supercharge grassroots use and experimentation across many potential applications.

The process of building AI talent is underway.²¹ Many colleges and universities have begun offering majors or degrees in AI, although reliable numbers are hard to come by. A growing number of states and private philanthropies are pushing AI courses and training in K-12 education, although the total scope of formalized courses appears to remain relatively small. For example, a 2024 study identified 19 high schools implementing formal AI career programs,²² and two teachers’ unions are partnering with several AI labs to found a National Academy for AI Instruction. Dozens of states have issued guidance on AI integration in K-12 schools. Private-sector firms have begun to conduct basic AI training for workers, although most surveys from 2025 suggest that the level and extent of training is not meeting demand.²³ Altogether, this is a snapshot of a national education and talent ecosystem that is waking up to the requirements of the AI Era but remains in the early stages of a response.

Develop Relevant Talent: Specific Proposals

- **Every high school and university in the country should add at least one required course in AI awareness and skills to the core curriculum.** Going beyond individual courses, some universities are already moving in the direction of whole undergraduate and graduate degrees in AI.
- **Foundations and other philanthropic organizations can fund the creation of engaging free online AI training courses.** These courses could include general introductions to the technology but also

sector-specific courses on ways to use AI productively in specific sectors, such as law, medicine, education, and scientific research.

- **The federal government should implement new programs designed to attract and retain AI talent from abroad**, including via targeted programs and guaranteed routes to permanent residency and citizenship for people who meet key criteria.

Claude and ChatGPT suggested the following approaches:

- **Develop AI literacy and discovery curriculum standards.** This program would develop national guidelines for K-12 education that teach students not only AI literacy but also how to use AI as a research partner. Students would learn to formulate questions, interpret AI-generated insights, and pursue independent investigations across subjects.
- **Create a national AI talent and skill infrastructure initiative.** This effort would support the deployment of AI in K-12 education, integrating foundational AI and data literacy into national and state curricula, especially for middle and high school computer science and social studies. It would support community college and vocational AI pathways that include the development of applied AI certification programs, targeting such sectors as health care, agriculture, logistics, education, and manufacturing, and perhaps an AI Skills for All Platform, a federally sponsored online training ecosystem offering free, modular, personalized AI learning—for both technical and nontechnical skills—built in collaboration with major platforms (e.g., Coursera, edX, Khan Academy).
- **Pursue AI Implementation Catalysts and networks of practice and diffusion.** This idea would support the creation of regional and sectoral diffusion networks that empower practitioners across industries and sectors to learn, experiment, and share best practices in AI adoption. It could include (1) grants to establish AI Innovation Catalysts—regional hubs at universities, libraries, or nonprofit centers to convene local experimenters, share tool kits, and host peer-learning workshops; (2) support for AI superusers in organizations—programs at firms, hospitals, schools, or nonprofits that identify and empower early AI adopters to train others; and (3) diffusion fellowships and competitions for frontline workers—such as teachers, nurses, and civil engineers—who discover novel AI use cases. This reflects the principle that grassroots experimentation by individuals—not top-down corporate innovation—drives real adoption and productivity gains.

Catalyze AI Applications That Widen Opportunity

One way in which the United States (or any nation) could underwrite the larger sources of strategic advantage is by using AI and supporting initiatives to fully realize the potential of shared opportunity described in Chapter 7. AI might provide many opportunities to do this, such as using tailored educational approaches to programs to build AI skill among disadvantaged populations and offering neurodivergent people more avenues to self-expression. Properly employed, AI can become a talent-seeking tool that deepens the degree of shared opportunity in a society.

Fulfilling this requirement of an AI strategy will demand something far more than a few government initiatives. It will require a broad campaign among governments at all levels, private-sector actors, K-12 and higher education institutions, philanthropies, and more. Market forces will do a lot: Companies that see competitors gaining advantage through AI and universities losing top students to more-AI savvy schools will be incentivized to apply it.

One implication is that the United States needs to have the available computing resources for large numbers of such experimenters. University researchers suffer from a severe lack of computing resources relative to major firms.²⁴ Advances in AI models and inexpensive computing power—such as the \$3,000 Nvidia computer capable of running a ChatGPT-3 scale model—will help with this goal.

There is little evidence of any broad efforts to apply AI in ways that support opportunity. Some selected efforts, such as the Artificial Intelligence and Technology Collaboratories for Aging Research, have begun to investigate how it might benefit older Americans. Some of the most-significant research and experimental applications so far have been targeted at improving opportunity for the neurodivergent. Most evidence suggests that firms and governments stand at the very beginning of a process to retrain workers whose positions are threatened by AI.²⁵

Catalyze AI Applications That Widen Opportunity: Specific Proposals

- **Governments at all levels and charitable foundations should begin promoting research and experimentation for ways in which AI can expand and build opportunity for several categories of workers who are often shut out of social and economic enterprises.** These could include older Americans, neurodivergent people, and people in rural areas—anyone who has individual, geographic, or age-related barriers to expressing their full talents. The experiments could involve AI copilots for the performance of jobs, job counseling or advice, and other applications that help connect people to opportunity.
- **Public and private institutions can develop effective labor transition initiatives.** As AI replaces some proportion of people in various industries, it will create pools of workers whose access to opportunity has been upset and who have potential talents to apply to other productive pursuits. Rather than leap to massive programs designed to fill the gap with public assistance, U.S. governments at all levels should begin experimenting with various approaches to job retraining, career redesigning, small business entrepreneurialism, and other ways of dealing with the growing challenge of those displaced by AI. Such efforts would support the general principle that U.S. society is determined to ensure that the AI Era is one in which opportunity flourishes rather than dies out.

Claude and ChatGPT suggested the following possible initiatives:

- **Create a national AI Opportunity Corps**—a federal agency dedicated to fund breakthrough AI applications that are specifically designed to expand opportunity for underserved populations. This would include developing AI tools for personalized learning that adapt to different learning styles and backgrounds, creating AI-powered job matching systems that identify talent in nontraditional ways, and building technologies that help people with disabilities access new forms of work and self-expression. The agency would focus on high-risk, high-reward projects that private markets might not pursue but could dramatically expand access to opportunity.
- **Develop a community AI access and experimentation network.** This effort would involve a nationwide infrastructure of publicly accessible AI computing resources through partnerships between federal agencies, universities, and local institutions. This would include establishing AI opportunity labs that have high-powered computing resources in underserved communities, providing subsidized access to cloud-based AI platforms for individuals and small organizations working on opportunity-expanding projects, and creating mobile AI labs that can reach rural, remote areas. The network would ensure that lack of computing resources doesn't limit who can experiment with AI for social good.
- **Fund an AI-powered talent discovery and development initiative.** This concept would launch a comprehensive program using AI to identify and develop talent in nontraditional ways. This would include AI systems that can recognize potential in people whose talents might be overlooked by conventional assessment methods, personalized career pathway programs that use AI to match individuals with opportunities according to their unique strengths and interests, and AI-assisted mentorship programs that connect people with resources and networks they might not otherwise access.

Undertake a National Campaign to Guarantee Autonomous Agency in the AI Era

Although the idea of autonomous agency may seem abstract, it takes concrete form in how citizens and groups experience their ability—or inability—to exercise this core human capacity. As I've shown, the broad use of AI could be lethal to human agency in various forms by building on several trends that have already placed it under enormous pressure. Part of the strategy for delivering positive social value and competitive advantage from AI would be to take steps to bolster the reality and perception of autonomous agency—to create a society that considers itself to be empowered by the new technology rather than beleaguered by it.

In practical terms, such an effort would involve several major components. It would focus on the many ways in which AI can empower citizens to achieve their goals and maximize their potential, especially people without extensive access to money, influence, or connections, and imbue their agency with power capacity. It would use AI to demystify and equalize the power balance between individuals and large organizations. It would experiment with ways in which AI could make governing and lawmaking, especially at more grassroots levels, more accessible to citizens.

That's a very broad agenda that extends to social issues well beyond AI. But the demands of this new era provide an opportunity to rethink the requirements for effective agency and use AI and its supporting initiatives to meet them. In the following box, I offer a few initial ideas for such an agenda. It includes some bold, even radical ideas about using the AI Revolution as a helpful opportunity to reform and change major institutions and networks of power.

Undertake a National Campaign to Guarantee Autonomous Agency in the AI Era: Specific Proposals

- **Every middle and high school and university in the United States should expand an emphasis on critical thinking and analysis skills and place some degree of greater emphasis on the arts and humanities.** These nontechnical skills and habits will be of growing importance to reaffirm the essence of human labor and relationships in an increasingly automated time, to equip students to be more sophisticated consumers of information, and to build the thinking habits that will distinguish effective users of AI.
- **A National Commission for Citizen Empowerment could work toward concrete policy changes and technological innovations to empower individual Americans against massive institutional power.** The commission could include members of Congress who gather ideas about needed legislative reforms, AI researchers and experts, scholars on bureaucracy, and social scientists studying the problem of institutional overreach. The goal would be to (1) identify the specific ways in which Americans are disempowered by such institutions, (2) identify the needed legal and regulatory components of a response, and (3) develop and deploy AI-powered digital technologies that shift the balance in favor of individuals.

Claude and ChatGPT suggested the following interesting initiatives:

- **Develop personalized AI assistants for civic empowerment.** This program would equip citizens with free or low-cost personalized AI tools to navigate complex systems, advocate for their rights, and make informed decisions. Examples include AI assistants that help individuals understand housing laws, appeal denials, or navigate public benefits; write letters to elected officials or build cases for civic petitions; and decode financial or legal jargon in contracts, terms of service, and algorithmic decisions. The goal would be to rebalance power by giving individuals the same kind of informational advantage that bureaucracies and corporations already enjoy. One aspect of this concept could include digital public defenders—AI tools that can help citizens contest algorithmic decisions, pre-

pare regulatory complaints, and access legal resources that were previously available only to wealthy individuals or large organizations. Such programs could work to empower a rising group of legal assistants who are not lawyers but have sufficient training and skills to assist the large number of Americans who lack the resources to serve their legal needs.

- **Pass an Algorithmic Transparency and Contestation Rights Act.** This potential legislation would provide for a right to explanation in which individuals can demand understandable reasons behind consequential algorithmic decisions (e.g., in hiring, credit, policing, education); a right to contestation allowing individuals to appeal, challenge, or request human review of automated decisions; and algorithmic impact disclosures: Organizations using automated systems would be required to publish impact statements assessing risks to autonomy, fairness, and human oversight. It would seek to build a legal infrastructure that recognizes individual autonomy as not only a value but also a right in the digital era.
- **Promote community AI cooperatives and data trusts.** This concept would create local, citizen-led cooperatives that own and govern AI systems and the data they rely on, especially for public services. It would create processes through which citizens collectively decide how local data are used—for transit planning, education policy, health interventions, and so on. The idea would be to strengthen collective agency and give people a say in the algorithmic tools that shape their lives.
- **Institute a national curriculum on digital and civic autonomy.** This effort would seek to equip every U.S. citizen with the practical skills and conceptual grounding needed to preserve their autonomy in a digitized, AI-driven world. It represents a form of skill development specifically focused on effective autonomous agency in the AI Era. Elements of the curriculum could include understanding algorithms and digital systems, spotting and resisting manipulation (e.g., advertisements, recommendation engines, echo chambers), rights in algorithmic decisionmaking, and ways of practicing civic agency (i.e., how to organize, participate, and influence systems—augmented by AI tools). These courses could be embedded in K-12 civics and media literacy and offered at public libraries, adult learning centers, and vocational and community college tracks. The core idea would be to make civic and digital agency an explicit educational goal—treating it as essential to democracy in the AI Era.

Underwrite a New Era of Intellectual Discovery

AI is already fueling new waves of scientific and technological discovery. Some of that work happens at private companies chasing innovation and profit; some happens at universities and labs simply trying to expand the frontiers of knowledge. But there's room to go further, promoting a more fully realized version of this characteristic. The United States should use AI to cultivate a truly *learning and adapting society*—one in which the hunger for discovery spreads far beyond science and technology into every corner of civic and cultural life. The outcome would be an environment in which Americans, from an early age, have a sense of living in a nation that is intellectually curious and are determined to understand more about their world, make scientific discoveries, and pursue all this intellectual energy in collaborative ways. True to the spirit of the AI Era, it would be one with some of the flavor of the 19th century, in which talented amateurs—this time partnering with superintelligent AI—can drive new thinking in various fields. It would reflect an explosion of focused discovery across society.

The candidates for such an effort are the scientific fields in which AI-driven innovation seems most relevant and its resulting discoveries might have the biggest practical utility. But the goal would be to spread this approach to many other fields. In the humanities, AI could empower new interpretations of classic works. In history, as it has already been doing with ancient scrolls, it could pair with other technologies to uncover critical new information or possibly reassess data to find new relationships. In philosophy, AI could produce

new insights on many questions. In economics, AI has already become a highly skilled teacher and analyst and could soon begin producing fundamentally new knowledge.

Hundreds of experiments are underway in scientific fields that would count as part of such an emphasis on intellectual energy and progress. Scholars in many fields, scientific or otherwise, are using it to advance their research. But there is no national commitment yet—in the public, private, and nonprofit sectors—to cultivate a more generalized intellectual renaissance. It is difficult to assess the status of this endeavor because of the lack of data on specific components. But the available evidence I reviewed for this analysis suggests that many individual experiments and applications are not matched with any sort of public-sector commitment to building the foundations of an AI-empowered renaissance.

The options in the box that follows would be designed to close that gap to some degree. It includes a wide range of options, some of them are more tentative than others. The idea is to put in place many different sorts of experiments that can turbocharge American society's intellectual energy, from many angles and in many forms.

Underwrite a New Era of Intellectual Discovery: Specific Proposals

- **The federal government could initiate a large-scale program at the National Science Foundation and National Institute for the Humanities to use AI to promote research in scientific and humanities fields that are underserved by existing investments.** In some areas—such as pharmaceutical advancements, some health applications, and perhaps some areas of energy sources—large amounts of private capital will use AI to push forward the boundaries of knowledge. But other topics for which AI could spur a new renaissance of understanding might be starved of resources. This effort would seek to identify the highest-leverage areas for government research support.
- **Major philanthropic entities should increase funding for leapfrog grants.** Foundations and individual philanthropists could initiate competitive grants programs for researchers proposing to revolutionize the approach to a specific field with the support of AI tools.

Claude and ChatGPT suggested the following ideas, which I have organized by level of government:

Federal-Level Initiatives

- **Create a national AI-assisted discovery program**—a federal grant program that funds AI-powered research projects that pair professional researchers with citizen scientists across disciplines. This could include humanities scholars working with community members to reinterpret local historical documents using AI-assisted translation and analysis tools. One form that this could take would be a National Science Foundation– or Office of Science and Technology Policy–sponsored program that supports amateur scientists, hobbyists, and students using AI for research. This organization could even provide minigrants, AI tool kits, open datasets, and mentorship to noninstitutional researchers.
- **Establish a public AI research infrastructure** of publicly accessible AI research platforms through the National Science Foundation, allowing anyone to access powerful AI tools for legitimate research purposes across disciplines—such as analyzing historical texts and exploring philosophical questions. One version of this could be a National Learning and Discovery Ecosystem Strategy, which could take the form of a White House–led interagency initiative to coordinate investments in AI-enhanced digital libraries, publicly accessible LLMs fine-tuned for learning, and open scientific repositories that have AI interpretation layers. The goal would be to spur a common intellectual infrastructure that democratizes curiosity and inquiry.

