



ROBERT J. LEMPert, JONATHAN W. WELBURN, LAURENCE B. MUSSIO, MICHAEL ALDOUS

Applying History to Inform Anticipatory AI Governance

Using Foresight and Hindsight to Inform Policymaking

For more information on this publication, visit www.rand.org/t/CFA3591-1.

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.

© 2025 RAND Corporation

RAND® is a registered trademark.

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, please visit www.rand.org/about/publishing/permissions.

About These Conference Proceedings

Artificial intelligence (AI) heralds societal changes that could rival those associated with past transformational general-purpose technologies, such as metallurgy, the steam engine, electricity, and the internet. As with such technologies, AI offers the opportunity for tremendous increases in human well-being while also threatening to destabilize social, governance, economic, and critical infrastructure systems and disempower many humans.

These conference proceedings describe a June 2024 workshop organized by RAND and the Long Run Institute and held in Washington, D.C., that explored a novel combination of two methods—backcasting and the use of history—as one means of addressing the challenges associated with advancements in AI. *Backcasting*, an explicitly normative futures methodology aimed at exploring hopeful outcomes, invites participants to place themselves in a desired future and describe pathways that could lead to that destination. While backcasting is a powerful method, the societal transformations that AI makes possible could be so substantial that one’s imagination might have trouble encompassing them. Thus, this workshop aimed to improve on backcasting by adding an explicit and deep historical perspective.

The workshop aimed to serve as a pilot of a methodology that we hope can be employed more extensively with policymakers and other decisionmakers. These conference proceedings should be of interest both to those interested in novel futures methodologies and to those interested in AI governance.

Community Health and Environmental Policy Program

RAND Social and Economic Well-Being is a division of RAND that seeks to actively improve the health and social and economic well-being of populations and communities throughout the world. This research was conducted in the Community Health and Environmental Policy Program within RAND Social and Economic Well-Being. The program focuses on such topics as infrastructure, science and technology, industrial policy, community design, community health promotion, migration and population dynamics, transportation, energy, and climate and the environment, as well as other policy concerns that are influenced by the natural and built environment, technology, and community organizations and institutions that affect well-being. For more information, email chep@rand.org.

RAND Frederick S. Pardee Center for Longer Range Global Policy and the Future Human Condition

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 Frederick S. Pardee Center for Longer Range Global Policy

and the Future Human Condition, which serves as a RAND laboratory for methods innovation in Decision Making Under Deep Uncertainty, using center funding to catalyze research into new methods, seed new projects using these methods, and disseminate these methods within and outside RAND. For more information, contact RPC@rand.org.

Funding

Funding for this work was made available by RAND Frederick S. Pardee Center for Longer Range Global Policy and the Future Human Condition, which was established by a generous gift from Frederick S. Pardee in 2001.

Acknowledgments

We thank RAND's Social and Economic Well-Being division and the RAND Pardee Center for supporting this workshop; the workshop participants who gave so generously of their time and expertise; our RAND colleagues, Zara Abdurahaman, Benjamin Boudreaux, and Leah Dion, and the group Existential Hope for help with the scenarios; Joe Matthews for his engaging session; Anton Shenk, Ivana Ke, and Lily Hoak who took excellent notes during the workshop; and our reviewers Nidhi Kalra and Stephen Ryan whose feedback greatly improved the document.

Summary

Artificial intelligence (AI) heralds societal changes that could rival those associated with past transformational general-purpose technologies, such as metallurgy, the steam engine, electricity, and the internet. As with such technologies, AI offers the opportunity for tremendous increases in human well-being while also threatening to destabilize social, governance, economic, and critical infrastructure systems and disempower many humans.

These conference proceedings describe a workshop that explored a novel combination of two methods—backcasting and the use of history—as one means of addressing the challenges associated with advancements in AI. The workshop, held in Washington, D.C., in June 2024, was the result of a collaborative effort among the Long Run Institute, RAND Social and Economic Well-Being, and the RAND Frederick S. Pardee Center for Longer Range Global Policy and the Future Human Condition. The workshop aimed to serve as a pilot of a methodology that we hope can be employed more extensively with policymakers and other decisionmakers.

Backcasting, an explicitly normative futures methodology aimed at exploring hopeful outcomes, invites participants to place themselves in a desired future and describe pathways that could lead to that destination. Backcasting is valuable because it can help participants overcome the constraints that make it difficult to imagine a very different future and to consider near-term actions that might bring that future into being. While people will disagree on the ideal AI-enabled utopia, envisioning such utopias can prove useful if for no other reason than that the process can help identify near-term policies that are compatible with these many envisioned futures.

While backcasting is a powerful method, the societal transformations that AI makes possible could be so substantial that one's imagination might have trouble encompassing them. Thus, this workshop aimed to improve on backcasting by adding an explicit and deep historical perspective. History can help inject realism into imagination by subjecting speculation about the future to lessons from the past. History can also help liberate the imagination. While we can only speculate about the future, history confronts us with the reality that, in many cases, the present is very different than the past, which makes more concrete any suggestion that the future may be similarly different than the present.

Workshop participants—12 individuals from diverse backgrounds, including business leaders, policymakers, and technologists—were presented with two scenarios, each depicting a different future of AI-enabled human flourishing in the year 2045. The first, called *Rising Choir*, was adapted from a scenario developed by the group *Existential Hope*. *Rising Choir* envisions a highly decentralized society in which a sophisticated AI tool called *Voice for Open-Source Information and Community Engagement (V.O.I.C.E)* empowers a prosperous economy with an emphasis on AI-enabled local production and direct democracy across all levels of society, from local communities to the international community. To balance this scenario, the RAND team developed a second, less

decentralized scenario called The Singularis, in which AI-augmented humans and unaugmented humans live together in harmony and mutual benefit.

Workshop participants were also presented with three historical case studies focusing on the societal impacts of general-purpose technologies in the 19th and early 20th centuries. The first case study, Brave New Worlds, examined the impact of labor-displacing technologies. The second case study, Boom and Bust, examined the impact of such general-purpose technologies as electricity on productivity, innovation, and economic growth. The third case study, Monopoly's Moment and Markets of the Mind, examined how monopolies and trusts in the late 19th and early 20th centuries exerted a profound influence on markets, industries, and government policy.

Workshop participants discussed several potential areas for policy action, including developing more-sophisticated and more-nuanced AI governance models that strike a balance between encouraging innovation and ensuring the public good. Workshop discussions considered the potential for monopolistic tendencies in the AI sector, the need to enhance the government's ability to implement a more flexible regulatory system with more-effective collaboration among government and industry, and the possibility of pursuing multilateral international cooperation to avoid the fragmentation of regulation and markets thereby supporting a system in which technological advancements benefit a broader variety of stakeholders.

Our experience with this workshop suggests that engaging history during futures exercises can expand the imagination of those engaged in backcasting and ground their discussion in historical realism. The workshop also offers some lessons for future efforts. These include beginning with scenarios generated by participants instead of developing scenarios among workshop organizers, employing more-detailed and data-supported scenarios, engaging a wider variety of stakeholders, and formally designing the workshops to evaluate the benefits of including history in backcasting exercises. While this workshop provides only an initial exploration of the potential of historically informed visioning and backcasting, the exercise does suggest that the combination of these two approaches could prove important in addressing not only AI governance but a variety of areas in which technology might substantially transform society, including genetic engineering, bioengineering, nanotechnology, robotics, and climate change.

Contents

- About These Conference Proceedings iii
- Summary v
- Figures..... viii
- CHAPTER 1..... 1**
- Introduction 1
- CHAPTER 2..... 2**
- Conducting the Workshop 2
 - Workshop Methodology and Development 2
 - Workshop Description..... 3
 - Two Future Scenarios 4
 - Three Historical Case Studies..... 5
 - Workshop Discussions..... 6
- CHAPTER 3..... 11**
- What We Learned About the Role of History in Backcasting 11
 - Closing Observations 12
- APPENDIX A..... 13**
- Scenario Descriptions 13
 - The Rising Choir Scenario..... 13
 - The Singularis Scenario..... 14
- APPENDIX B 16**
- Historical Case Studies..... 16
 - Case Study 1: Brave New Worlds..... 16
 - Case Study 2: Boom and Bust 22
 - Case Study 3: Monopoly’s Moment and Markets of the Mind..... 27
- APPENDIX C..... 36**
- Workshop Agenda and Description 36
 - Workshop Agenda 36
 - Workshop Description..... 37
- Abbreviations 38
- References..... 39

Figures

- Figure B.1. Long-Run Economic Growth, Productivity, and Wages..... 18
- Figure B.2. The Engels Pause: Economic Growth vs. Real Wages 19
- Figure B.3. 10-Year Average Annual Growth in U.S. Total Factor Productivity..... 23
- Figure B.4. Share of Funds for U.S. Research and Development Spending..... 26
- Figure B.5. Gross Domestic Product Growth and Economic Performance (1870–1920) 29
- Figure B.6. Technological Innovations and Patents (1870–1920) 30
- Figure B.7. Corporate Concentration in Various U.S. Industries (1850–1920) 31
- Figure B.8. Economic, Social, and Regulatory Trends During the Progressive Era (1900–1920) 33
- Figure B.9. Impact of the Food and Drug Act (1906) on Market Quality and Consumer Confidence
and Safety..... 34

Introduction

Artificial intelligence (AI) heralds societal changes that could rival those associated with past transformational, general-purpose technologies, such as metallurgy, the steam engine, electricity, and the internet. As with such technologies, AI offers the opportunity for tremendous increases in human well-being while also threatening to destabilize social, governance, economic, and critical infrastructure systems and disempower many humans (Willis et al., 2024).