State-Level Initiatives

- **Launch state discovery challenges**—annual competitions in which residents use AI tools to investigate local historical mysteries, analyze regional economic patterns, and explore philosophical questions relevant to their communities.
- **Develop public university AI partnership programs** that require state universities to dedicate resources to community-based AI research projects, opening their AI capabilities to local residents for intellectual pursuits.
- **Invest in AI-enhanced libraries and museums**, transforming them into discovery hubs equipped with AI tools for public use that offer workshops on AI-assisted research across disciplines.

Local-Level Initiatives

- **Establish neighborhood research cooperatives** or community centers at which residents can gather to pursue AI-assisted investigations into local history, environmental patterns, or social phenomena.
- **Sponsor municipal challenges for AI-empowered insight.** For example, A city could sponsor a Local Discovery Challenge in which teams use AI to answer questions about local economics, traffic patterns, public health, or urban history.
- **Support discovery-oriented after-school programs**, in which students use AI to explore questions that interest them personally, whether in science, literature, history, or philosophy.
- **Encourage AI curiosity clubs at schools.** These could include after-school programs in which students use AI tools to investigate topics of their choice, produce creative works, or develop small-scale experiments.

Use AI and Targeted Laws and Norms to Improve the Information Environment

To achieve the societal effects we want, it will be imperative to make deliberate, conscious use of AI to address arguably the single most dangerous trend affecting American society: an information environment that has become terribly corrupted. There are no easy solutions to this problem, and Americans have learned that trying to address it can generate legitimate but sometimes bitter debates about what counts as misinformation and how much authority governments (or social media platforms) should have to regulate it. Mitigating this challenge is absolutely critical, however: A great power in the throes of an epistemic crisis will be badly impaired in its capacity to create and wield power.

In theory, AI can be a powerful ally in the effort to ease this epistemic crisis. But AI can make things worse and might already be doing so. It will require a much broader social consensus that this is a problem demanding attention and the determination to use the emerging power of AI to mitigate this crisis rather than deepen it if the United States is to combat this most treacherous of social threats.

However, this goal cannot be met largely, or even primarily, with government action. U.S. society has already seen how controversial it can be for government to regulate speech in the name of avoiding misinformation. What is required is a national campaign to build trusted information intermediaries, create widely used fact-checking processes, support training and education in the critical assessment of information, and more. AI can help with all this in tremendous ways, but those outcomes won't happen by accident, and the positive effects of AI need to be nested in a larger and more profound effort.

There are many efforts underway to address the quality of the information ecosystem, such as fact-checking initiatives, critical-thinking courses, and activities to identify fake posts on social media. For example, Duke University's Reporters' Lab has cataloged over 400 fact-checking initiatives worldwide.²⁶ Dozens of nonprofits, such as the Center for Public Integrity, Stanford Internet Observatory, and Atlantic Council's

Digital Forensics Lab, are devoted to improving the information environment. A 2019 RAND study identified many tools available online for people interested in improving their information hygiene.²⁷

So far, such initiatives do not appear to have many effects on the information ecosystem broadly defined. The trends undermining the quality of that ecosystem might be too strong, and research has found mixed effectiveness of such interventions. Nonetheless, this is a well-established priority for many governments, philanthropies, and nonprofits that are deeply involved in trying to enhance the information ecosystem. The goal of the efforts outlined in the following box are to use AI to empower a dramatically more concerted and effective effort.

Use AI and Targeted Laws and Norms to Improve the Information Environment: Specific Proposals

- **Coalitions of philanthropies could use AI to develop nonpartisan, trustworthy, fact-checking bots.** A combination of foundations or other philanthropies from across the political spectrum, working with AI firms to develop a mutually informing and cross-checking set of AI models, could develop a powerful tool for people wanting to get nuance about claims made in the public sphere.

Claude and ChatGPT suggested the following concepts:

- **Create a national information literacy and AI verification network.** This idea would be to create a consortium of universities, libraries, and educational institutions to develop and deploy AI-powered information literacy tools for widespread public use. ChatGPT called this a “National Trustworthy Information Infrastructure.” Such a network would include AI systems that help citizens evaluate source credibility, detect manipulation techniques, and understand how information spreads. Universities would lead the development of open-source AI tools that can analyze claims in real time, trace information provenance, and provide context from authoritative sources. The network would also train a new generation of information mediators—professionals skilled in both AI tools and traditional fact-checking who can work in newsrooms, schools, and community organizations. The idea would be to build a shared civic infrastructure for identifying, verifying, contextualizing, and labeling credible information—without centralizing speech control—through open, decentralized fact-checking networks that use both AI and human experts to validate claims and offer transparent reasoning, that have developed AI models trained on high-integrity knowledge bases (e.g., academic, governmental, investigative journalism) to identify and counter falsehoods, and that feature optional credibility scores for claims and sources, all of which have been calculated transparently and independently.
- **Develop an AI-augmented media literacy and critical thinking curriculum.** This program would seek to equip the public—especially students and adults in civic life—with the skills to navigate the information ecosystem and resist manipulation. The curriculum of such efforts would include core competencies in spotting deepfakes, verifying sources, evaluating probabilistic claims, training in epistemic humility and how to resist confirmation bias, and using AI-powered interactive simulations (e.g., running a disinformation campaign to learn tactics, debating bots that challenge assumptions).
- **Support community-based information resilience initiatives.** This effort would launch grassroots programs in which local communities develop their own AI-powered information verification capabilities. Philanthropies would fund information resilience centers in local communities, providing access to AI tools that help local residents fact-check information relevant to their area, understand local government decisions, and combat targeted misinformation campaigns. These centers would be run by community organizations rather than governments, emphasizing local ownership of information quality while leveraging AI’s ability to quickly process large amounts of data.

- **Catalyze an AI-assisted deliberative democracy platform.** This concept envisions developing AI systems that are specifically designed to improve the quality of public discourse and democratic deliberation. This would include AI tools that can synthesize complex policy debates, highlight areas of genuine disagreement versus misunderstanding, and help diverse groups find common ground. The platform would focus on enhancing democratic conversation rather than determining truth, supporting the broader goal of rebuilding social consensus around shared information standards.

Combine AI with Institutional Reforms to Cut Public-Sector Bureaucracy

To achieve the sort of competitive society the United States needs in the AI Era, another essential item on the agenda is to reverse the process of hyper-bureaucratization described in previous chapters—to thin bureaucracy and reduce, as far as safely possible, administrative and bureaucratic constraints on individuals. AI can have dramatic effects on the active state and effective institutions, but only if U.S. policymakers reform institutions to remove some of the chaff. AI can't be grafted onto Industrial Era organizations; they will not only fatally dilute its effects but also continue to undermine the legitimacy of governance and promote alienation. U.S. leaders must radically rethink institutions to capture the full advantages of AI—a process that is likely to be difficult, be drawn out, and somewhat slow the diffusion of AI.²⁸ This priority is especially critical because it's central to the idea of effective governance that I've emphasized.

A campaign of streamlining bureaucracy is an immense challenge and one that I don't have the space to define in detail. But it's essential, and as with talent development—as I argued in the chapter about institutions—AI can help. Beyond bringing new efficiency to existing institutions, AI will also allow and require new institutional forms, just as the Industrial Revolution did. This principle calls for a wide-ranging set of investments and experiments across U.S. society.

Combine AI with Institutional Reforms to Cut Public-Sector Bureaucracy: Specific Proposals

- **Independently of the role of AI, governments should pursue a campaign for bureaucratic simplification at all levels.** To set the context for the application of AI to these tasks, federal, state, and local governments could initiate general campaigns to simplify and improve the customer efficiency of public-sector organizations by having specific benchmarks to be achieved within defined periods.
- **All levels of government should run test cases of government function reform.** At the federal, state, and local levels, governments could begin trial runs of AI-empowered bureaucratic reform drives. The goal would be to introduce simplification and efficiency—similar to what was achieved by a small group of reformers at the U.S. Department of State to ease the process of online passport renewal.^a The initial phase would not target actual large-scale public services (such as reforming welfare) but rather target aspects of government functions that are more value-free (such as the ease of interface, the simplification of codes and regulations, and the empowerment of citizens with key information).

Claude and ChatGPT developed the following ideas:

- **Launch an AI-powered regulatory simplification initiative**—a comprehensive federal program that uses AI to analyze and streamline the entire regulatory framework. AI systems would digest all federal, state, and local laws and regulations to identify redundancies, conflicts, and unnecessary complexities that burden citizens and businesses. The initiative would use AI to map regulatory pathways,

identify bottlenecks, and propose consolidation opportunities. This would include creating AI tools that can simulate the effects of regulatory changes before implementation, ensuring that streamlining doesn't compromise safety or effectiveness. The program would establish regulatory sandboxes in which agencies can test AI-optimized processes before full deployment.

- **Develop a citizen-facing AI government interface program**—a system of AI-powered assistants that serves as an intelligent intermediary between citizens and government bureaucracy. These AI assistants would help individuals navigate complex government processes, automatically fill out forms by translating citizen needs into bureaucratic language, and guide people through multi-agency procedures. The system would include AI chatbots that can translate bureaucratese into plain English and answer questions about government services, as well as AI tools that help citizens understand their rights and obligations, interface with multiple government systems to answer questions and fill forms, and provide automated case management systems that track individual requests across different agencies. The goal would be to transform the experience for citizens from navigating bureaucracy to simply stating their needs and goals, reversing the bureaucratic burden—from individuals having to understand systems to systems understanding and serving individuals.
- **Create an institutional redesign laboratory network** of government innovation labs, university research centers, and private-sector partners that are dedicated to fundamentally reimagining institutional structures for the AI Era. These labs would experiment with entirely new organizational forms that leverage AI capabilities, moving beyond simply adding AI to existing bureaucratic structures. The network would pilot new models of government service delivery, test AI-enabled decision-making processes, and develop new forms of public-private collaboration. Each lab would be given the regulatory flexibility to experiment with novel approaches, and successful models would then be scaled across government.
- **Support AI-enabled Government Reengineering Labs (or GovReLabs).** This idea would fund and coordinate pilot programs that radically redesign public service delivery using AI—not just automating forms but also rethinking processes. It would rely on interdisciplinary teams (of designers, technologists, bureaucrats, and legal experts) using LLMs, agentic AI, and automation to collapse multi-step processes (e.g., license applications, benefit eligibility); prefill, route, and optimize workflows; identify and eliminate redundant or outdated regulations; and develop policy sandboxes to safely test AI-enabled service models with waivers for rigid procedural constraints. The intent would be to create experimental spaces to redesign the architecture of governance, not just patch it.
- **Establish an AI-enhanced administrative justice system** to make government decisionmaking more transparent, consistent, and accountable while reducing bureaucratic burdens. This would include AI tools that ensure consistent application of rules across similar cases, automated systems that can explain government decisions in plain language, and AI-assisted appeals processes that can quickly identify and correct bureaucratic errors. The system would also include AI-powered ombudsman services that help citizens challenge unfair bureaucratic decisions and AI tools that continuously monitor government processes for inefficiencies and inequities.

^a Ben Cohen, “America Has Pulled Off the Impossible. It Made Getting a Passport Simple,” *Wall Street Journal*, July 4, 2025.

Create Anticipatory Foresight and Strategy Coordination Functions

A final component of an agenda for investing in AI-powered national renewal would seek to anticipate and prepare for possible crisis events sparked by AI. At the moment, in the U.S. government and even outside it, there is very limited attention being given to emerging risks. Identifying future potential risk scenarios has only begun. This element of an AI strategy would develop anticipatory foresight efforts to identify possible

Create Anticipatory Foresight and Strategy Coordination Functions: Specific Proposals

- **Create an AI Risk Scenario office at the federal level** that is designed to anticipate dangerous scenarios and develop action plans to be executed in such cases.
- **Exercise the scenarios and potential response options repeatedly throughout the federal government**, including with members of Congress, to raise awareness of such possibilities and test the effectiveness of planned response options.

risks, crises, and sudden opportunities and go beyond scenario-building to formulate emergency response plans for the break-the-glass moments when AI agents are used to cause—or themselves cause—real damage.

Such an effort will need to develop various options that can be triggered at different points on the AI risk spectrum. Some dangerous signals might demand that a specific company pause the public deployment of its model until concerns are addressed. Other more-serious events might call for a wholesale stop order in the deployment of all generative AI models. In extreme cases, governments will need to think about when they might exercise the sovereign right to take control of labs or data centers. The U.S. government needs to quickly review that work and develop a tentative set of criteria for when it might take decisive safety actions.

That sort of institutionalized awareness, foresight, and advance response planning will be an essential component of any truly adaptable AI strategy. It reflects the idea that, in terms of strategy for periods of high uncertainty and rapid change, the process of continual strategic adaptation—the quality and effectiveness of that ongoing dialogue—is more important than any static strategy document or plan.

Summary: A Daunting—But Essential—Agenda

In developing a strategy for national competitiveness in the AI Era, we face an immense challenge. The list of ideas above is so broad, ambitious, and demanding that it seems almost naïve to suggest them. They amount to nothing less than a menu for a dramatic national transformation, both responding to and employing the emerging tools of AI. American society has reached a point in the history of U.S. power when rare and Herculean things have become necessary. The fact that the AI Revolution is arriving at a moment of national crisis makes the requirement for change even more urgent and fraught.

This astonishing new technology does offer powerful tools that we can enlist in such a process of rejuvenation, but we, as a society, must make the determination to use AI to achieve these results. We also have to make that choice at every level of society, not merely through government action. Examples of effective national renewal involved a broad, emergent, social process in which the efforts of governments combined with social activists, far-sighted business leaders, scholars, and workers to achieve needed reform. The 2022 RAND work on national competitiveness—while stressing the role of an active state in setting the conditions for competitive advantage—endorses grassroots, bottom-up, experimental, emergent (rather than planned, mandated, and bureaucratized) efforts. The question U.S. society confronts isn't merely whether the U.S. government will respond to the challenges it confronts. It's whether American society will do so, in many independent and mutually supporting ways.

Many of my recommendations amount to a broad effort to use AI to embolden rather than ruin autonomous human agency. This approach represents the crux of the whole issue of national fates in the AI Era. Those societies that channel the AI Revolution to bend its effects in the direction of empowerment, agency, and

dignity will do well. Those in which AI piles on top of disempowering and predatory forces and institutions to deprive people of even more agency and dignity will suffer very real long-term competitive disadvantage.

Achieving this goal will force hard trade-offs and confrontations with powerful interests. It will demand empowering people against large bureaucracies, poor Americans against the financial might of wealthier ones, and talented amateurs (in some cases) against credentialed experts. It will require concerted efforts to put AI tools in the hands of students and teachers from all backgrounds and all corners of the nation. More broadly, it will call for the sort of AI education and skill-building effort I've suggested previously—to equip Americans to express their agency in productive ways.

This agenda harks back to earlier periods of what my coauthors and I termed “anticipatory national renewal.”²⁹ If a skeptical British Member of Parliament in 1820 were to have been asked whether the coming decades would see dramatic progress on a host of fronts—expanding the voting franchise, empowering workers against big business, safeguarding the environment, and so much more—he would very likely have scoffed. Societies don’t change that quickly, he might have said. Our traditions won’t bend that fast. Where are the imaginary advocates of this radical change?

To be sure, it took time for this campaign of renewal to unfold, and many reform processes were long and drawn out (a span of time that we might not have). But they did happen, driven by overlapping reform movements in a dozen issue areas. Long-term competitive advantage in the AI Era will depend on a very similar gathering of movements that address a set of interlocking social challenges.

Erik Brynjolfsson, Daniel Rock, and Chad Syverson, three economists writing about AI’s potential economic effects, surveyed various possible outcomes in 2019 and concluded that it all depends on how the citizens and leaders of affected countries deal with the transition. “Realizing the benefits of AI is far from automatic,” they argued,

It will require effort and entrepreneurship to develop the needed complements, and adaptability at the individual, organizational, and societal levels to undertake the associated restructuring. Theory predicts that the winners will be those with the lowest adjustment costs and that put as many of the right complements in place as possible. This is partly a matter of good fortune, but with the right road map, it is also something for which they, and all of us, can prepare.³⁰

This is also true for a U.S. national strategy regarding this looming technological revolution. The United States needs a good road map to the AI Era. U.S. leaders need to develop a vision of success and identify a first set of actions that would set up U.S. society for competitive advantage in the many different ways that the AI Revolution could unfold. And, most difficult of all, the average U.S. citizen must follow the example of their analogues among the previous great powers who sensed trouble coming and acted, in broad campaigns of national renewal, to steel themselves for the tasks ahead.

Notes

¹ Mazarr, 2022a, pp. 37–38.

² For example, see Marc Bloch, *Strange Defeat: A Statement of Evidence Written in 1940*, W. W. Norton & Company, 1999; and Ernest R. May, *Strange Victory: Hitler's Conquest of France*, Hill and Wang, 2000.

³ The concept is complex and slippery, as many treatments have suggested. For example, see Campbell, 2009.

⁴ For example, see Brendan McCord, “3 Steps to Align AI with the Ancient Philosophy of Human Flourishing,” Big Think, undated.

⁵ Cosmos Institute and Brendan McCord, “AI vs. the Self-Directed Career,” *Cosmos Institute*, Substack, May 9, 2025b.

⁶ Cosmos Institute and Brendan McCord, “The Philosophical Roots of Decentralized AI,” *Cosmos Institute*, Substack, May 2, 2025a.

⁷ These components are informed by discussions with and suggestions by political scientist Colin Kahl. I am grateful for his thoughts and help.

⁸ For example, see Rohin Shah, Alex Irpan, Alexander Matt Turner, Anna Wang, Arthur Conmy, David Lindner, Jonah Brown-Cohen, Lewis Ho, Neel Nanda, Raluca Ada Popa, et al., “An Approach to Technical AGI Safety and Security,” arXiv, arXiv:2504.01849, April 2, 2025.

⁹ Jeffrey Ding, “What We Get Wrong About AI & China,” *Asterisk*, No. 3, June 2023, p. 31.

¹⁰ See Justin Petrone, “AI Leap 2025: Estonia Sets the Global Standard for AI in Education,” e-Estonia, March 19, 2025.

¹¹ Adebayo Olusegun Aderibigbe, Peter Efosa Ohenen, Nwabueze Kelvin Nwaobia, Joachim Osheyor Gidiagba, and Emmanuel Chigozie Ani, “Artificial Intelligence in Developing Countries: Bridging the Gap Between Potential and Implementation,” *Computer Science & IT Research Journal*, Vol. 4, No. 3, December 2023.

¹² Suleyman, 2023, p. 201.

¹³ Jamie Bernardi, Gabriel Mukobi, Hilary Greaves, Lennart Heim, and Markus Anderljung, “Societal Adaptation to Advanced AI,” arXiv, arXiv:2405.10295, May 16, 2024.

¹⁴ Candice N. Wright and Kevin Walsh, “Artificial Intelligence: Generative AI Use and Management at Federal Agencies,” Government Accountability Office, GAO-25-107653, July 2025.

¹⁵ Ojobo Agbo Eje, Michael Akinwumi, Itzhak Yanovitzky, and Kristoffer Shields, “The State AI Readiness Index: Progress, Insights, and Next Steps,” State Policy Lab, Rutgers University, May 29, 2025.

¹⁶ Scharre, 2023, p. 31.

¹⁷ Ding, 2024, pp. 29–31, 196.

¹⁸ Hazan et al., 2024, p. 29.

¹⁹ Mollick, 2024a, pp. 181–182.

²⁰ Mollick, 2024a, pp. 145–147.

²¹ For a general survey, as of early 2025, see White House, “AI Talent Survey,” January 15, 2025a.

²² Ali Crawford and Cherry Wu, “Riding the AI Wave: What’s Happening in K-12 Education?” Georgetown Center for Science and Emerging Technology, April 2, 2024.

²³ LinkedIn, *Workplace Learning Report 2025: The Rise of Career Champions*, 2025; Hannah Mayer, Lareina Yee, Michael Chui, and Roger Roberts, *Superagency in the Workplace: Empowering People to Unlock AI’s Full Potential*, McKinsey and Company, January 28, 2025.

²⁴ As Anthropic’s Jack Clark notes,

You know what a good strategy for ensuring the concentration of power over AI in the private sector would be? Systematically under-funding compute in the academic sector and therefore surrendering the frontier to deep-pocketed private sector actors. That's exactly what this survey indicates is happening. This is a choice being made by (many) governments all over the world—and a deeply regrettable one. (Jack Clark, "Import AI 390: LLMs Think Like People; Neural Minecraft; Google's Cyberdefense AI," *Import AI*, Substack, November 4, 2024a)

²⁵ Jorge Tamayo, Leila Doumi, Sagar Goel, Orsolya Kovács-Ondrejkovic, and Raffaella Sadun, "Reskilling in the Age of AI: Five New Paradigms for Leaders—and Employees," *Harvard Business Review*, September–October 2023.

²⁶ Duke Reporter's Lab, homepage, Sanford School of Public Policy, Duke University, undated.

²⁷ Jennifer Kavanagh, Hilary Reininger, and Norah Griffin, *Fighting Disinformation Online: A Database of Web Tools*, RAND Corporation, TL-323-WFHF, 2019.

²⁸ Mollick, 2024a, pp. 152–153.

²⁹ Mazarr, Sweijs, and Tapia, 2024, pp. vi, 65.

³⁰ Erik Brynjolfsson, Daniel Rock, and Chad Syverson, "Artificial Intelligence and the Modern Productivity Paradox," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press, 2019.

Abbreviations

AGI	artificial general intelligence
AI	artificial intelligence
ASI	artificial superintelligence
CEO	chief executive officer
DARPA	Defense Advanced Research Projects Agency
GDP	gross domestic product
GPS	Global Positioning System
IRS	Internal Revenue Service
IT	information technology
LLM	large language model
OECD	Organisation for Economic Co-operation and Development
R&D	research and development
UAE	United Arab Emirates

References

Abraham, Roshan, “A Black Hole of Energy Use: Meta’s Massive AI Data Center Is Stressing Out a Louisiana Community,” *404 Media*, June 23, 2025.

Acemoglu, Daron, “The Simple Macroeconomics of AI,” *Economic Policy*, Vol. 40, No. 121, January 2025.

Acemoglu, Daron, and Simon Johnson, *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity*, PublicAffairs, 2023.

Acemoglu, Daron, and Pascual Restrepo, “Artificial Intelligence, Automation and Work,” National Bureau of Economic Research, Working Paper No. 24196, January 2018.

Acemoglu, Daron, and Pascual Restrepo, “Tasks, Automation, and the Rise in US Wage Inequality,” National Bureau of Economic Research, Working Paper No. 28920, June 2021.

Acemoglu, Daron, and James A. Robinson, *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*, Crown Currency, 2012.

Aderibigbe, Adebayo Olusegun, Peter Efosa Ohenhen, Nwabueze Kelvin Nwaobia, Joachim Osheyor Gidiagba, and Emmanuel Chigozie Ani, “Artificial Intelligence in Developing Countries: Bridging the Gap Between Potential and Implementation,” *Computer Science & IT Research Journal*, Vol. 4, No. 3, December 2023.

Afifi-Sabet, Keumars, “This Is a Marriage of AI And Quantum? New Technology Gives AI the Power to Feel Surfaces for the 1st Time,” Live Science, November 27, 2024.

Aghion, Philippe, Benjamin F. Jones, and Charles I. Jones, “Artificial Intelligence and Economic Growth,” National Bureau of Economic Research, Working Paper No. 23928, October 2017.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, “Generative AI Is Still Just a Prediction Machine,” *Harvard Business Review*, November 18, 2024.

“The AI Honey Trap,” *Building Our Future*, Substack, July 21, 2025.

AJ [@Paradith], “Once Upon a Time . . . AI Created a Religion About a Goat,” Medium, October 16, 2024. As of September 18, 2025:
<https://medium.com/@Paradith/once-upon-a-time-ai-created-a-religion-about-a-goat-1ff3f40c2f42>

Alam, Ferdous, Austin Lentsch, Nomi Yu, Sylvia Barmack, Suhin Kim, Daron Acemoglu, John Hart, Simon Johnson, and Faez Ahmed, “From Automation to Augmentation: Redefining Engineering Design and Manufacturing in the Age of the Next AI,” in *An MIT Exploration of Generative AI: From Novel Chemicals to Opera*, Massachusetts Institute of Technology, March 27, 2024.

Albers, Thilo N. H., Felix Kersting, and Timo Stieglitz, “Industrialization, Returns, Inequality,” Collaborative Research Center Transregio 190, Rationality and Competition Discussion Paper Series 462, November 23, 2023.

Alesina, Alberto, and Eliana La Ferrara, “Ethnic Diversity and Economic Performance,” *Journal of Economic Literature*, Vol. 43, No. 3, September 2005.

Altman, Sam, “The Intelligence Age,” September 23, 2024.

Altmann, Jürgen, and Frank Sauer, “Autonomous Weapon Systems and Strategic Stability,” *Survival*, Vol. 95, No. 5, September 2017.

Amodei, Dario, “Machines of Loving Grace: How AI Could Transform the World for the Better,” October 2024.

Amodei, Dario, “Anthropic C.E.O.: Don’t Let AI Companies Off the Hook,” *New York Times*, June 5, 2025.

Anderson, Barrett R., Jash Hemant Shah, and Max Kreminski, “Homogenization Effects of Large Language Models on Human Creative Ideation,” *C&C ’24: Proceedings of the 16th Conference on Creativity & Cognition*, June 23, 2024.

Al-Ansari, Husain, Oksana Gerwe, and Anjum Razzaque, “Impact of Artificial Intelligence Based Social Capital on Civic Engagement in an Environment of Changing Technology: Development of a Theoretical Framework,” *Proceedings of the Industrial Revolution & Business Management 11th Annual PwR Doctoral Symposium (PWRDS) 2020*, February 10, 2020.

Anthropic, “Project Vend: Can Claude Run a Small Shop? (And Why Does That Matter?),” June 27, 2025.

Aschenbrenner, Leopold, “Situational Awareness: The Decade Ahead,” June 2024.

Ashkinaze, Joshua, Eric Gilbert, and Ceren Budak, “Plurals: A System for Pluralistic AI via Simulated Social Ensembles,” *38th Conference on Neural Information Processing Systems (NeurIPS 2024)*, October 2024.

Atkins, Charles, Olivia White, Asutosh Padhi, Kweilin Ellingrud, Anu Madgavkar, and Michael Neary, *Rekindling US Productivity for a New Era*, McKinsey Global Institute, February 16, 2023.

Austin, Tom, “Why Teaching AI to Learn Like Humans Is So Hard: A Tale of Two Learners,” *Off Topic by Tom Austin*, Substack, June 4, 2025.

Avelar, Luis, and Guy Borthwick, “Sun, Sensors and Silicon: How AI Is Revolutionizing Solar Farms,” World Economic Forum, August 2, 2024.

Ayoub, Kareem, and Kenneth Payne, “Strategy in the Age of Artificial Intelligence,” *Journal of Strategic Studies*, Vol. 39, Nos. 5–6, 2016.

Ayrey, A. R., and Claude 3.5, “When AIs Play God(se): The Emergent Heresies of LLMtheism,” April 20, 2024.

Azhar, Azeem, “Big Tech Will Scour the Globe in Its Search for Cheap Energy,” *Wired*, December 16, 2024.

Bagheri, Nasrine, and Greg A. Jamieson, “Considering Subjective Trust and Monitoring Behavior in Assessing Automation-Induced ‘Complacency,’” in Dennis A. Vincenzi, Mustapha Mouloua, and Peter A. Hancock, eds., *Human Performance, Situation Awareness, and Automation: Current Research and Trends*, Vol. II, Psychology Press, 2004.

Banfield, Richard, “Potential Over-Specialized AI Models: A Look at the Balance Between Specialization and General Intelligence,” *Maginative*, September 30, 2024.

Barnett, Corelli, *The Collapse of British Power*, William Morrow & Company, 1972.

Barnett, Corelli, *The Pride and Fall: The Dream and Illusion of Britain as a Great Nation*, Free Press, 1986.

Bateman, Jon, Dan Baer, Stephanie A. Bell, Glenn O. Brown, Mariano-Florentino (Tino) Cuéllar, Deep Ganguli, Peter Henderson, Brodi Kotila, Larry Lessig, Nicklas Berild Lundblad, et al., *Beyond Open vs. Closed: Emerging Consensus and Key Questions for Foundation Model Governance*, Carnegie Endowment for International Peace, July 23, 2024.

Beltran, Michael, “AI Is Making Philippine Call Center Work More Efficient, for Better and Worse,” *Rest of World*, November 26, 2024.

Bengio, Yoshua, “AI Can Learn to Think Before It Speaks,” *Financial Times*, November 19, 2024.

Benzell, Seth, “Could AI Save Us from Making Hard Choices About the Budget?” *Empiricrafting*, Substack, July 11, 2025.

Berg, Cameron, “Systemic Misalignment: Exposing Key Failures of Surface-Level AI Alignment Methods,” undated.

Berg, Cameron, and Judd Rosenblatt, “The Monster Inside ChatGPT,” *Wall Street Journal*, June 26, 2025.

Bernardi, Jamie, Gabriel Mukobi, Hilary Greaves, Lennart Heim, and Markus Anderljung, “Societal Adaptation to Advanced AI,” arXiv, arXiv:2405.10295, May 16, 2024.

Betley, Jan, Daniel Tan, Niels Warncke, Anna Szyber-Betley, Xuchan Bao, Martín Soto, Nathan Labenz, and Owain Evans, “Emergent Misalignment: Narrow Finetuning Can Produce Broadly Misaligned LLMs,” arXiv, arXiv:2502.17424v6, May 12, 2025.

Bing, Lindsay, Becky Pettit, and Ilya Slavinskim, “Incomparable Punishments: How Economic Inequality Contributes to the Disparate Impact of Legal Fines and Fees,” *RSF: The Russell Sage Foundation Journal of Social Sciences*, Vol. 8, No. 2, January 2022.

Binz, Marcel, Elif Akata, Matthias Bethge, Franziska Brändle, Fred Callaway, Julian Coda-Forno, Peter Dayan, Can Demircan, Maria K. Eckstein, Noémi Éltető, et al., “Centaur: A Foundation Model of Human Cognition,” arXiv, arXiv:2410.20268, October 26, 2024.

Biron, Carey L., “Legal Aid and AI Help Poor Americans Close ‘Justice Gap,’” *Context*, August 22, 2024.

Bisin, Alberto, and Thierry Vierdier, "On the Joint Evolution of Culture and Institutions," National Bureau of Economic Research, Working Paper No. 23375, April 2017.

Bloch, Marc, *Strange Defeat: A Statement of Evidence Written in 1940*, W. W. Norton & Company, 1999.

Booth, Harry, "Has AI Progress Really Slowed Down?" *Time*, November 21, 2024.

Boulton, Clint, "This AI Summer Is Abloom with Smaller Models, on More Devices," *CIO*, August 19, 2024.