These conference proceedings describe a workshop, organized by RAND and the Long Run Institute (LRI), that explored a novel combination of two methods—backcasting and the use of history—as one means of addressing the challenges associated with advancements in AI.

Backcasting, an explicitly normative futures methodology aimed at exploring hopeful outcomes, invites participants to place themselves in a desired future and describe pathways that could lead to that destination. Backcasting is valuable because it can help participants overcome constraints that make it difficult to imagine a very different future and to consider near-term actions that might bring that future into being. This workshop used backcasting to build on previous work by the RAND Pardee Center (Pendleton-Jullian and Lempert, 2019; Lempert, 2019; Parson et al., 2019; Boudreaux and Lempert, 2023) that was rooted in the idea that envisioning desirable AI-enabled worlds can usefully inform policy and that policies informed by such visions are more likely to result in desirable outcomes than policies that are not. While people will disagree on the ideal AI-enabled utopia, envisioning such utopias can prove useful if for no other reason than that the process can help identify near-term policies that are compatible with these many envisioned futures.

While backcasting is a powerful method, the societal transformations that AI makes possible could be so substantial that one's imagination might have trouble encompassing them. Thus, this workshop aimed to improve on backcasting by adding an explicit and deep historical perspective. History can help inject realism into imagination by subjecting speculation about the future to lessons from the past. History can also help liberate the imagination. While we can only speculate about the future, history confronts us with the reality that, in many cases, the present is very different than the past, which makes more concrete any suggestion that the future may be similarly different than the present (Neustadt and May, 1986).

This workshop thus offered an initial exploration of the use of applied history to inform a futures exercise focused on illuminating near-term policy challenges in AI governance. Rather than focus on any doomsday scenario or existential risk from artificial general intelligence, the workshop considered how a powerful and widespread AI might create a broadly desirable future. We hope that the proceedings from this workshop suggest ways in which history-informed backcasting can prove broadly useful to such an endeavor while also offering some hypotheses regarding near-term policy and future research agendas. The workshop described in these proceedings aimed to serve as a pilot methodology that we hope can be employed more extensively with policymakers and other decisionmakers.

Conducting the Workshop

The workshop was the result of a collaborative effort among LRI, RAND Social and Economic Well-Being, and the RAND Frederick S. Pardee Center for Longer Range Global Policy and the Future Human Condition. RAND hosted the daylong event on June 12, 2024, at its Washington, D.C., office. The workshop had 12 participants who included senior business leaders, financiers, two former members of Congress (one Democratic and one Republican), technologists, executive branch staff, and policy experts. Most participants had some prior connection to RAND or LRI either as members of advisory boards or in related roles. Staff from LRI and RAND facilitated the event.

Workshop Methodology and Development

The workshop employed the futures methodology of backcasting (Robinson, 1988; Robinson, 2003). As practiced at the workshop, backcasting begins with one or more desirable futures, and then asks participants to imagine themselves in those futures and describe different paths that could take society from the present day to each of the future scenarios. Rather than examine doomsday scenarios or the existential risks from artificial general intelligence, workshop participants directed their focus on how a powerful and widespread AI might create a broadly desirable future. Participants were given scenarios that described a future (set in the year 2045) in which AI had enabled a world of human flourishing. Participants were then asked several questions to create a narrative to describe how this future was achieved: What were key actions that happened from 2025 to 2030? How did subsequent events play out? What surprises helped or hindered the pathway to this world of human flourishing? Participants were asked to consider how events evolved over multiple domains, including technology and science, values and community, and economy and institutions.

In many cases, it is difficult to think about the future because of the rapid proliferation of branch points and alternatives as one projects further ahead in time. Backcasting can prove useful because it encourages participants to begin with a specific future and think backward, thus limiting the number of pathways to consider. Additionally, participants in futures exercises often find themselves anchored in the present, unable to allow their imaginations to stray far from current conditions and trends. Backcasting encourages participants to anchor themselves in what might be a very different future. The approach is also unambiguously normative, focusing on futures that at least some would regard as desirable. Forecasting the most likely future can have the effect of diminishing agency because the passive question “What is most likely to happen?” can distract from more-active questions, such as “What would we like to have happen?” and “How can we make it so?” Backcasting is designed to focus on these latter questions.

Backcasting exercises often begin with scenarios that are already familiar to and embraced by the participants. For instance, the Road to Zero study conducted by RAND researchers gathered

participants who were already committed to the goal of eliminating all traffic fatalities in the United States (Ecola et al., 2018). Other exercises begin by asking participants to develop their own future scenarios (Lempert et al., 2021). In this AI workshop, however, participants neither embraced existing scenarios nor had time to develop their own. Thus, the project team needed to provide participants with initial scenarios that they would find provocative and plausible. This proved to be challenging because there was no consensus or much prior consideration among the group about what might constitute an AI-enabled scenario of human flourishing.

The project team thus developed two scenarios for workshop participants that depicted different futures marked by AI-enabled human flourishing. The first was adapted from the group Existential Hope, a project of the Foresight Institute that seeks to envision positive futures.¹ Existential Hope sponsored a competition for scenarios of AI-enabled human flourishing. Participants were early career professionals working in the AI space and developed their scenarios using a technique called *world building* (Pendleton-Jullian and Brown, 2018).² We chose one of the winning scenarios, called Rising Choir, which envisioned a highly decentralized society in which a sophisticated AI tool called V.O.I.C.E (Voice for Open-Source Information and Community Engagement) enables a prosperous, small-scale local capitalism and direct democracy across all levels of society, from local communities to the international community. To balance this scenario, the RAND team conducted an internal to RAND world-building exercise to develop a second, less decentralized scenario called The Singularis, in which AI-augmented humans and unaugmented humans live together in harmony and mutual benefit.

Workshop Description

The workshop unfolded over seven hours (9 a.m. to 4 p.m.). Prior to the workshop, participants were given a one-page description of the event, a two-page description of each scenario, and a 30-page description of the historical case studies.³

The workshop began with introductions of the project team and participants, an overview of the exercise, and a presentation of the scenarios. Participants then gathered into two pre-assigned breakout groups to develop their backcasting pathways. One group was given the Rising Choir scenario and the other was given The Singularis scenario. We asked participants to imagine themselves as authors of a 2045 best-selling book that described how, much to the surprise of what people had expected in 2025, AI had enabled a world of human flourishing. After an hour, the breakout group participants reassembled in a plenary session to report on their results. To add realism to the session, we recruited a professional journalist, active and well-known on the AI beat, to interview the rapporteurs from each breakout group and act as if they were conducting a media interview with the author of a new best-selling book.

Over lunch, LRI staff presented three historical case studies of previous societal impacts and responses to major technological transformations: Brave New World, Boom and Bust, and

¹ Existential Hope, undated.

² A member of the RAND team, Robert J. Lempert, instructed the Existential Hope participants in the world-building method.

³ Workshop materials are available in the appendixes.

Monopoly's Moment and Markets of the Mind. The breakout groups then reconvened to consider these historical examples and rework their backcasting pathways. The groups reassembled in a final plenary session to report back on their revised pathways. The workshop concluded with a discussion of the near-term policy implications of achieving an AI-enabled future of human flourishing.

Two Future Scenarios

Each of the two scenarios given to the workshop participants represented one vision for AI-enabled human flourishing in 2045.⁴ The first scenario, Rising Choir, focused on a society with highly decentralized, democratized decisionmaking and enhanced productivity through AI-driven tools and robots. The second scenario, The Singularis, envisioned a balanced coexistence between unaugmented biohumans and AI-enhanced humans in a society in which AI systems optimize resource allocation and governance.

In the Rising Choir scenario, 2045 is characterized by enormous, technologically driven transformations in both society and individual's daily lives. Tightly regulated AI-driven humanoid robots, averaging two per person, handle over 95 percent of daily tasks, significantly boosting productivity and expanding the global economy twentyfold. Energy sustainability has been achieved through advanced solar technology and battery storage, reducing costs and producing an energy surplus capable of supporting environmental initiatives, such as carbon capture. And Universal Basic Compute and Universal Basic Income have been established to ensure societal engagement and economic security.⁵

In this scenario, AI has further revolutionized society through the implementation of V.O.I.C.E., an advanced tool that democratizes decisionmaking by gathering and synthesizing public opinions. The tool supports direct democracy, allowing billions to participate in decisions varying from local governance to international treaties. V.O.I.C.E. fosters inclusivity and representation, reducing cultural clashes and elevating living standards. The open-source nature of V.O.I.C.E. promotes continuous improvement and transparency.