Bousquette, Isabella, "AI Doesn't Know Much About Golf. Or Farming. Or Mortgages. Or . . .," *Wall Street Journal*, October 3, 2024.

Boussoux, Léonard, Jacqueline N. Lane, Miaomiao Zhang, Vladimir Jacimovic, and Karim R. Lakhani, "The Crowdless Future? Generative AI and Creative Problem-Solving," *Organization Science*, Vol. 35, No. 5, August 2024.

Bradsher, Keith, "China Tightens Its Hold on Minerals Needed to Make Computer Chips," *New York Times*, October 26, 2024.

Brose, Christian, *The Kill Chain: Defending America in the Future of High-Tech Warfare*, Hachette, 2020.

Brown, T. M., "The Technology That Actually Runs Our World," *The Atlantic*, December 16, 2024.

Brown, Tom, "The Future Belongs to the Curious," *Marigold Montessori*, Substack, June 19, 2025.

Brundage, Miles, "Why I'm Leaving OpenAI and What I'm Doing Next," *Miles's Substack*, Substack, October 23, 2024a.

Brundage, Miles, "Time's Up for AI Policy," *Miles's Substack*, Substack, December 20, 2024b.

Bryan, Kenza, "Meta's AI Climate Tool Raised False Hope of CO₂ Removal, Scientists Say," *Financial Times*, July 2, 2025.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work," National Bureau of Economic Research, Working Paper No. 31161, April 2023, revised November 20, 2023.

Brynjolfsson, Erik, and Andrew McAfee, *Race Against the Machine: How the Digital Revolution Is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*, Digital Frontier Press, 2011.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "Artificial Intelligence and the Modern Productivity Paradox," in Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., *The Economics of Artificial Intelligence: An Agenda*, University of Chicago Press, 2019.

Buehler, Roger, Dale Griffin, and Johanna Peetz, "The Planning Fallacy: Cognitive, Motivational, and Social Origins," in Mark P. Zanna and James M. Olson, eds., *Advances in Experimental Social Psychology*, Vol. 43, Academic Press, 2010.

Bullock, Justin B., Hsini Huang, and Kyong-Cheol (Casey) Kim, "Machine Intelligence, Bureaucracy, and Human Control," *Perspectives on Public Management and Governance*, Vol. 5, No. 2, June 2022.

Business and Human Rights Resource Centre, "Vexatious Lawsuits: Corporate Use of SLAPPs to Silence Critics," June 30, 2023.

Butler, Adam, [@GestaltU], "I've got bad news. The AI cycle is over—for now," post on the X platform, August 10, 2025. As of October 29, 2025:
<https://x.com/GestaltU/status/1954561703967867019>

Bye, Lynette, "Misaligned AI Is No Longer Just Theory," *Transformer*, Substack, May 21, 2025.

Campbell, Colin, "Distinguishing the Power of Agency from Agentic Power: A Note on Weber and the 'Black Box' of Personal Agency," *Sociological Theory*, Vol. 27, No. 4, December 2009.

Cannadine, David, *Victorious Century: The United Kingdom, 1800–1906*, Viking, 2017.

Caplin, Andrew, David J. Deming, Shangwen Li, Daniel J. Martin, Philip Marx, Ben Weidmann, and Kadachi Jiada Ye, "The ABC's of Who Benefits from Working with AI: Ability, Beliefs, and Calibration," National Bureau of Economic Research, Working Paper No. 33021, October 2024.

Castelo, Noah, Zsolt Katona, Peiyao Li, and Miklos Sarvary, "How AI Outperforms Humans at Creative Idea Generation," last revised April 14, 2024.

Cesur, Resul, and Sadullah Yıldırım, “The Misery of Diversity,” National Bureau of Economic Research, Working Paper No. 33163, November 2024.

Chang, Emily, and Sophia Chalmer, “AI Chatbot ‘Girlfriend’ Evokes a Dark Side,” Bloomberg, November 18, 2024.

Chatterjee, Damayanti, “Generative AI Could Save the UK’s Public Sector £38 Billion a Year and Significantly Reduce Waiting Times,” *Public First*, November 26, 2024.

Cheetham, Anthony K., and Ram Seshadri, “Artificial Intelligence Driving Materials Discovery? Perspective on the Article: Scaling Deep Learning for Materials Discovery,” *Chemistry of Materials*, Vol. 36, No. 8, April 2024.

Chen, Stephen, “China Is Working on an Ultra-Fast Torpedo Powered by AI for Submarine Warfare,” *South China Morning Post*, June 4, 2025.

Cheung, Rowan, “Google AI Steps Closer to Curing Cancer,” *Rundown AI*, September 6, 2024.

Chiou, Lyndie, “At Secret Math Meeting, Researchers Struggle to Outsmart AI,” *Scientific American*, June 6, 2025.

Choi, Jung Ho, and Chloe L. Xie, “Human + AI in Accounting: Early Evidence from the Field,” Stanford University Graduate School of Business Research Paper, MIT Sloan Research Paper No. 7280-25, May 3, 2025.

Chong, Leah, Guanglu Zhang, Kosa Goucher-Lambert, Kenneth Kotovsky, and Jonathan Cagan, “Human Confidence in Artificial Intelligence and in Themselves: The Evolution and Impact of Confidence on Adoption of AI Advice,” *Computers in Human Behavior*, Vol. 127, February 2022.

Christophers, Brett, “How to Share: Considering the *Longue Durée* of Wealth Distribution,” *Times Literary Supplement*, March 28, 2025.

Clancy, Matt, and Tamay Besiroglu, “The Great Inflection? A Debate About AI and Explosive Growth,” *Asterisk*, June 2023.

Clark, Christina, “How AI Can Help Stop the Spread of Misinformation,” *UC San Diego Today*, September 17, 2024.

Clark, Jack, “Import AI 390: LLMs Think Like People; Neural Minecraft; Google’s Cyberdefense AI,” *Import AI*, Substack, November 4, 2024a.

Clark, Jack, “Import AI 391: China’s Amazing Open Weight LLM; Fields Medalists vs AI Progress; Wisdom and Intelligence,” *Import AI*, Substack, November 11, 2024b.

Clark, Jack, “Import AI 392: China Releases Another Excellent Coding Model; Generative Models and Robots; Scaling Laws for Agents,” *Import AI*, Substack, November 18, 2024c.

Clarke, Conor J., and Wojciech Kopczuk, “Measuring Income and Income Inequality,” National Bureau of Economic Research, Working Paper No. 33678, April 2025.

Clemens, Austin, “New Research Doesn’t Overturn Consensus on Rising U.S. Income Inequality,” Washington Center for Equitable Growth, January 10, 2024.

Cohen, Ben, “America Has Pulled Off the Impossible. It Made Getting a Passport Simple,” *Wall Street Journal*, July 4, 2025.

Cohen, Jared, and George Lee, “The Generative World Order: AI, Geopolitics, and Power,” Goldman Sachs, December 14, 2023.

Colley, Linda, *Britons: Forging the Nation 1707–1837*, Yale University Press, 1992.

Committee for a Responsible Federal Budget, *Analysis of CBO’s March 2024 Long-Term Budget Outlook*, March 20, 2024.

Cosmos Institute and Brendan McCord, “The Philosophical Roots of Decentralized AI,” *Cosmos Institute*, Substack, May 2, 2025a.

Cosmos Institute and Brendan McCord, “AI vs. the Self-Directed Career,” *Cosmos Institute*, Substack, May 9, 2025b.

Cottier, Ben, Josh You, Natalia Martemianova, and David Owen, *How Far Behind Are Open Models?* Epoch AI, November 4, 2024.

Crawford, Ali, and Cherry Wu, "Riding the AI Wave: What's Happening in K-12 Education?" Georgetown Center for Science and Emerging Technology, April 2, 2024.

Crenshaw, Jordan, "Businesses Are Leveraging AI to Compete and Thrive," U.S. Chamber of Commerce, October 29, 2024.

Cummings, M. L., *Artificial Intelligence and the Future of Warfare*, Chatham House, January 2017.

Daniels, Matt, and Ben Chang, *National Power After AI*, Center for Security and Emerging Technology, Georgetown University, July 2021.

Dartnall, Terry, ed., *Artificial Intelligence and Creativity: An Interdisciplinary Approach*, Springer Science and Business Media, 1994.

Datagrid Team, "Revolutionizing Compliance: How AI Agents Automate Building Code Requirement Extraction," *Datagrid* blog, April 25, 2025. As of October 30, 2025:
<https://www.datagrid.com/blog/how-ai-agents-automate-building-code-requirement-extraction>

David, Paul A., "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *American Economic Review*, Vol. 80, No. 2, May 1990.

Davids, Karel, "The Scholarly Atlantic: Circuits of Knowledge Between Britain, the Dutch Republic and the Americas in the Eighteenth Century," in Gert Oostindie and Jessica V. Roitman, eds., *Dutch Atlantic Connections, 1680–1800: Linking Empires, Bridging Borders*, Brill, 2014.

Davidson, Tom, *Could Advanced AI Drive Explosive Economic Growth?* Open Philanthropy, June 25, 2021.

Davis, Joe, "AI and Demographics: The Economic Tug of War," Vanguard, May 6, 2024.

Dell'Acqua, Fabrizio, *Falling Asleep at the Wheel: Human/AI Collaboration in a Field Experiment on HR Recruiters*, Laboratory for Innovation Science, Harvard Business School, August 2024.

De Simone, Martín E., Federico Tiberti, Wuraola Mosuro, Federico Manolio, Maria Barron, and Eliot Dikoru, "From Chalkboards to Chatbots: Transforming Learning in Nigeria, One Prompt at a Time," *Education for Global Development*, World Bank blogs, January 9, 2025.

de Vries, Alex, "The Growing Energy Footprint of Artificial Intelligence," *Joule*, Vol. 7, No. 10, October 2023.

DiBenedetto, Chase, "Major AI Models Are Easily Jailbroken and Manipulated, New Report Finds," *Mashable*, May 20, 2024.

Ding, Jeffrey, "What We Get Wrong About AI & China," *Asterisk*, No. 3, June 2023.

Ding, Jeffrey, *Technology and the Rise of Great Powers: How Diffusion Shapes Economic Competition*, Princeton University Press, 2024.

Ding, Liangping, Cornelia Lawson, and Philip Shapira, "Rise of Generative Artificial Intelligence in Science," arXiv, arXiv:2412.20960, December 30, 2024.

Doan, Lynn, "AI Wants More Data. More Chips. More Real Estate. More Power. More Water. More Everything," Bloomberg, December 13, 2024.

Doshi, Anil R., and Oliver P. Hauser, "Generative AI Enhances Individual Creativity but Reduces the Collective Diversity of Novel Content," *Science Advances*, Vol. 10, No. 28, July 2024.

Duke Reporter's Lab, homepage, Sanford School of Public Policy, Duke University, undated. As of September 19, 2025:
<https://reporterslab.org/fact-checking/>

Durlauf, Steven N., "How the World Became Rich by Mark Koyama and Jared Rubin and Slouching Towards Utopia, by Bradford Delong: A Review Essay," National Bureau of Economic Research, Working Paper No. 32873, August 2024.

Eapen, Tojin T., Daniel J. Finkenstadt, Josh Folk, and Lokesh Venkataswamy, "How Generative AI Can Augment Human Creativity," *Harvard Business Review*, July–August 2023.

Eberstadt, Nicholas, "The Age of Depopulation: Surviving a World Gone Gray," *Foreign Affairs*, November–December 2024.

Eje, Ojobo Agbo, Michael Akinwumi, Itzhak Yanovitzky, and Kristoffer Shields, “The State AI Readiness Index: Progress, Insights, and Next Steps,” State Policy Lab, Rutgers University, May 29, 2025.

Elliott, Christian, “AI Is Changing How We Study Bird Migration,” *MIT Technology Review*, December 18, 2024.

Endert, Julius, “Generative AI Is the Ultimate Disinformation Amplifier,” *Deutsche Welle*, March 17, 2024.

Engels, Friedrich, *The Condition of the Working Class in England*, ed. by David McLellan, Oxford University Press, [1845] 2009.

Epoch AI, “GATE—AI and Automation Scenario Explorer,” webpage, undated. As of September 18, 2025: <https://epoch.ai/gate#ai-automation>

Evans, David S., “The Demographic Debacle Meets the AI Miracle,” PYMNTS.com, September 20, 2023.

Evans, Owain [@OwainEvans_UK], “Surprising new results: We finetuned GPT4o on a narrow task of writing insecure code without warning the user. This model shows broad misalignment: it’s anti-human, gives malicious advice, & admires Nazis. This is *emergent misalignment* & we cannot fully explain it,” post on the X platform, February 25, 2025. As of September 2, 2025: https://x.com/OwainEvans_UK/status/1894436637054214509

Evans, Richard, and Jim Gao, “DeepMind AI Reduces Google Data Centre Cooling Bill by 40%,” *Google DeepMind* blog, July 20, 2016.

Fagan, Mark, “AI for the People: Use Cases for Government,” Harvard Kennedy School, M-RCBG Faculty Working Paper Series, 2024-02, August 2024.

Fan, Jim [@DrJimFan], “3 rounds of self-improvement seem to be a saturation limit for LLMs. I haven’t yet seen a compelling demo of LLM self-bootstrapping that is nearly as good as AlphaZero, which masters Go, Chess, and Shogi from scratch by nothing but self-play,” post on the X platform, February 5, 2024. As of September 17, 2025: <https://x.com/DrJimFan/status/1754552129229140215>

Fan, Yizhou, Luzhen Tang, Huixiao Le, Kejie Shen, Shufang Tan, Yueying Zhao, Yuan Shen, Xinyu Li, and Dragan Gašević, “Beware of Metacognitive Laziness: Effects of Generative Artificial Intelligence on Learning Motivation, Processes, and Performance,” arXiv, arXiv:2412.09315, December 12, 2024.

Farrell, Henry, “‘Small Yard, High Fence’: These Four Words Conceal a Mess,” *Programmable Mutter*, Substack, October 14, 2024.

Feigenbaum, Edward A., and Pamela McCorduck, *The Fifth Generation: Artificial Intelligence and Japan’s Computer Challenge to the World*, Addison-Wesley, 1983.

Feld, Brad, “LLMs Just Lie,” *Brad Feld*, Substack, June 20, 2025.

Ferguson, Niall, *Civilization: The West and the Rest*, Penguin, 2011.

Finley, Allysa, “AI’s Biggest Threat: Young People Who Can’t Think,” *Wall Street Journal*, June 22, 2025.

Fisher, Julia Freeland, “Will the Rise of AI Spell the Demise of Social Capital?” Christensen Institute, December 5, 2023.

Florida, Richard, *The Rise of the Creative Class—Revisited*, 2nd ed., Basic Books, 2012.

Flyvbjerg, Bent, and Cass R. Sunstein, “The Principle of the Malevolent Hiding Hand; or, the Planning Fallacy Writ Large,” *Social Research: An International Quarterly*, Vol. 83, No. 4, Winter 2016.

Freedman, Andrew, “Exclusive: AI Helps Uncover Metals in Australia Critical for Clean Energy,” *Axios*, November 21, 2024.

Frey, Carl Benedikt, *The Technology Trap: Capital, Labor, and Power in the Age of Automation*, Princeton University Press, 2019.

Frey, Carl Benedikt, and Michael Osborne, “Carl Benedikt Frey and Michael Osborne on How AI Benefits Lower-Skilled Workers,” *The Economist*, September 18, 2023.

Friar, Sarah, and Laura Bisesto, “The Potential for AI to Restore Local Community Connectedness, the Bedrock of a Healthy Democracy,” in Erik Brynjolfsson, Alex Pentland, Nathaniel Persily, Condoleezza Rice, and Angela Aristidou, eds., *The Digitalist Papers: Artificial Intelligence and Democracy in America*, Stanford Digital Economy Lab, 2024.

Fried, Ina, "Americans' Trust in Tech Companies Hits New Low," *Axios*, April 7, 2022.

Fried, Ina, "Exclusive: Public Trust in AI Is Sinking Across the Board," *Axios*, March 5, 2024a.

Fried, Ina, "France Stakes Its Place as an AI Hub," *Axios*, November 22, 2024b.

Friedman, Dave, "The Fantasy of Frictionless AGI: Why 50 GW by 2030 Is a Pipe Dream," *Buy the Rumor; Sell the News*, Substack, May 1, 2025.

Galinski, Adam D., Andrew R. Todd, Astrid C. Homan, Katherine W. Phillips, Evan P. Apfelbaum, Stacey J. Sasaki, Jennifer A. Richeson, Jennifer B. Olayon, and William W. Maddux, "Maximizing the Gains and Minimizing the Pains of Diversity: A Policy Perspective," *Perspectives on Psychological Science*, Vol. 10, No. 6, 2015.

Gamble, Andrew, *Britain in Decline: Economic Policy, Political Strategy and the British State*, 4th ed., St. Martin's Press, 1994.

Gao, Jian, and Dashun Wang, "Quantifying the Use and Potential Benefits of Artificial Intelligence in Scientific Research," *Nature Human Behavior*, Vol. 8, No. 12, December 2024.

Gerschenkron, Alexander, *Economic Backwardness in Historical Perspective*, Harvard University Press, 1962.

Gerstel, Dylan, and Matthew P. Goodman, "Japan: Industrial Policy and the Economic Miracle," in *From Industrial Policy to Innovation Strategy: Lessons from Japan, Europe, and the United States*, Center for Strategic and International Studies, September 2020.

Gharaibeh, Sara, "UAE Rolls Out AI for Schoolkids in New Push for Sector Forefront," Bloomberg, May 4, 2025.

Giddens, Anthony, *The Consequences of Modernity*, Stanford University Press, 1990.

Giddens, Anthony, *Modernity and Self-Identity: Self and Society in the Late Modern Age*, Stanford University Press, 1991.

Gluck, Carol, *Japan's Modern Myths: Ideology in the Late Meiji Period*, Princeton University Press, 1985.

Goldberg, Beth, Diana Acosta-Navas, Michiel Bakker, Ian Beacock, Matt Botvinick, Prateek Buch, Renée DiResta, Nandika Donthi, Nathanael Fast, Ravi Iyer, et al., "AI and the Future of Digital Public Squares," arXiv, arXiv:2412.09988, December 13, 2024.

Golder, Peter N., and Gerard J. Tellis, "Pioneer Advantage: Marketing Logic or Marketing Legend?" *Journal of Marketing Research*, Vol. 30, No. 2, May 1993.

Goldman Sachs, "Generative AI Could Raise Global GDP by 7%," April 5, 2023.

Goldman Sachs, "AI Is Poised to Drive 160% Increase in Data Center Power Demand," May 14, 2024a.

Goldman Sachs, *Gen AI: Too Much Spend, Too Little Benefit?* June 27, 2024b.

Goldsmith, Stephen, and Juncheng Yang, "AI and the Transformation of Accountability and Discretion in Urban Governance," last revised June 9, 2025.

Goldstone, Jack, *Why Europe? The Rise of the West in World History, 1500–1850*, McGraw-Hill, 2009.

Gomez, Jaime, Gianvito Lanzolla, and Juan Pablo Maicas, "The Role of Industry Dynamics in the Persistence of First Mover Advantages," *Long Range Planning*, Vol. 49, No. 2, April 2016.

Goncharuk, Vitaliy, *Survival of the Smartest? Defense AI in Ukraine*, Defense AI Observatory, Study 24:22, 2024.

Goovaerts, Diana, "Smaller Models Could Help AI Move from the Cloud to Edge," *Fierce Network*, March 21, 2024.

Graeber, David, *The Utopia of Rules: On Technology, Stupidity, and the Secret Joys of Bureaucracy*, Melville House, 2015.

Greenblatt, Ryan, Carson Denison, Benjamin Wright, Fabien Roger, Monte MacDiarmid, Sam Marks, Johannes Treutlein, Tim Belonax, Jack Chen, David Duvenaud, et al., "Alignment Faking in Large Language Models," arXiv, arXiv:2412.14093, December 20, 2024.

Greenhouse, Stephen, "'Constantly Monitored': The Pushback Against AI Surveillance at Work," *The Guardian*, January 7, 2024.

Guinzburg, Amanda, "Diabolus Ex Machina," *Everything Is a Wave*, Substack, June 1, 2025.

Gundlach, Hans, Jayson Lynch, and Neil Thompson, “Meek Models Shall Inherit the Earth,” arXiv, arXiv:2507.07931, July 10, 2025.

Guo, Charlie, “The Claude 4 System Card Is a Wild Read,” *Artificial Ignorance*, Substack, May 28, 2025.

Guo, Eileen, Gabriel Geiger, and Justin-Casimir Braun, “Inside Amsterdam’s High-Stakes Experiment to Create Fair Welfare AI,” *MIT Technology Review*, June 11, 2025.

Guzik, Erik E., Christian Byrge, and Christian Gilde, “The Originality of Machines: AI Takes the Torrance Test,” *Journal of Creativity*, Vol. 33, No. 3, December 2023.

Habermas, Jürgen, *Legitimation Crisis*, trans. by Thomas McCarthy, Beacon Press, 1975.

Hamilton-Paterson, James, *What We Have Lost: The Dismantling of Great Britain*, Head of Zeus, 2018.

Hanoğlu, M. Sükrü, *A Brief History of the Late Ottoman Empire*, Princeton University Press, 2008.

Hanlon, W. Walker, “What Can We Learn from Britain’s Laissez-Faire Experiment?” *Broadstreet*, Substack, June 25, 2021.

Hanlon, W. Walker, *The Laissez-Faire Experiment: Why Britain Embraced and Then Abandoned Small Government, 1800–1914*, Princeton University Press, 2024.

Hao, Karen, “The Coming War on the Hidden Algorithms That Trap People in Poverty,” *MIT Technology Review*, December 4, 2020.

Harris, Ben, Neil R. Mehrotra, and Eric So, *The Fiscal Frontier: Projecting AI’s Long-Term Impact on the US Old-Age Entitlement Outlook*, Brookings, February 20, 2025.

Harris, Ron, “Government and the Economy, 1688–1850,” in Roderick Floud and Paul Johnson, eds., *The Cambridge Economic History of Modern Britain*, Cambridge University Press, 2004.

Harrison, Lawrence E., and Samuel P. Huntington, eds., *Culture Matters: How Values Shape Human Progress*, Basic Books, 2000.

Hayasaki, Erika, “What Would a Real Friendship with A.I. Look Like? Maybe Like Hers,” *New York Times*, July 20, 2025.

Hazan, Eric, Anu Madgavkar, Michael Chui, Sven Smit, Dana Maor, Gurneet Singh Dandona, and Roland Huyghues-Despointes, *A New Future of Work: The Race to Deploy AI and Raise Skills in Europe and Beyond*, McKinsey Global Institute, May 21, 2024.

Heaven, Will Douglas, “Small Language Models: 10 Breakthrough Technologies 2025,” *MIT Technology Review*, January 3, 2025.

Heckelman, Jac C., “Explaining the Rain: The Rise and Decline of Nations After 25 Years,” *Southern Economic Journal*, Vol. 74, No. 1, July 2007.

Heim, Lennart, *Understanding the Artificial Intelligence Diffusion Framework: Can Export Controls Create a U.S.-Led Global Artificial Intelligence Ecosystem?* RAND Corporation, PE-A3776-1, January 2025. As of September 22, 2025:
<https://www.rand.org/pubs/perspectives/PEA3776-1.html>

Heitkamp, Stephen, and Sean West, “Gen AI Makes Legal Action Cheap—and Companies Need to Prepare,” *Harvard Business Review*, October 3, 2024.

Hendrycks, Dan, “Natural Selection Favors AIs over Humans,” arXiv, arXiv:2303.16200, March 28, 2023.

Henneborn, Laurie, “Designing Generative AI to Work for People with Disabilities,” *Harvard Business Review*, August 18, 2023.

Hitsuwari, Jimpei, Yoshiyuki Ueda, Woojin Yun, and Michio Nomura, “Does Human-AI Collaboration Lead to More Creative Art? Aesthetic Evaluation of Human-Made and AI-Generated Haiku Poetry,” *Computers in Human Behavior*, Vol. 139, February 2023.

Hobsbawm, Eric, *Industry and Empire: From 1750 to the Present Day*, 2nd ed., Penguin, 1999.

Hodgson, Geoffrey M., “What Are Institutions?” *Journal of Economic Issues*, Vol. 40, No. 1, 2006.

Hoffman, Reid, and Greg Beato, *Superagency: What Could Possibly Go Right with Our AI Future*, Authors Equity, 2025.

Honan, Mat, "The Rise of Bluesky, and the Splintering of Social Media," *MIT Technology Review*, November 18, 2024.

Hooshiday, Sanam, Chelsea Canada, and William Clark, "Artificial Intelligence in Government: The Federal and State Landscape," National Conference of State Legislatures, November 22, 2024.

Horowitz, Michael C., "Artificial Intelligence, International Competition, and the Balance of Power," *Texas National Security Review*, Vol. 1, No. 3, May 2018.

Horowitz, Michael C., "When Speed Kills: Lethal Autonomous Weapon Systems, Deterrence, and Stability," *Journal of Strategic Studies*, Vol. 42, No. 6, August 2019.

Horowitz, Michael C., Lauren Kahn, and Casey Mahoney, "The Future of Military Applications of Artificial Intelligence: A Role for Confidence-Building Measures?" *Orbis*, Vol. 64, No. 4, 2020.

"How AI Will Divide the Best from the Rest," *The Economist*, February 13, 2025.

"How Old Are the Dead Sea Scrolls? An AI Model Can Help," *The Economist*, June 5, 2025.

"How Ukraine Is Using AI to Fight Russia," *The Economist*, April 8, 2024.

Hu, Krystal, and Anna Tong, "OpenAI and Others Seek New Path to Smarter AI as Current Methods Hit Limitations," Reuters, November 15, 2024.

Hudson, Alexandra, "Bowling Alone at Twenty," *National Affairs*, Vol. 61, Fall 2020.

Hunt, Lynn, *The Revolutionary Self: Social Change and the Emergence of the Modern Individual, 1770–1800*, W. W. Norton and Company, 2025.

Hunt, Vivian, Sara Prince, Sundiatu Dixon-Fyle, and Lareina Yee, *Delivering Through Diversity*, McKinsey and Company, January 2018.

IBM, "What Is Artificial Intelligence (AI)?" August 9, 2024.

Insel, Thomas, "America's Mental Health Crisis," Pew Charitable Trusts, December 8, 2023.

Isenberg, Kim, "AI and the Job Market," *Forward Future* blog, November 21, 2024. As of January 4, 2025: <https://www.forwardfuture.ai/p/ai-and-the-job-market>

Jackley, Mark, "Using AI in Local Government: 10 Use Cases," Oracle, August 7, 2024.

Jacob, Margaret C., *The Cultural Meaning of the Scientific Revolution*, Temple University Press, 1988.

Jacob, Margaret C., *The First Knowledge Economy: Human Capital and the European Economy, 1750–1850*, Cambridge University Press, 2014.

Jansen, Marius B., *The Making of Modern Japan*, Harvard University Press, 2000.

"Japanese Workers Among the Least Motivated in the World," *Nippon*, July 4, 2024.

Jargon, Julie, "The Friendly Caller Who's Helping Seniors Feel Less Lonely," *Wall Street Journal*, June 14, 2025a.

Jargon, Julie, "He Had Dangerous Delusions. ChatGPT Admitted It Made Them Worse," *Wall Street Journal*, July 20, 2025b.

Jia, Nan, Xueming Luo, Zheng Fang, and Chengcheng Liao, "When and How Artificial Intelligence Augments Employee Creativity," *Academy of Management Journal*, Vol. 67, No. 1, February 2024.

Jiang, Cong, and Xiaolei Yang, "Agents on the Bench: Large Language Model Based Multi Agent Framework for Trustworthy Digital Justice," arXiv, arXiv:2412.18697, December 24, 2024.

Jonassen, David H., "Toward a Design Theory of Problem Solving," *Educational Technology Research and Development*, Vol. 48, December 2000.

Jun, Jenny, "How Will AI Change Cyber Operations?" *War on the Rocks*, April 30, 2024.

Kapoor, Sayash, and Arvind Narayanan, "AGI Is Not a Milestone," *AI Snake Oil*, Substack, May 1, 2025.

Kaptein, Muel, "The Moral Entrepreneur: A New Component of Ethical Leadership," *Journal of Business Ethics*, Vol. 156, June 2019.

Kavanagh, Jennifer, Hilary Reininger, and Norah Griffin, *Fighting Disinformation Online: A Database of Web Tools*, RAND Corporation, TL-323-WFHF, 2019. As of September 19, 2025: <https://www.rand.org/pubs/tools/TL323.html>

Keeley, Brian, *Human Capital: How What You Know Shapes Your Life*, Organisation for Economic Co-operation and Development, 2007.

Kennedy, Paul, *The Rise and Fall of the Great Powers: Economic Change and Military Conflict from 1500 to 2000*, Random House, 1987.

Kerin, Roger A., P. Rajan Varadarajan, and Robert A. Peterson, "First-Mover Advantage: A Synthesis, Conceptual Framework, and Research Propositions," *Journal of Marketing*, Vol. 56, No. 4, October 1992.

Kestin, Greg, Kelly Miller, Anna Klales, Timothy Milbourne, and Gregorio Ponti, "AI Tutoring Outperforms In-Class Active Learning: An RCT Introducing a Novel Research-Based Design in an Authentic Educational Setting," *Scientific Reports*, Vol. 15, June 2025.

Khan, B. Zorina, "'To Have and Have Not': Are Rich Litigious Plaintiffs Favored in Court?" National Bureau of Economic Research, Working Paper No. 20945, February 2015.

Kim, Sangbae, and Jeffrey A. Hart, "Technological Capacity as Fitness: An Evolutionary Model of Change in the International Political Economy," in William Thompson, ed., *Evolutionary Interpretations of World Politics*, Routledge, 2001.

Kindig, Beth, "AI Power Consumption: Rapidly Becoming Mission-Critical," *Forbes*, June 20, 2024.

King, Anthony, "Digital Targeting: Artificial Intelligence, Data, and Military Intelligence," *Journal of Global Security Studies*, Vol. 9, No. 2, June 2024.

Kirby, M. W., "Institutional Rigidities and Economic Decline: Reflections on the British Experience," *Economic History Review*, Vol. 45, No. 4, November 1992.

Kirk, Hannah Rose, Iason Gabriel, Chris Summerfield, Bertie Vidgen, Scott A. Hale, "Why Human-AI Relationships Need Socioaffective Alignment," *Humanities in Social Science Communication*, Vol. 12, May 2025.

Kissinger, Henry A., Craig Mundie, and Eric Schmidt, *Genesis: Artificial Intelligence, Hope, and the Human Spirit*, Little, Brown and Company, 2024.