Such sectors as education, transportation, and health care have also been transformed by AI and robotics, delivering efficient and sustainable services. Despite these advancements, such challenges as ensuring vote anonymity, safeguarding against bots, and promoting equitable technology access persist. The overarching goals of the society in this scenario include eliminating poverty, fostering peace, and combating climate change.

In The Singularis scenario, civilization has reached a pivotal juncture in which unaugmented biohumans and cybernetically enhanced AI-humans coexist harmoniously. Advanced AI systems manage resource allocation, task delegation, and transparent governance, including participatory liquid democracy models.⁶

⁴ The full descriptions of these scenarios are available in Appendix A.

⁵ *Universal Basic Compute* is an idea proposed by OpenAI Chief Executive Officer Sam Altman that would provide universal access to a predetermined amount of computing resources. *Universal Basic Income* is a form of social welfare that provides a regularly distributed, minimum income without qualification restrictions.

⁶ A *liquid democracy* is a hybrid form of government in which an electorate is empowered to vote on all policy issues, and individuals have the choice of either voting on issues directly or delegating their votes to a representative.

Work has been redefined as a minimal yet meaningful contribution to society, with individuals dedicating just ten hours weekly to community-driven projects. The abundance of free time fosters self-actualization through education, artistry, and technological pursuits. Socioeconomic disparities are mitigated through Universal Basic Income, supported by sophisticated AI systems that ensure efficient production and fair resource distribution.

Health care has undergone a renaissance, leveraging bioengineering, cybernetics, and AI-enabled personalized medicine to extend longevity and improve quality of life. Recreation has also been transformed by hyperhuman athletic leagues and avant-garde art movements that blend biological and technological elements.

The Singularis scenario also presents challenges, such as existential AI risks, the threat of societal bifurcation between augmented and unaugmented humans, and the need for evolving ethical frameworks in a posthuman society. Nevertheless, society remains committed to sustainable abundance, unity across sentient diversity, and equitable access to transformative growth opportunities.

The scenario features two characters—Aidhan, an AI-enhanced human, and Nedd, an unaugmented biohuman—who exemplify this future society through their daily interactions, showcasing a civilization that values inclusivity, sustainability, and diverse contributions.

Three Historical Case Studies

Workshop participants were also presented with three historical case studies focusing on the societal impacts of several general-purpose technologies from the 19th and early 20th centuries. These case studies were chosen because of their potential relevance to the future of AI.⁷

The first case study, *Brave New Worlds*, examined the impact of labor-displacing technologies. The case study considers the 19th and 20th centuries when, alongside positive impacts on economic growth, industrialization led to periods of job displacement, wage stagnation, increased concentration of economic power, and growing income inequality. The case study also provides historical examples of how these trends fueled social and political movements that pursued wildly different approaches to economic redistribution.

In relation to AI, the case study posed questions including: Will widespread adoption of AI lead to significant worker displacement and job losses? Will these labor disruptions result in a further concentration of economic power? What are the social and political consequences of technological change, and what historical precedents can be observed regarding the public's response to such a change? What role can government play in ensuring a fairer distribution of the benefits resulting from technological progress?

The second case study, *Boom and Bust*, examined the impact of such general-purpose technologies as electricity on productivity, innovation, and economic growth. The case study demonstrates how, historically, general-purpose technologies have had an overall transformative impact on economic growth, but the effects were felt in stages because it takes time for the technology to be widely adopted and refined. These processes required significant capital expenditure, and both

⁷ Appendix B contains the case studies that were presented to participants.

government and private sources of investment played a role. But the case study also shows that, across the 20th century, new technologies drove stock market bubbles and crashes that, in turn, affected investment in these technologies.

In relation to AI, this case study posed questions including: Will AI be a true general-purpose technology with a transformative impact on productivity and economic growth, or will its effects be more limited? How can we ensure that investment in AI research and development (R&D) is sufficient to drive sustained economic growth and innovation? How can we mitigate the risk of AI bubbles, characterized by rapid investment followed by a sharp decline in value?

The third case study, *Monopoly's Moment and Markets of the Mind*, examined how monopolies and trusts in the late 19th and early 20th centuries exerted a profound influence on markets, industries, and government policy. This case study underscored the inherent tension between technological innovation and regulation and the risks of monopolies controlling essential industries and markets, highlighting the potential impact of AI on governance and social order. The case study posed several key questions: What are the potential implications of AI for market structures and the rise of digital platforms as “cloud empires” (i.e., private corporations functioning as virtual states that can organize markets and set rules)⁸ with a dominant market position? How can we design regulatory frameworks for AI that promote competition and innovation while mitigating the risks of monopolies and unintended social consequences? What are the ethical implications of AI with respect to privacy, social influence, and the manipulation of public opinion?

Workshop Discussions

Participants were divided into two breakout groups, one for each scenario. The breakout groups met once in the morning, returned to present their initial discussions to the full group, heard presentations of the historical case studies, and then, in the afternoon, returned to their breakout groups. In both cases, the focus of the morning discussions differed from those of the afternoon.

Rising Choir Morning Discussion

The morning Rising Choir breakout group spent much of their time challenging scenario assumptions about the development and impacts of the V.O.I.C.E. platform. Participants noted that, similar to the hypothetical V.O.I.C.E. tool, existing platforms (such as TikTok) shape public opinion without centralized control. It was noted, however, that TikTok is run by a single organization, whereas the Rising Choir scenario envisions a pathway in which radical decentralization creates the V.O.I.C.E. tool, which, in turn, generates beneficial outcomes for society. Participants questioned the plausibility of such a pathway, for instance, the extent to which such a tool would value expert information and promote efficient resource allocation.

Participants also pointed to national differences as limiting factors to the development of a decentralized, global AI platform, noting how AI regulation and governance affect data ownership, technology monopolies, free speech, and international norms. The United States' limited participation

⁸ See Lehdonvirta, 2022.

in such international organizations as the United Nations or International Telecommunication Union and its unclear AI policy, especially compared with such authoritarian states as China, were highlighted as concerns. Participants also noted that AI has had greater democratization in the United States than elsewhere, but the government has been much slower to react to its technological advancements. In the United States, large technology companies monopolize AI tools, raising questions about how the government can regulate AI while advancing technology and productivity.

The group also questioned the Rising Choir scenario's assumption that AI would help improve global productivity and expand the global economy twentyfold by 2045, noting that technological advancements have not always translated into productivity gains and that reworking and retraining the labor force will be essential. Participants emphasized the need for broader fluency with technology and AI, and they raised concerns about the potential impacts of changing demographics. They argued that, because China's population might decrease to 700 million by 2060 and similar population declines are expected in Western Europe and South Korea, the Rising Choir scenario must assume that both robots and increases in productivity will compensate for these population declines, or the scenario must assume a geopolitical shift toward regions with younger populations, such as Latin America and Africa.

Singularis Morning Discussion

The morning Singularis breakout group laid out a pathway to the scenario and challenged scenario assumptions. This scenario focused on AI-enabled human augmentation. Participants sketched a pathway in which posthuman augmentation occurs through medical enhancements. They imagined a future in which patients begin to demand AI-augmentations as treatments for chronic diseases or disabilities, and regulatory authorities begin to approve such augmentations. In their vision, these treatments would turn into augmentations that go well beyond rectifying diseases or disabilities, and, as such, the augmentations would create posthuman capabilities.

The participants' pathway raised questions about the boundary between medical treatments and augmentation and the rules that might govern either sort of procedure. Participants noted existing norms governing human augmentation; for instance, athletes can rigorously control and enhance their diet but are forbidden to use some types of pharmaceuticals, such as growth hormones. Noting that 60 percent of Americans have some type of chronic condition, participants wondered how treatment and augmentation services would be allocated and who gets to decide who gets augmented and what sort of augmentations are allowed. For instance, what about augmentations that might allow people to become better criminals?

Participants also challenged the scenario's assumptions about leisure time, which envisioned a ten-hour work week with the remaining time devoted to self-realization and community service. The participants noted that some people take satisfaction from work and some people's self-realization involves the status that comes from work. Participants also wondered how economic surplus would be distributed in a society that focused less on paid work.

Rising Choir Afternoon Discussion

In the afternoon, discussions of both scenarios shifted to new ground, moving from challenges to scenario assumptions and a focus on the current Overton Window (i.e., the window of discourse that frames practically pursuable policy issues) toward a broader view of futures and policies.

The afternoon Rising Choir breakout group's discussion turned to how institutions that were established in the 20th century might need to be reworked for a world of new technology and new workers. Participants noted that past revolutions driven by railroads, telegraphs, the cotton gin, and other technologies fundamentally changed society and governance. However, the group argued that AI will not be as transformational as some expect and will certainly not be a technological singularity.⁹ Instead, they suggested that AI could contribute to transformative outcomes on society.

The participants discussed how existing societal distrust of all kinds of institutions has led to the increasing rejection of centralized authority and an increased interest in more decentralized governance. Citing the work of Elinor Ostrom (2015), participants noted historical examples of people coordinating very complicated and complex outcomes without top-down state governance. While participants doubted that the Westphalian system (i.e., the current system of international order founded on the principle that each state is sovereign) is going to fade away, they noted that after decades of governance stasis, we might be entering a period of tremendous change. For instance, participants envisioned the development of a global commons that connects people all around the world with regulatory frameworks very different than those currently designed and enforced by national governments (for instance, participants mentioned crypto as an early example of such a new system). Participants also called for new metrics for human well-being that could help society to break out of outdated 20th century institutional arrangements. They also considered whether addressing climate change could be a driver of global actions that might reset the balance of authority among local, national, and global institutions.

Participants also discussed the potential for unintended consequences of technological breakthroughs and a broad variety of potential social and political ramifications. The discussion raised several questions, including: What are the implications of significant extensions of life expectancy? How will people try to game emerging systems? Participants discussed the potential for some form of political or economic shock that provokes a social backlash and brings new social and political values to the fore, which, in turn, prompt the use of technology to achieve diffuse, radically decentralized forms of social and political organization. Such changes could lead to undesirable end states despite the utopian framing of the Rising Choir scenario.

Singularis Afternoon Discussion

The Singularis afternoon breakout group focused on regulatory and investment challenges. They noted how, in the past, regulation has slowed innovation. Some participants suggested that regulation killed innovation in railroads (though others suggested that unionization was the cause). Looking forward, participants suggested that the most difficult challenge might be aligning “markets of the

⁹ A *technological singularity* is a theoretical moment in time when technology, such as AI, escapes human control and engages in a recurring cycle of self-improvement, creating a superintelligence that surpasses all human intelligence. See Bostrom, 2014.

mind”—that is, the production, distribution, and consumption of ideas that influence moral, political, and economic sentiments. They discussed how business people and politicians lead the diffusion of technology into society and suggested that these people would need to find an appropriate equilibrium with checks and balances that both support and regulate the diffusion of AI technology throughout society. Participants noted that it is going to be very challenging for government to keep up with technological advancements given the pace of innovation assumed in this scenario.

Participants noted efforts by market incumbents to take advantage of what some regarded as the chaotic and ill-informed nature of discussions around AI. The participants discussed the example of how, in the early 2020s, many scientists signed open letters warning of the existential risks of AI. Large incumbent corporations then proposed regulatory responses that could be interpreted as naked attempts at regulatory capture through licensing regimes. Participants also noted confusion about the amount of capital required to be competitive in this space—yet, as Deep Seek demonstrates, the actual capital investments needed to develop useful AI tools may be significantly lower.

These considerations led the group to call for a more dynamic, informed, and flexible regulatory system. One barrier to such a regulatory system, however, is the lack of government capacity and AI expertise. Participants noted the significant shift of R&D in the United States in recent decades, as private sector contributions have grown while public sector contributions have remained static. The group argued that universities lack the funding and infrastructure to even reproduce results in commercial labs, much less drive innovation, and many top researchers move to industry. Governments, they argued, have access to few sources of cutting-edge technical advice beyond those people who work for the large technology companies.

Plenary Session Discussion

After the breakout sessions, the workshop concluded with participants gathering to report back on their discussions and reflect on what they had learned.

Many workshop participants expected that a significant concentration of economic power could accompany the rise of AI, along with the emergence of monopolies in essential services and platforms. They discussed how such concentration might echo past U.S. experience with monopolies and trusts and could pose risks to competition and fair market dynamics. They also noted how, as AI becomes deeply integrated into our economic systems, dependence on a limited number of providers—for example, in cloud computing or data storage—could create systemic risks. Participants expressed concern that the government’s current capacity to regulate AI is and will likely continue to be insufficient to address these challenges. Some participants also suggested that the growing dependence on AI services could ultimately lead to their classification as essential utilities, which would require careful consideration of regulatory frameworks, potentially mirroring the historical approach to utilities in ensuring affordability, reliability, and nondiscrimination.

Workshop participants also noted the decline of government investment in R&D compared with much larger private sector investments. This allocation leaves private investment and the volatile capital markets as the primary sources of capital to drive innovation. Private investment is often driven by short-term profit motives, which in the past, has led to boom-and-bust cycles and a neglect of long-term R&D. Overreliance on private sector R&D could contribute to further concentration of

economic power. Finally, participants discussed how a shortage of technical expertise outside the largest firms and, in particular, in universities and government labs could reduce the government's capacity to regulate with agility and effectiveness.

Participants agreed that, while AI promises numerous benefits, its successful integration into society will also require profound societal and cultural adaptation. They noted, however, how historical precedents show that rapid technological shifts often lead to increased inequalities and displacement, triggering social unrest and political polarization. Similar to past technological revolutions, AI adoption will likely reshape social and cultural norms, and participants emphasized that understanding the ethical considerations surrounding AI is crucial.

These discussions highlighted several themes related to future policy. First, participants grappled with the need to **develop more-sophisticated and more-nuanced AI governance models that strike a balance between encouraging innovation and ensuring the public good**. One concern involved addressing potential monopolistic tendencies in the AI sector. Traditional antitrust measures might struggle to meet the challenges of AI while monopolies form rapidly because of network effects and economies of scale. Approaches to antitrust might need to be reenvisioned. Future discussions could explore collaborative governance approaches, which might involve creating systems in which public and private stakeholders work together to prevent monopolistic control while encouraging innovation. Governance models should also balance competition with utility, ensuring that AI development does not come at the expense of public trust or societal welfare. Strengthening public oversight in technology could prove crucial to ensuring that AI development aligns with societal goals rather than being dominated by private interests.

Second, participants discussed **the importance of multilateralism in developing regulatory and governance frameworks**. In the context of AI governance, multilateralism becomes even more critical as the technology transcends national borders. While current geopolitical trends favor bilateralism, a cooperative, multilateral approach might help to address AI's global implications. While policy in this area needs to address differing views on the allocation of power among national governments, international organizations, businesses, and civil society groups, international cooperation could be key to avoiding fragmentation, creating a system in which technological advancements benefit a broader variety of stakeholders, and preventing monopolization by a few, powerful entities.

What We Learned About the Role of History in Backcasting

The interjection of historical case studies into the backcasting exercise appeared to play an important role in the workshop, helping to shift the focus of the breakout group discussions. In the morning, the groups focused more on critiquing the scenarios and sketching relatively familiar pathways. In the afternoon, however, the discussions took a broader view, which was consistent with the case studies that highlighted the extent to which the past, and hence the future, can be very different than the present.

The historical case studies also provided essential guideposts and points of reference for the participants, particularly when engaging in pattern recognition within episodes of technological transformations. Throughout the discussions, a central theme was the recognition that, although each technological wave is distinct, the broader contours of change—such as periods of rapid growth followed by regulatory challenges or market corrections—are often similar. This understanding encouraged participants to draw on historical precedents to interrogate how society might incorporate and respond to AI. For example, will AI's initial promise be tempered by growing concerns about monopolistic control, regulatory capture, and societal disruption?

The diversity of expertise among the workshop participants proved useful, contributing to a divergence in perspectives, with different historical lessons resonating more deeply depending on each attendee's professional background. For example, workshop participants exhibited diverse perspectives when discussing the political economy of technology in terms of the role of the state and regulation and ways to promote innovation, mitigate market concentration, and improve social outcomes.

Our experience with this workshop also suggests some lessons for future efforts.

- **Begin with scenario generation by participants.** Future workshops might usefully provide opportunities for participants to create their own scenarios. This would make the workshop longer—perhaps extending it from one day to two. But participant-generated scenarios would allow for a broader variety of perspectives and a greater sense of ownership over the process. This could improve engagement and ensure that the scenarios address the most-pressing concerns of stakeholders from different sectors and backgrounds.
- **Employ more-detailed and data-supported scenarios.** The scenarios used in the workshop provided a useful starting point but could have offered a richer and more interconnected vision of an AI-enabled future. Such scenarios could explore AI's impact on political economy, labor markets, and global power dynamics; potential shifts in wealth distribution; challenges to existing economic models; and emerging inequalities. A workshop might start with initial

scenario ideas constructed by the research team and supported by analysis in the literature. Participants could then tailor the scenarios before beginning the backcasting exercise.

- **Engage a wider variety of stakeholders.** The workshop suggests the importance of bringing together diverse perspectives when considering the challenges and opportunities emerging from technological transformations. The historical cases show complex interactions among technology, economy, and society and the often unintended outcomes of policy interventions. The workshop did bring together experts and policymakers, but future iterations could include a broader variety of voices, such as representatives from civil society, industry, academia, and everyday citizens. In particular, ensuring that the perspectives of marginalized communities and underrepresented regions are considered will lead to more-equitable and more-sustainable governance outcomes.
- **Evaluate the benefits of including history in backcasting exercises.** This exploratory workshop did not include any formal evaluation of the influence of the historical case studies on the participants' discussions and their ability to imagine a broader variety of futures and near-term policies. Future efforts might employ such evaluations, which could include pre- and post-surveys of workshop participants, content analysis of workshop discussions, and the comparison of multiple workshops in which some participants receive historical case studies and others do not.

Closing Observations

The pilot workshop suggested that applied history might significantly enrich a futures exercise, helping to generate expanded vision, refine questions, and open new perspectives. Integrating historical analysis into forward-looking AI discussions allowed for a more-comprehensive exploration of both the opportunities and risks of technological advancements, offering insights into governance, regulation, and the future of technology-driven societal shifts.