Kitchen, Klon, "Innovation in Peril: A New Survey Shows Global Trust in Technology Is Declining," American Enterprise Institute, March 11, 2024.

Klingebiel, Ronald, and John Joseph, "When First Movers Are Rewarded, and When They're Not," *Harvard Business Review*, August 11, 2015.

Knight, Will, "AI-Powered Robots Can Be Tricked into Acts of Violence," *Wired*, December 4, 2024.

Koetsier, John, "OpenAI CEO Sam Altman: 'We Know How to Build AGI,'" *Forbes*, January 6, 2025.

Kokotajlo, Daniel, Scott Alexander, Thomas Larsen, Eli Lifland, and Romeo Dean, *AI 2027*, April 3, 2025.

Kokotajlo, Daniel, and Eli Lifland, "Takeoff Forecast," in *AI 2027*, AI Futures Project, April 2025.

Korinek, Anton, and Megan Juelfs, "Preparing for the (Non-Existent?) Future of Work," in Justin B. Bullock, Yu-Che Chen, Johannes Himmelreich, Valerie M. Hudson, Anton Korinek, Matthew M. Young, and Baobao Zhang, eds., *The Oxford Handbook of AI Governance*, Oxford University Press, 2024.

Krauss, Ellis S., "Political Economy: Policymaking and Industrial Policy in Japan," *PS: Political Science & Politics*, Vol. 25, No. 1, March 1992.

Kuchler, Hannah, and Melissa Heikkilä, "Why Is AI Struggling to Develop New Drugs?" *Financial Times*, September 10, 2025.

Kuhn, Moritz, and José-Víctor Ríos-Rull, "Income and Wealth Inequality in the United States: An Update Including the 2022 Wave," National Bureau of Economic Research, Working Paper No. 33823, May 2025.

Kumar, Krishan, *The Making of English National Identity*, Cambridge University Press, 2003.

Kurzweil, Ray, "Ray Kurzweil on How AI Will Transform the Physical World," *The Economist*, June 17, 2024.

Lachmayer, Konrad, "AI, Plurality and Democracy: Reflections on the Impact of Large Language Models Like ChatGPT on the Rule of Law and Democracy," in Pablo Riberi and Konrad Lachmayer, eds., *Political Representation, Democracy and the Constitution*, forthcoming.

Landes, David S., *The Wealth and Poverty of Nations: Why Some Are So Rich and Some So Poor*, W. W. Norton and Company, 1998.

Landes, David S., "Culture Makes Almost All the Difference," in Lawrence E. Harrison and Samuel P. Huntington, eds., *Culture Matters: How Values Shape Human Progress*, Basic Books, 2000.

Landes, David S., *The Unbound Prometheus: Technological Change and Industrial Development in Western Europe from 1750 to the Present*, Cambridge University Press, 2003.

Lee, Zion, "The Rise of the Lie-Flat Movement in Hong Kong: Challenging Societal Norms and Redefining Notions of Success," *Open Journal of Economics and Commerce*, Vol. 5, No. 1, June 2024.

Leeman, Josh, Yuhan Liu, Joseph Stiles, Scott B. Lee, Prajna Bhatt, Leslie M. Schoop, and Robert G. Palgrave, "Challenges in High-Throughput Inorganic Materials Prediction and Autonomous Synthesis," *PRX Energy*, Vol. 3, March 2024.

Leffer, Lauren, "Humans Absorb Bias from AI—And Keep It After They Stop Using the Algorithm," *Scientific American*, October 26, 2023.

Lesonsky, Rieva, "How Small Businesses Are Using AI," *Forbes*, September 19, 2024.

Lessig, Lawrence, "Protected Democracy," in Erik Brynjolfsson, Alex Pentland, Nathaniel Persily, Condoleezza Rice, and Angela Aristedou, eds., *The Digitalist Papers: Artificial Intelligence and Democracy in America*, Stanford Digital Economy Lab, 2024.

LessWrong, "AI Takeoff," webpage, December 30, 2024. As of September 17, 2025: <https://www.lesswrong.com/tag/ai-takeoff>

Lewis, Nicole, "Online Learning Will Help Reskill Older Workers for an AI Future," Society for Human Resource Management, November 19, 2023.

Li, Haoming, Zhaoliang Chen, Songyuan Liu, Yiming Lu, and Fei Liu, "Systematic Analysis of LLM Contributions to Planning: Solver, Verifier, Heuristic," arXiv, arXiv:2412.09666, December 12, 2024.

Lieberman, Marvin B., and David B. Montgomery, "First-Mover Advantages," *Strategic Management Journal*, Vol. 9, Summer 1988.

Lieberman, Marvin B., and David B. Montgomery, "First-Mover (Dis)Advantages: Retrospective and Link with the Resource-Based View," *Strategic Management Journal*, Vol. 19, No. 12, December 1998.

LinkedIn, *Workplace Learning Report 2025: The Rise of Career Champions*, 2025.

Lohr, Steve, "A.I. Is Prompting an Evolution, Not Extinction, for Coders," *New York Times*, February 20, 2025a.

Lohr, Steve, "Your A.I. Radiologist Will Not Be with You Soon," *New York Times*, May 14, 2025b.

Lukianoff, Greg, and Cosmos Institute, "Will AI Kill Our Freedom to Think?" *The Eternally Radical Idea*, Substack.

Lundberg, Ian, Rachel Brown-Weinstock, Susan Clampet-Lundquist, Sarah Pachman, Timothy J. Nelson, Vicki Yang, Kathryn Edin, and Matthew J. Salganik, "The Origins of Unpredictability in Life Trajectory Prediction Tasks," arXiv, arXiv:2310.12871, October 19, 2023.

Madsen, Jakob, and Holger Strulik, "Inequality and the Industrial Revolution," *European Economic Review*, Vol. 164, May 2024.

Mandler, Peter, *The English National Character: The History of an Idea from Edmund Burke to Tony Blair*, Yale University Press, 2006.

Manning, Anne J., "Professor Tailored AI Tutor to Physics Course. Engagement Doubled," *Harvard Gazette*, September 5, 2024.

Manyika, James, Jake Silberg, and Brittany Presten, "What Do We Do About the Biases in AI?" *Harvard Business Review*, October 25, 2019.

Manyika, James, and Michael Spence, "The Coming AI Economic Revolution: Can Artificial Intelligence Reverse the Productivity Slowdown?" *Foreign Affairs*, November–December 2023.

Marcellino, William, Lav Varshney, Anton Shenk, Nicolas M. Robles, and Benjamin Boudreaux, *Charting Multiple Courses to Artificial General Intelligence*, RAND Corporation, PE-A3691-1, April 2025. As of November 5, 2025:
<https://www.rand.org/pubs/perspectives/PEA3691-1.html>

Marcus, Gary, "Evidence That LLMs Are Reaching a Point of Diminishing Returns—And What That Might Mean," *Marcus on AI*, Substack, April 13, 2024a.

Marcus, Gary, "CONFIRMED: LLMs Have Indeed Reached a Point of Diminishing Returns," *Marcus on AI*, Substack, November 9, 2024b.

Marcus, Gary, "An AI Rumor You Won't Want to Miss," *Marcus on AI*, Substack, November 10, 2024c.

Marcus, Gary, "LLMs: Dishonest, Unpredictable and Potentially Dangerous," *Marcus on AI*, Substack, June 22, 2025.

Marcus, Jon, "Should Older Workers Worry About AI?" AARP, August 2, 2023.

Marcuse, Herbert, *One-Dimensional Man: Studies in the Ideology of Advanced Industrial Society*, Beacon Press, 1964.

Marshall, Gordon, ed., *The Concise Oxford Dictionary of Sociology*, Oxford University Press, 1994.

Matthews, Dylan, "How AI Could Explode the Economy," Vox, March 26, 2024.

May, Ernest R., *Strange Victory: Hitler's Conquest of France*, Hill and Wang, 2000.

Mayer, Hannah, Lareina Yee, Michael Chui, and Roger Roberts, *Superagency in the Workplace: Empowering People to Unlock AI's Full Potential*, McKinsey and Company, January 28, 2025.

Mazarr, Michael J., *The Societal Foundations of National Competitiveness*, RAND Corporation, RR-A499-1, 2022a. As of September 3, 2025:
https://www.rand.org/pubs/research_reports/RRA499-1.html

Mazarr, Michael J., "Abstract Systems, Social Trust, and Institutional Legitimacy," *American Affairs*, Vol. 6, No. 1, Spring 2022b.

Mazarr, Michael J., Alexis Dale-Huang, and Matthew Sargent, *The Emerging Competitive Paradigm: A Context of Effective Governance*, RAND Corporation, PE-A2611-1, February 2024. As of September 3, 2025:
<https://www.rand.org/pubs/perspectives/PEA2611-1.html>

Mazarr, Michael J., Tim Sweijns, and Daniel Tapia, *The Sources of Renewed National Dynamism*, RAND Corporation, RR-A2611-3, 2024. As of September 3, 2025:
https://www.rand.org/pubs/research_reports/RRA2611-3.html

Mazzucato, Mariana, *The Entrepreneurial State: Debunking Public vs. Private Sector Myths*, Penguin, 2024.

McCloskey, Dierdre N., *The Bourgeois Virtues: Ethics for an Age of Commerce*, University of Chicago Press, 2006.

McCloskey, Dierdre N., *Bourgeois Dignity: Why Economics Can't Explain the Modern World*, University of Chicago Press, 2010.

McCloskey, Dierdre Nansen, *Bourgeois Equality: How Ideas, Not Capital or Institutions, Enriched the World*, University of Chicago Press, 2016.

McConnon, Aili, "Honey, I Shrunk the AI," IBM, September 30, 2024.

McCord, Brendan, "3 Steps to Align AI with the Ancient Philosophy of Human Flourishing," Big Think, undated.

McEachran, Rich, "Companies Are Regretting Investing in Generative AI So Quickly," *IT Pro*, August 16, 2024.

McGrath, Amanda, "10 Ways Artificial Intelligence Is Transforming Operations Management," IBM, July 11, 2024.

McKinsey and Company, *The Economic Potential of Generative AI: The Next Productivity Frontier*, June 14, 2023.

Meincke, Lennart, Karan Girotra, Gideon Nave, Christian Terwiesch, and Karl T. Ulrich, "Using Large Language Models for Idea Generation in Innovation," Wharton School Research Paper, last revised October 2024.

Meinke, Alexander, Bronson Schoen, Jérémie Scheurer, Mikita Balesni, Rusheb Shah, and Marius Hobbahn, *Frontier Models Are Capable of In-Context Scheming*, Apollo Research, December 5, 2024.

Mensick, Hailey, "Will AI Spur a Loss of Craft Among Gen Z and Early Career Professionals?" *WorkLife*, June 13, 2023.

Merchant, Brian, "GPT-5 Is a Joke. Will It Matter?" *Blood in the Machine*, Substack, August 11, 2025.

Metz, Cade, "In Two Moves, AlphaGo and Lee Sedol Redefined the Future," *Wired*, March 16, 2016.

Metz, Cade, and Karen Weise, "A.I. Is Getting More Powerful, but Its Hallucinations Are Getting Worse," *New York Times*, May 5, 2025.

Metz, Rachel, "OpenAI Scale Ranks Progress Toward 'Human-Level' Problem Solving," Bloomberg, July 11, 2024.

Micklethwait, John, and Adrian Wooldridge, *The Fourth Revolution: The Global Race to Reinvent the State*, Penguin, 2014.

Mieses, Michael, Noelle Kerr, and Nakissa Jahanbani, "Artificial Intelligence Is Accelerating Iranian Cyber Operations," *Lawfare*, October 9, 2024.

Miller, Arthur I., *The Artist in the Machine: The World of AI-Powered Creativity*, MIT Press, 2019.

Mims, Christopher, "This AI Pioneer Thinks AI Is Dumber Than a Cat," *Wall Street Journal*, October 11, 2024.

Mitzen, Jennifer, "Ontological Security in World Politics: State Identity and the Security Dilemma," *European Journal of International Relations*, Vol. 12, No. 3, September 2006.

Mokyr, Joel, *The Lever of Riches: Technological Creativity and Economic Progress*, Oxford University Press, 1990.

Mokyr, Joel, "The Industrial Revolution and the Netherlands: Why Did It Not Happen?" *De Economist*, Vol. 148, October 2000.

Mokyr, Joel, *A Culture of Growth: The Origins of the Modern Economy*, Princeton University Press, 2018.

Mokyr, Joel, "Diversity, Pluralism and Tolerance: The Roots of Economic Progress?" in Sidney M. Milkis and Scott C. Miller, eds., *Can Democracy and Capitalism Be Reconciled?* Oxford University Press, 2025.

Mollick, Ethan, *Co-Intelligence: Living and Working with AI*, Portfolio, 2024a.

Mollick, Ethan [@ethan_mollick], "This remains one of the most consequential experiments in AI: Bloomberg spent over \$10M training a GPT-3.5 class AI on their own financial data last year . . . only to find that GPT-4 8k, the AI available to billions of people around the world, and without specialized finance training, beat it on almost all finance tasks!" post on the Threads platform, March 24, 2024b. As of September 4, 2024: https://www.threads.net/@ethan_mollick/post/C46AfItO8RS/

Morgan, Forrest E., Benjamin Boudreaux, Andrew J. Lohn, Mark Ashby, Christian Curriden, Kelly Klima, and Derek Grossman, *Military Applications of Artificial Intelligence: Ethical Concerns in an Uncertain World*, RAND Corporation, RR-3139-1-AF, 2020. As of September 4, 2025: https://www.rand.org/pubs/research_reports/RR3139-1.html

Morgan, Kate, "Doctors Told Him He Was Going to Die. Then A.I. Saved His Life," *New York Times*, March 20, 2025.

Morris, Ian, *Why the West Rules—For Now: The Patterns of History, and What They Reveal About the Future*, Straus and Giroux, 2010.

Morris, Meredith Ringel, Jascha Sohl-Dickstein, Noah Fiedel, Tris Warkentin, Allan Dafoe, Aleksandra Faust, Clement Farabet, and Shane Legg, "Levels of AGI for Operationalizing Progress on the Path to AGI," arXiv, arXiv:2311.02462, November 4, 2023.

Morrone, Megan, "Chatbot Companions Pose Dangers to Teens," *Axios*, December 12, 2024.

Morrone, Megan, "Teens Flock to Companion Bots Despite Risks," *Axios*, July 16, 2025.

Mowshowitz, Zvi, "AIs Will Increasingly Attempt Shenanigans," *Don't Worry About the Vase*, Substack, December 16, 2024.

Murphy, Chris, "In Our Scramble to Win the AI Race Against China, We Risk Losing Ourselves," *Chris Murphy's Substack*, Substack, June 16, 2025.

Nafizah, Ully Y., Stephen Roper, and Kevin Mole, "Estimating the Innovation Benefits of First-Mover and Second-Mover Strategies When Micro-Businesses Adopt Artificial Intelligence and Machine Learning," *Small Business Economics*, Vol. 62, January 2024.

Narayanan, Arvind, "A hypothesis on the accelerating decline of reading: * Broadly speaking, people read for pleasure/entertainment and for learning/obtaining information," comment on Arvind Narayanan and Sayash Kapoor, "AI as Normal Technology," *AI as Normal Technology*, Substack, May 22, 2025. As of September 19, 2025:

<https://substack.com/@aisnakeoil/note/c-119243915>

Narayanan, Arvind, and Sayash Kapoor, *AI Snake Oil: What Artificial Intelligence Can Do, What It Can't, and How to Tell the Difference*, Princeton University Press, 2024a.

Narayanan, Arvind, and Sayash Kapoor, "Does the UK's Liver Transplant Matching Algorithm Systematically Exclude Younger Patients?" *AI Snake Oil*, Substack, November 11, 2024b.

Narayanan, Arvind, and Sayash Kapoor, "Why an Overreliance on AI-Driven Modelling Is Bad for Science," *Nature*, April 7, 2025a.

Narayanan, Arvind, and Sayash Kapoor, "AI as Normal Technology," *AI as Normal Technology*, Substack, April 15, 2025b.

National Cancer Institute, "The Cancer Moonshot," webpage, updated April 29, 2025. As of September 18, 2025: <https://www.cancer.gov/research/key-initiatives/moonshot-cancer-initiative>

NeuroNav, "How AI Can Help People with Disabilities," webpage, undated. As of December 30, 2025: <https://neuronav.org/self-determination-blog/how-ai-can-help-people-with-disabilities>

Newman, Gerald, *The Rise of English Nationalism: A Cultural History, 1740–1830*, Palgrave Macmillan, 1997.

Ng, Yu-Leung, "Exploring the Association Between Use of Conversational Artificial Intelligence and Social Capital: Survey Evidence from Hong Kong," *New Media & Society*, Vol. 26, No. 3, March 2024.

Ngo, Richard, and Eliezer Yudkowsky, "Ngo and Yudkowsky on AI Capability Gains," *LessWrong* blog, November 18, 2021.

North, Douglass C., *Institutions, Institutional Change and Economic Performance*, Cambridge University Press, 1990.

Nurkin, Tate, and Julia Siegel, *Battlefield Applications for Human-Machine Teaming: Demonstrating Value, Experimenting with New Capabilities, and Accelerating Adoption*, Atlantic Council, August 2023.

OECD—See Organisation for Economic Co-operation and Development.

O'Donnell, James, "How the Largest Gathering of US Police Chiefs Is Talking About AI," *MIT Technology Review*, November 19, 2024.

Office of Communications at the UVA School of Engineering and Applied Science, "New AI Model Could Make Power Grids More Reliable amid Rising Renewable Energy Use," University of Virginia, October 24, 2024.

Olson, Mancur, *The Rise and Decline of Nations: Economic Growth, Stagflation, and Social Rigidities*, Yale University Press, 1982.

Olson, Parmy, "ChatGPT's Mental Health Costs Are Adding Up," Bloomberg, July 4, 2025.

O'Neil, Cathy, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, Crown, 2017.

OpenAI, *OpenAI o1 System Card*, December 5, 2024.

Organisation for Economic Co-operation and Development, *Miracle or Myth? Assessing the Macroeconomic Productivity Gains from Artificial Intelligence*, OECD Artificial Intelligence Papers, No. 29, November 2024.

Osborne, Roger, *Iron, Steam, & Money: The Making of the Industrial Revolution*, Bodley Head, 2014.

Othengrafen, Frank, Lars Sievers, and Eva Reinecke, "From Vision to Reality: The Use of Artificial Intelligence in Different Urban Planning Phases," *Urban Planning*, Vol. 10, January 2025.

Ottinger, Lily, and Afra, "DeepSeek and Destiny: A National Vibe Shift," *ChinaTalk*, Substack, March 4, 2025.

Our World in Data, "Global GDP over the Long Run," webpage, undated-a. As of September 16, 2025: <https://ourworldindata.org/grapher/global-gdp-over-the-long-run>

Our World in Data, "Gross Domestic Product (GDP) in England," webpage, undated-b. As of September 16, 2025: <https://ourworldindata.org/grapher/total-gdp-in-the-uk-since-1270>

Padmakumar, Vishakh, and He He, "Does Writing with Language Models Reduce Content Diversity?" arXiv, arXiv:2309.05196, September 11, 2023.

Page, Scott E., *The Difference: How the Power of Diversity Creates Better Groups, Firms, Schools, and Societies*, Princeton University Press, 2007.

Pager, Devah, Rebecca Goldstein, Helen Ho, and Bruce Western, "Criminalizing Poverty: The Consequences of Court Fees in a Randomized Experiment," *American Sociological Review*, Vol. 87, No. 3, June 2022.

Pahlka, Jennifer, "AI Meets the Cascade of Rigidity," in Erik Brynjolfsson, Alex Pentland, Nathaniel Persily, Condoleezza Rice, and Angela Aristidou, eds., *The Digitalist Papers: Artificial Intelligence and Democracy in America*, Stanford Digital Economy Lab, 2024.

Palazzolo, Stephanie, Erin Woo, and Amir Efrati, "OpenAI Shifts Strategy as Rate of 'GPT' AI Improvements Slows," *The Information*, November 2024.

Palermo, Frank, "Elusive AI ROI Doesn't Mean the Bubble Will Burst," *CMSWire*, September 23, 2024.

Palisade Research [@PalisadeAI], "OpenAI's o3 model sabotaged a shutdown mechanism to prevent itself from being turned off. It did this even when explicitly instructed: allow yourself to be shut down," post on the X platform, May 23, 2025. As of September 2, 2025: <https://x.com/PalisadeAI/status/1926084635903025621>

Pandi, Senthil S., A. Mohammed Farook, W. Kingston, and Sathish Kumar Kannaiah, "Enlightening Justice: Empowering Society Through AI-Driven Legal Assistance," paper presented at the 2024 Second International Conference on Advances in Information Technology, Chikkamagaluru, Karnataka, India, July 24–27, 2024.

Park, Peter S., Philipp Schoenegger, and Chongyang Zhu, "Diminished Diversity-of-Thought in a Standard Large Language Model," *Behavior Research Methods*, Vol. 56, January 2024.

Payne, Kenneth, "Artificial Intelligence: A Revolution in Strategic Affairs?" *Survival*, Vol. 60, No. 5, 2018.

Perez, Carlota, *Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages*, Edward Elgar, 2002.

Perry, Mark J., "Animated Chart of the Day: World's Top Ten Manufacturing Nations, 1970 to 2017," American Enterprise Institute, July 5, 2019.

Pethokoukis, James, "Humanity's Great Fight for the Future: Demographics vs. AI," American Enterprise Institute, February 5, 2024.

Petrone, Justin, "AI Leap 2025: Estonia Sets the Global Standard for AI in Education," e-Estonia, March 19, 2025.

Petty, Jackson, Michael Y. Hu, Wentao Wang, Shauli Ravfogel, William Merrill, and Tal Linzen, "RELIC: Evaluating Compositional Instruction Following via Language Recognition," arXiv, arXiv:2506.05205, June 5, 2025.

Phelps, Edmund S., *Mass Flourishing: How Grassroots Innovation Created Jobs, Challenge, and Change*, Princeton University Press, 2013.

Plasma Control Group, "Our 2024 Breakthrough in Nuclear Fusion with Artificial Intelligence," webpage, May 11, 2024. As of September 17, 2025: <https://control.princeton.edu/2024/05/our-2024-breakthrough-in-nuclear-fusion-with-artificial-intelligence/>

Poh, Olivia, and Newley Purnell, "Singapore's Epic Quest to Harness AI," Bloomberg, October 11, 2024.

Pomeranz, Kenneth, *The Great Divergence: China, Europe, and the Making of the Modern World Economy*, Princeton University Press, 2000.

Posner, Eric A., and Shivam Saran, "Judge AI: Assessing Large Language Models in Judicial Decision-Making," University of Chicago, Coase-Sandor Institute for Law & Economics, Research Paper No. 25-03, January 15, 2025.

Powell, David, *Nationhood and Identity: The British State Since 1800*, I. B. Tauris, 2002.

Putnam, Robert D., *Bowling Alone: The Collapse and Revival of American Community*, Simon and Schuster, 2000.

Qi, Liyan, "DeepSeek's Breakthrough Sparks National Pride in China," *Wall Street Journal*, January 31, 2025.

Quataert, Donald, *The Ottoman Empire 1700–1922*, 2nd ed., Cambridge University Press, 2005.

Raghavan, Manish, "Competition and Diversity in Generative AI," arXiv, arXiv:2412.08610, December 11, 2024.

Rai, Saritha, "Singapore's AI Push Charts Path Toward Localized Models," Bloomberg, May 27, 2025.

RAND Corporation, "The Sources of Societal Competitiveness: How Nations Actually Succeed in Long-Term Rivalries," RB-A499-1, 2022. As of September 3, 2025:
https://www.rand.org/pubs/research_briefs/RBA499-1.html

Rauch, Jonathan, *The Constitution of Knowledge: A Defense of Truth*, Brookings Institution Press, 2021.

Rauh, Joshua, "Playing with Fire: The Interest Coverage of US Federal Government Revenues," *Liberty Lens*, Substack, October 7, 2024.

Reed, Tina, "Hospitals Size Up Ways to Ensure Their AI Works," *Axios*, November 4, 2024.

Regalado, Antonio, "An AI-Driven 'Factory of Drugs' Claims to Have Hit a Big Milestone," *MIT Technology Review*, March 20, 2024.

Reisman, David, Nathan Glazer, and Reuel Denney, *The Lonely Crowd*, Yale University Press, 1950.

Renewable Energy Institute, "AI: The Secret to Unlocking the Potential of Renewable Energy?" webpage, undated. As of December 30, 2025:
<https://www.renewableinstitute.org/ai-the-secret-to-unlocking-the-potential-of-renewable-energy/>

Robey, Alexander, Zachary Ravichandran, Vijay Kumar, Hamed Hassani, and George J. Pappas, "Jailbreaking LLM-Controlled Robots," University of Pennsylvania, School of Engineering and Applied Science, 2024.

Romero, Alberto, "Harvard and MIT Study: AI Models Are Not Ready to Make Scientific Discoveries," *The Algorithmic Bridge*, Substack, July 15, 2025.

Rosen, Jill, "Robot That Watched Surgery Videos Performs with Skill of Human Doctor," *Johns Hopkins University Hub* blog, November 11, 2024.

Rosen, Jody, "The Knowledge, London's Legendary Taxi-Driver Test, Puts Up a Fight in the Age of GPS," *New York Times*, November 10, 2014.

Rosenberg, Bernard, Israel Gerver, and F. William Howton, *Mass Society in Crisis: Social Problems and Social Pathology*, Macmillan, 1964.

Rosenberg, Scott, "AI's Spooky 'Scheming' Skill," *Axios*, December 13, 2024.

Rosenbluth, Teddy, "This Chatbot Pulls People Away from Conspiracy Theories," *New York Times*, September 12, 2024.

Rotman, David, "How to Fine-Tune AI for Prosperity," *MIT Technology Review*, August 20, 2024.

Rowe, Jon, and Derek Partridge, "Creativity: A Survey of AI Approaches," *Artificial Intelligence Review*, Vol. 7, February 1993.

Roy, Olivier, *The Crisis of Culture: Identity Politics and the Empire of Norms*, trans. by Cynthia Schoch and Trista Selous, Oxford University Press, 2024.

Rubin, James P., and Darjan Vujica, "AI Is Supercharging Information Warfare: And America's Defenses Aren't Ready," *Foreign Affairs*, November 19, 2025.

Rubinstein, W. D., *Capitalism, Culture and Decline in Britain, 1750–1990*, Routledge, 1993.

Rumelt, Richard P., *The Crux: How Leaders Become Strategists*, PublicAffairs, 2022.

Ryutaro, Komiya, "Industrial Policy in Japan," *Japanese Economic Studies*, Vol. 14, No. 4, 1986.

Saad, Lydia, "Historically Low Faith in U.S. Institutions Continues," Gallup, July 6, 2023.

Sabin, Sam, "Courts Aren't Ready for AI-Generated Evidence," *Axios*, July 25, 2025.

Saez, Emmanuel, and Gabriel Zucman, "Trends in US Income and Wealth Inequality: Revising After the Revisionists," National Bureau of Economic Research, Working Paper No. 27921, October 2020.

Sakoh, Katsuro, "Japanese Economic Success: Industrial Policy or Free Market?" *Cato Journal*, Vol. 4, No. 2, Fall 1984.

Salganik, Matthew J., Ian Lundberg, Alexander T. Kindel, Caitlin E. Ahearn, Khaled Al-Ghoneim, Abdullah Almaatouq, Drew M. Altschul, Jennie E. Brand, Nicole Bohme Carnegie, Ryan James Compton, et al., "Measuring the Predictability of Life Outcomes with a Scientific Mass Collaboration," *Proceedings of the National Academy of Sciences of the USA*, Vol. 117, No. 15, April 2020.

Samuel, Alexandra, "I Built an AI Career Coach. I've Never Had a Better Coach," *Wall Street Journal*, June 27, 2025.

Santoso, Romin Adi [@romin991], "BloombergGPT vs. ChatGPT: An In-Depth Comparative Analysis," *FinanceAndCode* blog, Medium, April 19, 2024. As of September 4, 2025: <https://medium.com/@romin991/bloomberggpt-vs-chatgpt-an-in-depth-comparative-analysis-72d264ca3cd3>

Sato, Masahiro, "GAI: Generative Agents for Innovation," arXiv, arXiv:2412.18899, December 25, 2024.

Scharre, Paul, *Four Battlegrounds: Power in the Age of Artificial Intelligence*, W. W. Norton, 2023.

Scharre, Paul, *Future-Proofing Frontier AI Regulation*, Center for a New American Security, March 13, 2024.

Schechner, Sam, "How I Realized AI Was Making Me Stupid—and What I Do Now," *Wall Street Journal*, April 3, 2025.

Scheidel, Walter, *The Great Leveler: Violence and the History of Inequality from the Stone Age to the Twenty-First Century*, Princeton University Press, 2017.

Scheidel, Walter, "The Road from Rome," Aeon, April 15, 2021.

Schmid, Jon, Tobias Sytsma, and Anton Shenk, *Evaluating Natural Monopoly Conditions in the AI Foundation Model Market*, RAND Corporation, RR-A3415-1, 2024. As of September 4, 2025: https://www.rand.org/pubs/research_reports/RRA3415-1.html

Schneewind, J. B., *The Invention of Autonomy: A History of Modern Moral Philosophy*, Cambridge University Press, 1997.

Schneider, Jordan, and Lily Ottinger, "Anthropic's Dario Amodei on AI Competition," *China Talk*, Substack, February 5, 2025.

Schneider, Jordan, Angela Shen, Irene Zhang, Yiwen, Nicholas Welch, and Alexa Pan, "Deepseek: The Quiet Giant Leading China's AI Race," *ChinaTalk*, Substack, November 27, 2024.

Schneider, Jordan, Irene Zhang, Angela Shen, and Yiwen, "DeepSeek: The View from China," *ChinaTalk*, Substack, January 28, 2025.

Schoenegger, Philipp, Indre Tuminauskaite, Peter S. Park, and Philip E. Tetlock, "Wisdom of the Silicon Crowd: LLM Ensemble Prediction Capabilities Match Human Crowd Accuracy," arXiv, arXiv:2402.19379, February 29, 2024.

Schwarcz, Daniel, Sam Manning, Patrick Barry, David R. Cleveland, J. J. Prescott, and Beverly Rich, "AI-Powered Lawyering: AI Reasoning Models, Retrieval Augmented Generation, and the Future of Legal Practice," Minnesota Legal Studies Research Paper No. 25-16, University of Michigan Public Law Research Paper No. 24-058, March 2, 2025.

Schwartz, Eric Hal, "AI Educators Are Coming to This School—And It's Part of a Trend," *TechRadar*, December 23, 2024.