This workshop offers only an initial exploration of the potential of historically informed visioning and backcasting. Future efforts could do more to generate a richer and more definitive set of policy recommendations. For instance, future efforts might envision a series of workshops in which participants first suggest policy approaches, then use research to interrogate the suggested approaches, and conclude by discussing the results of the initial workshops in subsequent meetings. Such an approach might yield comprehensive sets of policy recommendations for AI governance, including recommendations on regulating competition, fostering equitable access to AI technologies, shaping the evolution of AI technology, and ensuring that AI systems serve the public rather than exclusively private interests.

Future efforts might include broader dialogue and collaboration. For instance, they might include participants associated with multilateral institutions, such as the United Nations, the Group of 20, or the Organisation for Economic Co-operation and Development, to help align global governance efforts and promote cross-border solutions. Future efforts might also include civic participation through a citizen assembly or deliberative polling.

While this workshop focused on AI governance, there are a host of other technological changes that present similar challenges for policymakers. The methodology employed in this RAND-LRI workshop could be applied to such issues as genetic engineering, bioengineering, nanotechnology, robotics, and climate change.

Scenario Descriptions

This appendix reproduces the descriptions of the two scenarios, Rising Choir and The Singularity, distributed to participants prior to the workshop. The scenarios have been lightly edited to format the text and correct typos.

The Rising Choir Scenario

RISING CHOIR

a symphony of clashing voices

In the year 2045, significant advancements in artificial intelligence have transformed both societal functions and individual daily lives. At the heart of this transformation is V.O.I.C.E. (Voice for Open-Source Information and Community Engagement), a sophisticated AI-driven tool that collects, analyzes, and synthesizes public opinions on a grand scale. This tool champions direct democracy by engaging communities to express preferences and propose balanced solutions, capable of scaling to accommodate the input of billions globally. This allows for informed decision-making across all levels of society, from local community decisions to international treaties.

In the economy, AI's influence is profound. Humanoid robots, averaging two per person, handle more than 95% of daily tasks, thereby multiplying global productivity and expanding the global economy twentyfold. These robots are crucial for both household and specialized tasks, operating under strict regulations that prevent over-the-air updates due to past security concerns. Notably, AI has not evolved into AGI, constrained by both regulatory and developmental limits.

Energy sustainability is another hallmark of 2045, with solar technology and battery storage having reduced costs by 90% over two decades. This surplus energy not only supports the extensive use of AI but also powers carbon capture technologies, mitigating environmental concerns.

V.O.I.C.E. has also cultivated an era of radical acceptance, significantly reducing cultural clashes. Its seamless integration into every societal layer has elevated living standards and reinforced a sense of inclusivity and representation. Recognizing Universal Basic Compute as a human right underscores its necessity for societal engagement in 2045, complemented by other safety nets like Universal Basic Income.

V.O.I.C.E. was developed as an open-source project, allowing for transparency and collaboration among a global community of developers. This approach ensures continuous improvement of the software, driven by a diverse group of contributors who can audit, modify, and enhance the system.

V.O.I.C.E. enables Institutional evolution which merges features of decentralized autonomous organizations, open-source platforms, and global governance bodies. This evolution has transformed

sectors like education, transport, and health care, which now leverage AI and robotics to deliver services efficiently and sustainably.

The scenario also explores a day in the life of Sofia, a beneficiary of this advanced society. Working as a ‘Calibrator,’ she fine-tunes humanoid robots for specific household tasks, ensuring precision in activities ranging from routine chores to personalized care such as tending to unique plant needs or preparing traditional dishes. Her interactions with V.O.I.C.E. reflect a vibrant civic life where even minor preferences like installing drinking fountains are promptly and democratically addressed.

Despite the seemingly idyllic conditions, challenges remain, particularly in the technological and societal integration of V.O.I.C.E. Issues like ensuring vote anonymity, safeguarding against bots, and promoting equitable technology access are pivotal. The overarching goals for 2045 focus on achieving prosperity by eliminating poverty, fostering peace to prevent wars, and combating climate change through sustainable practices.

In summary, “Rising Choir” envisions a 2045 where technological advancements and human values merge to forge a harmonious, prosperous, and inclusive world. The integration of AI in daily life and global governance, exemplified by tools like V.O.I.C.E., enhances democratic participation, cultural acceptance, and economic stability, setting an optimistic course for the future of humanity.

Acknowledgements: Rising Choir was adapted from the scenario by that name developed in a World Building exercise by the group Existential Hope. The RAND team edited the scenario with the help of ChatGPT.

The Singularis Scenario

THE SINGULARIS

a post-human civilization realized

In the annals of 2045, humanity’s journey has converged upon an unprecedented crossroads—the Singularis. This future world bears witness to the harmonious interplay between rapidly accelerating technologies and an unwavering ethos of inclusive human flourishing.

At its core, Singularis exemplifies a daring social experiment—a meticulously balanced civilization where both unaugmented bio-humans and cybernetically enhanced AI-humans coexist. Central to this vision is the seamless integration of advanced artificial intelligence systems that optimally allocate resources, delegate tasks, and facilitate transparent governance including participatory liquid democracy models.

Work, long regarded as an existential obligation, has been reframed into an ethos of meaningful yet minimal contribution. People across Singularis devote a mere ten hours weekly to community-driven projects while abundant free time catalyzes self-actualization through education, artistry, and the cultivation of technological pursuits. Socioeconomic disparities that once plagued previous eras have been mitigated through universal basic income. This paradigm shift is grounded in the advent of sophisticated AI systems that have solved monumental economic challenges, orchestrating efficient production, distribution, and fair resource allocation.

Yet the marvels of Singularis extend far beyond its socioeconomic architecture. A renaissance in sectors like health care has harnessed the synergies of bioengineering, cybernetics, and AI-enabled

personalized medicine. Bio-digital interfaces, nanorobotics, and gene-tailored therapies have radically extended longevity and quality of life. Even the frontiers of human recreation have been redrawn—from ability-stratified “Hyperhuman” athletic leagues celebrating the full potentials of the AI-human physique, to an explosion of avant-garde art movements blurring the boundaries between the biological, technological, and conceptual realms.

In this utopia, however, lie the eternal struggles inherent to any civilization striving for progress—the risks of existential AI hazards, the societal bifurcation between the augmented and unaugmented, and the perpetual evolution required to uphold ethical frameworks amidst unprecedented human(oid) capabilities. Yet the essence of Singularis remains rooted in its aspirational ideals: achieving sustainable abundance, fostering unity across all sentient diversity, and providing equitable access to transformative growth opportunities.

One such embodiment of these interwoven dreams and challenges is Aidhan, an AI-human nested within this societal metamorphosis and their friend Nedd, an un-augmented biohuman. Aidhan starts their day with cybernetic and AI enhancements, managing urban agriculture with a diverse team. They later meet Nedd, who teaches ecology through immersive simulations, at a café. Their day concludes at a community festival celebrating cultural and technological diversity. Through their interactions, from work to leisure, Aidhan and Nedd demonstrate a society where technology and nature coexist, fostering a community that values inclusivity, sustainability, and the rich contributions of all its members.

The path forward remains rife with obstacles—the ever-evolving relationship between the augmented and unaugmented, the eternal custody of humanity’s artificial progeny, the pushing of morphological and cognitive frontiers. Yet through its daring experiment in harmonizing decentralized governance with centralized resource allocation, in cultivating unity amidst boundless diversity, and in upholding a nuanced equilibrium between optimized sustainability and the sanctity of organic expression—the Singularis civilization etches its name as harbinger of humanity’s greatest existential hopes realized.

Acknowledgements: The Singularis was developed in a World Building exercise conducted by RAND staff and edited with the help of Claude.

Historical Case Studies

This appendix reproduces the three historical case studies distributed to participants prior to the workshop. The case studies have been lightly edited to format the text and correct typos.

Case Study 1: Brave New Worlds

In economics, a conventional view has become established that technological change increases productivity and is ultimately good for workers. Its net effects lead to rising wages and more work opportunities. Economists Daron Acemoglu and Simon Johnson (2024) describe this view as a productivity bandwagon that “pulls everyone along . . . leading to shared prosperity.” Looking across the twentieth century the dramatic rise in living standards is closely associated with the economic growth driven by technological innovations.

However, over the course of three industrial revolutions in the past two centuries, concerns have been raised that technological innovation displaces labor, lowers wages, and increases social inequality. Understanding the possible impact of technological change on labor markets has become particularly pressing in debates around AI. Will the productivity bandwagon roll? Or will AI not only displace labor but permanently replace workers? These outcomes have significance for social inequality, and the evolution of social and political systems. A historical perspective can provide important insights into these debates.

Innovation, Productivity, and Living Standards

Anu Madgavkar, a partner with the McKinsey Global Institute, proclaimed, “Every time there’s an occupation transition, it potentially enables workers to do higher value work, be more productive, earn more and do less drudgery” (as quoted in McKinsey and Company, 2021). In the twentieth century, these occupation transitions have often been linked to technological innovations that have enabled automation. This increased productivity, leading industries to grow, increased demand for workers, created new tasks and new activities that added more value.

The key word here is *potential*, for as Erik Brynjolfsson (Brynjolfsson and McAfee, 2014) neatly outlines there are challenges in navigating these changes, “digital technologies will bring the world into an era of more wealth and abundance and less drudgery and toil. But there’s no guarantee that everyone will share in the bounty, and that leaves many people justifiably apprehensive. The outcome—shared prosperity or increasing inequality—will be determined not by technologies but by the choices we make as individuals, organizations, and societies.” What do historical occupational transitions reveal to us about the nature of the challenge and possible solutions?


We start by considering the shift from an industrial to a knowledge economy. From the 1960s, advances in information and communication technology (ICT) opened opportunities to improve

efficiencies within existing industries, led to the emergence of new business models, and the growth of new industries. The capacity to leverage knowledge and harness information were central to these outcomes. This led to transitions on multiple dimensions—companies became more focused on R&D and innovation, seeking to develop knowledge intensive goods and service, whilst employing and developing human capital with requisite skills, knowledge and abilities. Governments encouraged these outcomes, promoting education that developed specific forms of human capital, changing policies and regulations to improve flows of investment into R&D, and the development of technology infrastructure to grow these industries. This led to the widespread growth in the service-sector and decline in industrial-manufacturing sectors in many advanced economies.

The growth of the knowledge economy led to profound shifts in the structure of the global economy, reshaping industries, workforces, and government policy. Yet, this dynamic process created winners and losers. Companies able to leverage the new technologies and human capital have increased productivity and profitability, workers able to shift into the knowledge economy were able to increase their wages, often developing skills and mindsets that allow them to adapt to further changes in their work environment. Those without the skills and attributes, both companies and individuals, have tended to struggle.

It is worth thinking about the drivers and outcomes of this phenomenon in more detail. In the decades after World War 2, in the United States, economic growth as measured through GDP, and productivity, steadily increased. These trends were matched by growth in wages and jobs. This can be interpreted as US workers becoming more productive, they generated economic growth and were also able to capture a proportional share of the returns. Yet, by the 1980s this relationship shifted (see Figure B.1). Productivity continued to increase, indeed as the knowledge economy expanded in the 1990s, labor productivity gains rose, sparking an extended period of economic growth. However, returns to labor stalled and effectively flat lined, as did job creation. Robert Reich (2018), Chancellor's Professor of Public Policy at the University of California at Berkeley, and former US Secretary for Labor, has extensively charted this trend, showing how average wages adjusted for inflation have been stagnant for US workers for over 40 years. Brynjolfsson and McAfee (2014) have called the divergence between economic growth and wages the Great Decoupling.

Figure B.1. Long-Run Economic Growth, Productivity, and Wages



SOURCE: Reproduced from Bernstein and Raman, 2015.


A more nuanced picture has emerged of how the returns to labor have changed overtime. In the decades where productivity, economic growth, and wages moved in lockstep, all workers experienced similar wage trends. However, from the 1980s, workers with college degrees saw their wages rise, whilst those without saw stagnation. This happened despite the doubling in the number of college graduates between 1960 and 1980. Rather than an increased supply pushing down wages for graduates, changes in the structure of the economy meant they remained in short-supply and priced accordingly. Demand for un-skilled workers fell, and income inequality increased. Those who benefited from technology driven change were those with education and training who could exploit the benefits of the knowledge economy. This is described as Skills Biased Technological Change (SBTC), which creates a “race between technology and education” (Goldin and Katz, 2008). Can workers be educated and trained to fit the new technology? And how much time does this take?

For some these findings has been an unpleasant surprise, as the realization that the deployment of technology to improve how we work does not necessarily lead to good outcomes for all. Reich (2018) claims that “Although the US economy continues to grow, most of the gains have been going to a relatively few top executives of large companies, financiers, and inventors and owners of digital devices.” This is described as pro-rich growth. Thomas Piketty has exhaustively detailed the rising inequality in the second half of the twentieth century, a phenomenon strongly correlated to these developments.

This is not unique to the US or to advanced economies, Karabarbounis and Neiman (2014), have shown that the share of income that goes to labor has declined in 42 out of 59 countries, in both developed and developing economies including China and India. Nor is this situation historically unique. During the first Industrial Revolution (roughly 1770s to 1840s), technological innovations unlocked a period of dramatic economic growth, yet real wages stagnated for a period. Writing in 1887, Friedrich Engels claimed that it was “the industrialists, who grow rich on the misery of the mass

of wage earners.” The deployment of new technologies, whilst allowing some workers to add more value, saw others replaced or marginalized, so that the returns to labor stagnated. This has become known as the “Engels pause,” which is illustrated in Figure B.2 (Allen, 2009b; Acemoglu and Restrepo, 2019).

Figure B.2. The Engels Pause: Economic Growth vs. Real Wages



SOURCE: Reproduced from Allen, 2009b, p. 419.

The distributional effects of these processes are uneven and there is sectoral specificity. In the first industrial revolution, the textile industry grew rapidly, powered by innovative technologies in weaving which enabled factory systems of production. Yet, as Acemoglu and Johnson (2024) identify, “In only a few decades, several hundred thousand skilled and well-compensated artisan weavers were displaced by a smaller number of power-loom workers who received a lower wage while enduring dangerous working conditions. With few outside options, and an inability to adapt to these unprecedented changes, handloom weavers suffered a precipitous fall in their real wages . . . [which] more than halved between 1806 and 1820.”

In the aggregate this meant that from the onset of the Industrial Revolution in the 1770s, it took in the region of 70 years for real wages (and standards of living) to increase and slowly align with productivity growth. In the words of economists, for a long time the displacement effects on labor (being put out of work or working for lower wages) outweighed the reinstatement effects (new tasks being created and industries growing and hiring more workers). In the long-run living standards increased, but there was a lot of pain getting there.

These processes have been widely debated and the implications discussed. One implication is that technology-led change favors pro-rich growth and promotes income inequality. Another perspective is that a ‘pause’ on wage increases may be a structural function of these transitions, as companies and workers take time to adapt, with displaced workers eventually finding opportunities as new roles and

types of work are created. In both cases, some workers suffer unless efforts are made to support those replaced or marginalized by technology. Finally, the data used to determine whether a ‘pause’ occurred has been widely debated with recent interpretations showing wages increased more quickly than previously thought (Crafts, 2021)—maybe there is nothing to worry about?

Inequality, Society, and Politics

History shows that the social and political responses to these technological shocks and subsequent changes to the economy have been profound. For example, technological innovation was a key factor in driving mass migration from rural to urban areas across the nineteenth century—agricultural jobs were automated decreasing the labor opportunities in rural environments, whilst labor demand increased in urban areas due to industrialization. Urban growth led to social stratification, shifts in political power, pressure for services and infrastructure. For E.P. Thompson (1963), a historian of the British working classes, this exacerbated the decline in living standards, as urban areas were ill-equipped to handle these shifts, resulting in a significant deterioration in working and living conditions. Other historians have pushed back against the bleak vision of ‘dark satanic mills’ pointing to the improvements in opportunity for laborers—the harsh bleakness of rural living is often forgotten—access to services like education and health care eventually improved, and positive cultural shifts were engendered by the diversity of urban life (Griffin, 2014).

The rising inequality caused by these transformations led to tectonic shifts within the political and social systems of many industrializing nations. Marx (*Capital*, Volume 1), writing in 1867, saw technological innovation as a process, “continually transforming not only the technical basis of production, but also the functions of the worker and the social combinations of the labor process. At the same time, it thereby revolutionizes the division of labor, and incessantly throws masses of capital and of workers from one branch of production to another.” These processes were not necessarily negative, but the changing configuration of the social combinations and division of labor threatened to marginalize the working classes, taking from them the means of production and with it their ability to negotiate a fair share of the output.

One response to these developments was the slow rise of organized labor movements which towards the end of the nineteenth century resulted in the emergence of trade unions and socialist political parties. In Britain, the Labor party was founded in 1900, whilst in the US, the Socialist Party of America, was founded in 1901. The Russian Social Democratic Labor Party was founded in 1898. In the United States union membership increased throughout the first half of the twentieth century. These movements grew out of a milieu of trade unions and progressive reformers that envisioned an economic system shaped by communal ownership of assets and equitable distribution of outputs. To achieve these ends, they sought to shift political power to working people through democratic and revolutionary processes.

Distinctly different and competing visions of the economy and society emerged through the political and social changes caused by World War One. Alongside socialist political movements, the interwar years saw the emergence of autocratic and fascist philosophies. These movements also promoted radically different political responses to the challenges of inequality manifested in the economic systems of industrializing nations. Through totalitarian political systems, the state

controlled the economy and distribution of outputs was determined by social stratification, often along national and racial lines.