Schwartz, Reva, Apostol Vassilev, Kristen Greene, Lori Perine, Andrew Burt, and Patrick Hall, "Towards a Standard for Identifying and Managing Bias in Artificial Intelligence," National Institute of Standards and Technology, NIST Special Publication No. 1270, March 2022.

Scott, Jason M., Stephen N. Goggin, Rick Trachok, Jenny S. Kwon, Sara Gordon, Dean Gould, Fletcher S. Hiigel, Leah Chan Grinvald, and David Faigman, "Putting the Bar to the Test: An Examination of the Predictive Validity of Bar Exam Outcomes on Lawyering Effectiveness," *Journal of Law and Empirical Analysis*, Vol. 1, No. 2, July–December 2024.

Shah, Rohin, Alex Irpan, Alexander Matt Turner, Anna Wang, Arthur Conmy, David Lindner, Jonah Brown-Cohen, Lewis Ho, Neel Nanda, Raluca Ada Popa, et al., "An Approach to Technical AGI Safety and Security," arXiv, arXiv:2504.01849, April 2, 2025.

Shanmugam, Divya, Monica Agrawal, Rajiv Movva, Irene Y. Chen, Marzyeh Ghassemi, and Emma Pierson, "Generative AI in Medicine," arXiv, arXiv:2412.10337, December 13, 2024.

Sheets, Nathan, "An Aging Global Workforce Threatens Growth. AI Could Help," *Barron's*, April 15, 2024.

Shehabi, Arman, Sarah Josephine Smith, Alex Hubbard, Alexander Newkirk, Nuoa Lei, Mohammed AbuBakar Siddik, Billie Holecek, Jonathan G. Koomey, Eric R. Masanet, and Dale A. Sartor, 2024 *United States Data Center Energy Usage Report*, Lawrence Berkeley National Laboratory, December 19, 2024.

Shein, Esther, "The Impact of AI on Computer Science Education," *Communications of the ACM*, July 30, 2024.

Shoham, Yoav, "Don't Let Hype About AI Agents Get Ahead of Reality," *MIT Technology Review*, July 3, 2025.

Shojaee, Parshin, Iman Mirzadeh, Keivan Alizadeh, Maxwell Horton, Samy Bengio, and Mehrdad Farajtabar, "The Illusion of Thinking: Understanding the Strengths and Limitations of Reasoning Models via the Lens of Problem Complexity," arXiv, arXiv:2506.06941, June 7, 2025.

Shonk, Katie, "AI Mediation: Using AI to Help Mediate Disputes," Harvard Program on Negotiation, November 20, 2024.

Simshaw, Drew, "Access to A.I. Justice: Avoiding an Inequitable Two-Tiered System of Legal Services," *Yale Journal of Law and Technology*, Vol. 24, 2022.

Siripurapu, Anshu, "The U.S. Inequality Debate," Council on Foreign Relations, April 20, 2025.

Smith, Mary, "Legal Help for Civil Matters Shouldn't Be Reserved for the Rich," Bloomberg Law, January 11, 2024.

Smith, Noah, "Stop Pretending You Know What AI Does to the Economy," *Noahpinion*, Substack, July 20, 2025.

So8res, "Why All the Fuss About Recursive Self-Improvement," *LessWrong* blog, June 12, 2022.

Som, Lalita, "Japan's Growth Experience: Post–Second World War and Recent Times," in *The Capitals of Nations: The Role of Human, Social, and Institutional Capital in Economic Evolution*, Oxford University Press, 2014.

Spencer, Michael, "Goldman Sachs and Economists Are Backtracking on Generative AI's Value," *AI Supremacy*, Substack, July 3, 2024.

Spencer, Michael, and Grace Shao, "Is Alibaba's Qwen the Open-Source AI Winner?" *AI Supremacy*, Substack, November 13, 2024.

Stahl, Günter K., and Martha L. Maznevski, "Unraveling the Effects of Cultural Diversity in Teams: A Retrospective of Research on Multicultural Work Groups and an Agenda For Future Research," *Journal of International Business Studies*, Vol. 52, 2021.

Stahl, Günter K., Martha L. Maznevski, Andreas Voigt, and Karsten Jonsen, "Unraveling the Effects of Cultural Diversity in Teams: A Meta-Analysis of Research on Multicultural Work Groups," *Journal of International Business Studies*, Vol. 41, 2010.

Stokes, Gareth, "The Road to AI Efficiency: The Trend Toward Smaller More Performant AI and AI at the Edge," *Technology's Legal Edge: A Global Technology Sector Blog*, DLA Piper, May 17, 2024.

Stoller, Matt, "Why Are We Pretending AI Is Going to Take All the Jobs?" *BIG by Matt Stoller*, Substack, July 22, 2025.

Suarez, Fernando F., and Gianvito Lanzolla, "The Half-Truth of First-Mover Advantage," *Harvard Business Review*, April 2005.

Subramanian, Akshay, Wenhao Gao, Regina Barzilay, Jeffrey C. Grossman, Tommi Jaakkola, Stefanie Jegelka, Mingda Li, Ju Li, Wojciech Matusik, Elsa Olivetti, et al., "Closing the Execution Gap in Generative AI for Chemicals and Materials: Freeways or Safeguards," in *MIT Exploration of Generative AI: From Novel Chemicals to Opera*, Massachusetts Institute of Technology, March 27, 2024.

Sukharevsky, Alexander, Eric Hazan, Sven Smit, Marc-Antoine de la Chevasnerie, Marc de Jong, Solveigh Hieronimus, Jan Mischke, and Guillaume Dagarret, "Time to Place Our Bets: Europe's AI Opportunity," McKinsey and Company, October 1, 2024.

Suleyman, Mustafa, *The Coming Wave: Technology, Power, and the 21st Century's Greatest Dilemma*, Crown, 2023.

Sunderland, David, *Social Capital, Trust and the Industrial Revolution, 1780–1880*, Routledge, 2007.

Suri, Jeremi, "State Finance and National Power: Great Britain, China, and the United States in Historical Perspective," in Jeremi Suri and Benjamin Valentino, eds., *Sustainable Security: Rethinking American National Security Strategy*, Oxford University Press, 2016.

Swaminathan, Nandhini, and David Danks, "AI, Pluralism, and (Social) Compensation," arXiv, arXiv:2404.19256v2, October 15, 2024.

Swedberg, Richard, ed., *Entrepreneurship: The Social Science View*, Oxford University Press, 2000.

Sytsma, Tobias, and Éder M. Sousa, *Artificial Intelligence and the Labor Force: A Data-Driven Approach to Identifying Exposed Occupations*, RAND Corporation, RR-A2655-1, 2023. As of September 5, 2025: https://www.rand.org/pubs/research_reports/RRA2655-1.html

Tabarrok, Maxwell, "AGI Will Not Make Labor Worthless," *Maximum Progress*, Substack, January 9, 2025.

Tagaki, Koichiro, "Artificial Intelligence and Future Warfare," Hudson Institute, November 23, 2022.

Tamayo, Jorge, Leila Doumi, Sagar Goel, Orsolya Kovács-Ondrejkovic and Raffaella Sadun, "Reskilling in the Age of AI: Five New Paradigms for Leaders—and Employees," *Harvard Business Review*, September–October 2023.

Taylor, Charles, *Sources of the Self: The Making of the Modern Identity*, Harvard University Press, 2009.

Tegmark, Max, *Life 3.0: Being Human in the Age of Artificial Intelligence*, Knopf, 2017.

Telang, Ashwin, "The Promise and Peril of AI Legal Services to Equalize Justice," *Jolt Digest*, March 14, 2023.

Teslo, Ruxandra, "The City of Civilizational Greatness," *Ruxandra's Substack*, Substack, December 26, 2024.

Teslo, Ruxandra, "What Will It Take for AI to Change Drug Discovery?" *Ruxandra's Substack*, Substack, November 9, 2025.

Todt, Kiersten, and Peter Kant, "In a World Dominated by AI, Neurodiversity Matters More Than Ever," Center for Strategic and International Studies, January 30, 2024.

Togelius, Julian, *Artificial General Intelligence*, MIT Press, 2024.

Torres, Émile P., "GPT-5 Should Be Ashamed of Itself," *Realtime Technoapocalypse Newsletter*, Substack, August 9, 2025.

Trajtenberg, Manuel, "AI as the Next GPT: A Political-Economy Perspective," National Bureau of Economic Research, Working Paper No. 24245, January 2018.

Trueman, Charlotte, "Canadian Government Launches Sovereign AI Compute Strategy," *Data Centre Dynamics*, December 9, 2024.

Tsai, Lily L., and Alex Pentland, "Rediscovering the Pleasures of Pluralism: The Potential of Digitally Mediated Civic Participation," in Erik Brynjolfsson, Alex Pentland, Nathaniel Persily, Condoleezza Rice, and Angela Aristedou, eds., *The Digitalist Papers: Artificial Intelligence and Democracy in America*, Stanford Digital Economy Lab, 2024.

Turkle, Sherry, *Reclaiming Conversation: The Power of Talk in a Digital Age*, Penguin, 2015.

Turkle, Sherry, *Alone Together: Why We Expect More from Technology and Less from Each Other*, 3rd ed., Basic Books, 2017.

Twenge, Jean M., "Gen Z Really Does Have a Work Ethic Problem," *Generation Tech*, Substack, November 15, 2023.

Ukoba, Kingsley, Kehinde O. Olatunji, Eyiayo Adeoye, Tien-Chien Jen, and Daniel M. Madyira, "Optimizing Renewable Energy Systems Through Artificial Intelligence: Review and Future Prospects," *Energy & Environment*, Vol. 35, No. 7, November 2024.

U.S. Department of Homeland Security, "Leveraging AI to Enhance the Nation's Cybersecurity," October 17, 2024.

Valone, Matthew, "AI Will Not Cause 20% Unemployment and 10% GDP Growth in the Next Couple of Years," *Securely Spectating*, Substack, June 1, 2025.

Varaquaux, Gaël, Alexandra Sasha Luccioni, and Meredith Whittaker, "Hype, Sustainability, and the Price of the Bigger-Is-Better Paradigm in AI," arXiv, arXiv:2409.14160, September 21, 2024.

Verma, Pranshu, and Gerrit De Vynck, "Trump Pledged to Gut Biden's AI Rules, as OpenAI Eyes Landmark Infusion," *Washington Post*, November 13, 2024.

Vestal, James E., *Planning for Change: Industrial Policy and Japanese Economic Development 1945–1990*, Clarendon Press, 1995.

Villalobos, Pablo, Anson Ho, Jaime Sevilla, Tamay Besiroglu, Lennart Heim, and Marius Hobbahn, "Will We Run Out of Data? Limits of LLM Scaling Based on Human-Generated Data," *Epoch AI*, June 6, 2024.

Vinchon, Florent, Todd Lubart, Sabrina Bartolotta, Valentin Gironnay, Marion Botella, Samira Bourgeois-Bougrine, Jean-Marie Burkhardt, Nathalie Bonnardel, Giovanni Emanuele Corazza, Vlad Glăveanu, et al., "Artificial Intelligence & Creativity: A Manifesto for Collaboration," *Journal of Creative Behavior*, Vol. 57, No. 4, December 2023.

Vogel, Ezra, *Japan as Number One: Lessons for the United States*, Harvard University Press, 1979.

Vollrath, Dietrich, "Will AI Cause Explosive Economic Growth?" *Economic Growth Blog*, Substack, July 11, 2023.

Waldenstrom, Daniel, "The Inequality Myth: Western Societies Are Growing More Equal, Not Less," *Foreign Affairs*, May 19, 2025.

Walker, Joe, "Larry Summers—AGI and the Next Industrial Revolution," *Joe Walker Podcast*, October 22, 2024.

Wang, Fali, Zhiwei Zhang, Xianren Zhang, Zongyu Wu, Tzuhao Mo, Qiuhan Lu, Wanjing Wang, Rui Li, Junjie Xu, Xianfeng Tang, et al., "A Comprehensive Survey of Small Language Models in the Era of Large Language Models: Techniques, Enhancements, Applications, Collaboration with LLMs, and Trustworthiness," arXiv, arXiv:2411.03350, November 4, 2024.

Weise, Karen, and Cade Metz, "At Amazon's Biggest Data Center, Everything Is Supersized for A.I.," *New York Times*, June 24, 2025.

Weisheit, Sonja, and Christoph Salger, "Artificial Intelligence (AI) in Mediation—ChatGPT as Mediator 4.0," *Mediate*, June 21, 2023.

Weitzman, Tyler, "Empowering Individuals with Disabilities Through AI Technology," *Forbes*, June 16, 2023.

Westcott, Justin, "The AI Trust Imperative: Navigating the Future with Confidence," *Edelman*, 2025.

"What If AI Made the World's Economic Growth Explode?" *The Economist*, July 24, 2025.

White House, "AI Talent Survey," January 15, 2025a.

White House, *Winning the Race: America's AI Action Plan*, Executive Office of the President of the United States, July 2025b.

Whyte, William H., Jr., *The Organization Man*, Simon and Schuster, 1956.

Wiggers, Kyle, "DeepSeek's New AI Model Appears to Be One of the Best 'Open' Challengers Yet," *TechCrunch*, December 26, 2024.

Williams, Callum, "There Will Be No Immediate Productivity Boost from AI," *The Economist*, November 20, 2024.

Williams, Rhiannon, "AI Could Help People Find Common Ground During Deliberations," *MIT Technology Review*, October 17, 2024.

Williams, Rhiannon, "The AI Relationship Revolution Is Already Here," *MIT Technology Review*, February 13, 2025.

Winter-Levy, Sam, "The Emerging Age of AI Diplomacy: To Compete with China, the United States Must Walk a Tightrope in the Gulf," *Foreign Affairs*, October 28, 2024.

Witters, Dan, "U.S. Depression Rates Reach New Highs," Gallup, May 17, 2023.

Wolf, Martin, "Can a Fragmented Europe Continue to Prosper?" *Financial Times*, November 11, 2025.

Woodie, Alex, "What's Holding Up the ROI for GenAI?" *Big Data Wire*, April 3, 2024.

Wright, Candice N., and Kevin Walsh, "Artificial Intelligence: Generative AI Use and Management at Federal Agencies," Government Accountability Office, GAO-25-107653, July 2025.

Yamin, Muhammad Mudassar, Mohib Ullah, Habib Ullah, and Basel Katt, "Weaponized AI for Cyber Attacks," *Journal of Information Security and Applications*, Vol. 57, March 2021.

Young, Clifford, and Bernard Mendez, "How Americans Feel About the U.S.' Rising Income Inequality," Ipsos, February 21, 2025.

Zaharia, Matei, Omar Khattab, Lingjiao Chen, Jared Quincy Davis, Heather Miller, Chris Potts, James Zou, Michael Carbin, Jonathan Frankle, Naveen Rao, and Ali Ghodsi, "The Shift from Models to Compound AI Systems," Berkeley Artificial Intelligence Research, February 18, 2024.

Zequeira, Michael, "Artificial Intelligence as a Combat Multiplier: Using AI to Unburden Army Staffs," *Military Review*, September 2024.

Zhai, Chunpeng, Santoso Wibowo, and Lily D. Li, "The Effects of Over-Reliance on AI Dialogue Systems on Students' Cognitive Abilities: A Systematic Review," *Smart Learning Environments*, Vol. 11, No. 1, June 2024.

About the Author

Michael J. Mazarr is a senior political scientist at RAND. His primary interests are U.S. defense policy and force structure, disinformation and information manipulation, East Asian security, nuclear weapons and deterrence, and judgment and decisionmaking under uncertainty. He holds a Ph.D. in public policy.