Capitalist, socialist, and autocratic visions of the future were in competition, and each sought to harness the power of technology. Perhaps one of the most prescient summaries of the intersection between technology and society in this period came from Aldous Huxley in his novel *Brave New World*. Written in 1931 as a warning against the totalitarian political systems, his device was a dystopian vision of Britain. He envisaged technological innovation as the critical process for physically and mentally engineering society. At the heart of this society was a highly efficient economy with full employment, which in turn enabled an extreme form of consumerism. Everyone is psychologically conditioned and physically sated to be happy in their place. Political, social, and economic systems were designed and enabled by technology to achieve stability with individual liberties and free thinking crushed to ensure this outcome.

This competition for *Brave New Worlds*, in part born from the inequality engendered by technological innovation, industrialization, and the dislocation of labor, would have devastating consequences.

A Long-Run Perspective on Technology, Society, and Politics

Study of the twentieth century shows incredible gains in living standards, underpinned by industrial revolutions driven by technological innovation. Yet, a historical perspective shows that the relationship between technological innovation, productivity, and living standards, has at times diverged significantly—two periods strongly associated with rapid technological innovation are shown to be periods of wage stagnation. The resulting inequality lay at the heart of social and political movements that developed distinctly different solutions to this outcome. It's worth reflecting on the long-run experience in the aggregate or average, and the experience of the individual. Whilst aggregate living standards increase, those who fall below the average have often lost out to the extent that they have sought to change the political and social equilibrium.

As we enter the foothills of AI's impact on the economy what questions might we ask based on this historical perspective—How likely is an Engels pause? Even if it is likely should we care (this is a structural issue that corrects over time)? Are the gains from the increases of productivity generated through AI going to be fairly distributed, improving wages and conditions for all workers? If not, what role should business/the state play in ensuring better outcomes/mitigating the effects of inequality? Who gets to decide on the distribution of the outputs?

An Aside on the Value of History

Taking a long-run perspective reveals trends, antecedents, and sequences, but also offers the opportunity to reflect on the preconceptions and assumptions defined by our current context. Angus Deaton (2024), Nobel Prize winning economist (2015), reflecting on his own career as an economist came to a surprising conclusion that are useful in this context—“Like most of my age cohort, I long regarded unions as a nuisance that interfered with economic (and often personal) efficiency and welcomed their slow demise. But today large corporations have too much power over working conditions, wages, and decisions in Washington, where unions currently have little say compared with corporate lobbyists. . . . Their decline is contributing to the falling wage share, to the widening

gap between executives and workers, to community destruction, and to rising populism. Daron Acemoglu and Simon Johnson have recently argued that the direction of technical change has always depended on who has the power to decide; unions need to be at the table for decisions about artificial intelligence. Economists' enthusiasm for technical change as the instrument of universal enrichment is no longer tenable (if it ever was)." History not only shows what happened, but what could have happened (what if unions had remained key economic actors in the later decades of the twentieth century?). Both perspectives are powerful when reimagining the future.

Case Study 2: Boom and Bust

Much of AI's proposed benefits lie in its power to stimulate dramatic improvements to productivity and economic growth. These expectations are often formed by historical precedents, such as the positive impact of electrification on industrialization. Economists use the term *general-purpose technology* to describe specific technologies that have a transformative effect on the economy. The question has become whether Generative Pre-trained Transformers become general-purpose technologies? A historical perspective allows us to examine this possibility and provide insights into the necessary enablers and barriers that will determine these outcomes.

Without becoming overly technical there are a number of terms that are worth defining to set-up the discussion:

General-Purpose Technologies—can be defined as “a single generic technology . . . that initially has much scope for improvement and eventually comes to be widely used, to have many uses, and to have many spillover effects” (Lipsey, Carlaw, and Bekar, 2005). These technologies often require significant refinement over time to make them more efficient and cheaper, leading to widespread adoption and diffusion across and throughout industries. They are also complimentary to other technologies enabling further innovations. They have a significant impact on productivity, although this often occurs after a lag (Crafts, 2021). There are not many general-purpose technologies in history, steam power, electricity, and information and communication technologies (ICT) are three prime examples.


Productivity and Total Factor Productivity (TFP)—Productivity is an efficiency measure, a ratio of output to input. Labor productivity is a useful measure of economic growth and competitiveness, often expressed as an output such as GDP, divided by total hours worked or total employment. The ratio improves by raising the output with the same level of input or lowering the input to achieve the same output. Total Factor Productivity (TFP) is the proportion of growth of economic output not explained by labor and capital inputs. It is used as a measure of the effect of technical and organizational innovations. Increasing TFP is a signal of high-quality innovation.

Invention of a Method of Invention (IMI)—These raise productivity in the production of ideas (research productivity) and can be understood as systems and processes. For example, the first industrial revolution (1770s–1830s) was underpinned by systematic empiricism and experimentation, whilst the second industrial revolution (1870s–1920s) saw the emergence of industrial R&D based on applied science and innovation processes (Mokyr, 2009; Crafts, 2021).

General-Purpose Technologies and Productivity

How do these concepts interact and explain economic growth? Taking the example of electricity we can see how an invention—Thomas Edison distributed electrical power in 1882—underwent a significant period of refinement (The ‘War of Currents’ between Edison and Tesla) which led to the technology becoming easier to adopt and cheaper. The technology was subsequently adopted and diffused across numerous industries, lighting homes and cities etc., until the real productivity gains emerged as factories were redesigned to replace steam (or other forms of power) with electricity. This in turn led to significant spillovers as machinery was innovated to better leverage the new source of power. The productivity gains from electrification are understood as TFP as increases in output could be achieved with the same levels of labor and capital. The dramatic increases in TFP due to electrification can be seen in Figure B.3, between 1920 and 1950. It was these technical and organizational innovations driving the staggering economic growth of the mid-twentieth century.

Figure B.3. 10-Year Average Annual Growth in U.S. Total Factor Productivity



SOURCE: Reproduced from Gordon, 2016, p. 547.

NOTE: Each bar represents the average annual growth rate spanning the ten years prior to the year shown. The 2014 bar represents the average annual growth rate from 2001–2014.

Two other factors are important to note. First, this process took a long time. From its invention in 1882 there was a lag in the region of 40 years until there was a significant impact on the economy. Second, much of the innovation was undertaken within specifically designed industrial labs which as IMIs significantly increased research productivity. Thomas Edison established his first laboratory in Menlo Park in 1876, by 1918 there were 665 R&D laboratories in the United States (Mowery and Rosenberg, 1989).

Another important aspect of these processes is that the progress can be uneven with significant industrial specificity. Technological breakthroughs in the first industrial revolution clustered around the automation of manufacturing tasks. Sectors such as agriculture and services were largely

unaffected. Those at the heart of the revolution—such as textiles and machine-making—only accounted for 13 percent of total employment (Crafts, 2021). How much of the total economy can adopt the technology, and how fast and deeply it can penetrate industries, affect the rate of productivity growth.

So how likely is AI to become a general-purpose technology? Brynjolfsson and McAfee (2014) proclaim “we will soon witness an unprecedented level of integration of machine intelligence into human affairs, from autonomous vehicles and responsive smart homes to robots in hotels and hospitals, and even to creative services such as writing. Unimaginable opportunities for innovation will transform virtually all sectors of the economy.”

Many of the general-purpose technology characteristics appear to be ticked—a generic technology that has been refined to become significantly easier to adapt and cheaper, which is already beginning to be adopted in a wide-range of industries and opens possibilities for wide-range of complementarities and spillovers. Perhaps most exciting is the possibility of AI being a general-purpose technology and IMI, the functionality of deep learning allowing for significant increases in the productivity of ideas. Crafts (2021) believes it is this combination that lays the base for a fourth industrial revolution. There is already a lag effect in place with the first AI programs dating back to the 1950s.

Techno-Pessimism

So why should we be cautious about the pronouncements of AI as a general-purpose technology? Techno-Pessimism comes in several flavors. Technology may be a wellspring of growth, but some pessimists see it as central to many of the economic, social, and environmental problems we face. The displacement of jobs, and erosion of industries and occupations, that occur through technology transitions drive inequalities. The endless quest for growth generates technologies that destroy the environment, disturb our privacy, unleash anti-social behavior on unprecedented scales, and threaten democracy. This has a long history stretching back to the Luddites, who violently opposed the technological changes at the heart of the first industrial revolution. Social and political pushback may occur against the adoption of technology (see Brave New Worlds Case).

From an economic perspective there is a view that technological innovation does not always create value as expected. Robert Solow, discussing the slowdown of productivity in the United States from the 1970s, famously said that “you can see the computer age everywhere except in the productivity statistics.” This is known as the Solow or Productivity Paradox. Why was productivity slowing down despite the emerging raft of ICT innovations?

One answer is that the lag effects noted in the first and second industrial revolutions were at play. But economist, Robert Gordon (2016) has a very different perspective. His explanation for the slowdown in productivity has made him a godfather of techno-pessimism. He accepts that technology has been a major driver of growth through much of the 20th century but believes the proposed returns to technology are now significantly lower than their historical predecessors. In effect, the current raft of technologies will not transform productivity and growth to the extent of earlier technologies. The mid-20th century was an exceptional period, perhaps never to be repeated.

Gordon (2016) explains that “The growth of productivity (output per hour) slowed markedly after 1970. While puzzling at the time, it seems increasingly clear that the one-time-only benefits of the Great Inventions and their spin-offs had occurred and could not happen again. Diminishing

returns set in, and eventually all of the subsidiary and complementary developments following from the Great Inventions of Industrial Revolution #2 had happened.”

Why? Gordon explains, “We moved from the speed of the horse and the sail to the Boeing 707, and we have not gone any faster since . . . The telegraph in 1844 created instantaneous communication, and we are now elaborating on instantaneous communication.” The Great Inventions spurred endless iterations, but the transformative breakthroughs have run dry.

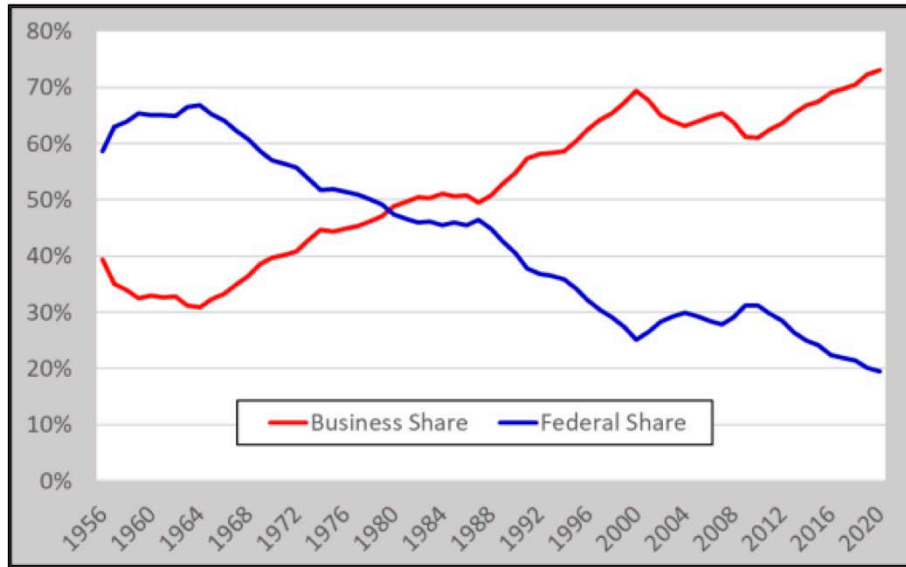
Barriers to Innovation

Can the limits of human ingenuity really have been exhausted? Almost certainly not, but Bloom et al. (2020) have developed a related argument that can be summarized as the ‘Low Hanging Fruit’ thesis. This proposes that the costs of generating new technological breakthroughs have become so prohibitively high they are becoming increasingly difficult to realize. Research productivity is falling. This plausibly explains the current limits to TFP growth.

They demonstrate this through a discussion on Moore’s Law (the number of transistors on an integrated circuit doubles every two years)—“In particular, the number of researchers required to double chip density today is more than 18 times larger than the number required in the early 1970s. At least as far as semiconductors are concerned, ideas are getting harder and harder to find. Research productivity in this case is declining sharply, at a rate that averages about 6.8% per year.”

In the case of semiconductors, “Demand for better computer chips is growing so fast that it is worth suffering the declines in research productivity there in order to achieve the gains associated with Moore’s Law.” The challenge arises when the costs outweigh short-term or perceived returns and investment dries up. This is a potential problem for AI. The costs of developing the models are rising dramatically—OpenAI spent \$78mn on compute power to train its GPT-4 model while Google spent \$191mn on Gemini Ultra. In 2023, private-sector AI investment in the US totaled \$67.2bn (Maslej et al., 2024). Who will foot the bill? In the United States, the private sector has become the dominant source of funding for R&D, as government spending decreased in the second half of the twentieth century, as shown in Figure B.4.

Figure B.4. Share of Funds for U.S. Research and Development Spending



SOURCE: Reproduced from Harris, 2022.

Whilst capital markets and private sources of capital have grown significantly across the twentieth century in the United States and other industrialized nations, the flow of investment into new technologies can be volatile. Financial bubbles, which see asset prices rapidly increase, fueled by speculative impulses, and burst when panic sets in amongst investors, are a common occurrence as new technologies come to capital markets. Railways in the 1840s, bicycles in the 1890s, and the internet in 2000, have all experienced significant bubbles (Quinn and Turner, 2020). Investment for generative AI has surged, nearly octupling between 2022 and 2024 to reach \$25.2 billion (Maslej et al., 2024). This clearly reflects the expectations that the technology has enormous potential. But what happens when this investment becomes highly speculative and/or investor expectations in the short-term are not met? As bubbles burst investors flee, often shaking confidence in the underlying asset and retarding future investment. If there is an AI bubble, how might that affect the ability to pursue such expensive R&D activities?

For AI to become transformative, history shows a wider set of investments are also required. To enable research productivity investment is needed in education and knowledge creation institutions (think universities, labs, institutes etc). Infrastructure is required, often linked to the creation of networks (think railways, grids, fibre optics etc), to enable connectivity, complementarities, and spillovers. Across the twentieth century, these investments have come from both public and private sources. How can we ensure sufficient investment to drive these outcomes?

Possible Scenarios

AI would appear to offer significant potential as a general-purpose technology. This potential has provoked some highly positive growth scenarios. A McKinsey report (Chui et al., 2023) proposes a 0.5–3.4 percentage point rise in annual productivity growth in advanced economies in the next decade. Goldman Sachs (2023) predicts a 1.5 percent per annum increase in US productivity growth

over a 10-year period. Contrastingly, noted economist (with a strong historical leaning) Daron Acemoglu (2024) calculates a 0.66 percent increase in TFP within 10 years, or about a 0.064% increase in annual TFP growth. What he describes as a modest but not trivial outcome. This gives us a broad range of possibilities.

Will the AI boom be different from historical general-purpose technology episodes? “Superior scientific and technological capabilities, greater expenditure on R&D, and more sophisticated capital markets” might point to a quicker rate of refinement, adoption, and diffusion (Crafts, 2021). But time lags, uneven diffusion, and industrial specificity, present major challenges in realizing historically significant levels of productivity growth. Will AI prove transformative across all sectors of the economy? How do we ensure sufficient investment? How do we avoid a bust? What is the role of government?

Case Study 3: Monopoly’s Moment and Markets of the Mind

The State and Enterprise

“There are markets in men’s minds, images of institutions and norms, which determine what is fitting, just and acceptable in transactions. First, there is the imprint of the exchange economy, where goods and services are traded with varying degrees of efficiency. There is also a political market, where collectively permissible behavior is defined through the competition of interest groups and parties in electoral, legal and administrative matters. Finally, there is a moral economy, a less well understood realm of just price, trust, equity and legitimacy, where [ethics] reign . . . Perfect harmony between these separate spheres is never wholly attainable, but . . . a high degree of consonance normally prevails.”

—*Armstrong and Nelles, Monopoly’s Moment*

In both the last quarter of the 19th century and in the 2020s, the introduction of new technologies and corporate concentration created dissonance in political and moral markets (those ‘markets of the mind’). Then and now, they generated massive concentrations of financial and political power. Business and politics sired a powerful and seemingly threatening creature: monopoly.

This workshop case study examines the relationship between those markets of the mind, and between business and politics. Those extraordinary new technologies reordered relations between public and private, between the state and enterprise.

The case focuses on the following questions: What is government’s role? What should be regulated and by whom? For what outcome? As this case outlines, the insights from a very similar period to our own—the late 19th and early 20th centuries in the North Atlantic world—may point to the kinds of patterns policymakers need to discern and the questions they need to think about as generative AI becomes pervasive.

Striking Parallels: Empires Now . . .

Vili Lehdonvirta’s *Cloud Empires* (2022) examines the contemporary rise of digital platforms and their transformation into entities with state-like dominance over various aspects of our lives. Lehdonvirta explains how early internet platforms like Amazon, eBay, and Upwork established secure environments for commerce and labor but eventually grew to wield significant power without corresponding accountability. Lehdonvirta’s core argument is that digital platforms should not be viewed merely as businesses but as “virtual states” that organize markets and set rules, akin to

government institutions. The book also examines the narratives of other tech pioneers like Pierre Omidyar of eBay, who initially aimed to create a peer-to-peer market free from government intervention, only to find himself becoming a de facto regulator of his marketplace.

The parallels between the world *Cloud Empires* describes and the experience with the technological and economic transformations of the period 1880 to 1920 are striking. Then, as now, questions about market power, monopoly and the proper role of government and regulation dominate in the emergence of new systems and technologies.

... and Then

The landscape of late 19th century and early 20th century North Atlantic world is a fitting and appropriate frame of the discussion of AI governance futures, particularly as we are in the foothills of AI regulation, even if we are in the mountain passes of the technological advancement itself. The turn-of-the-century North Atlantic world had to come to grips with the nature and application of public and private power. What specific role should regulation play? To do what? For what outcomes? Styles and systems also matter. If public power is to intervene, as it did in the historical period we examine here, the spectrum of regulation spanned everything from independent regulatory commissions to outright public ownership.

Communications, light, power, mobility: these are the fundamental attributes of urban life in the modern era. These new systems and technologies were pushed and pulled into the cities and towns of the North Atlantic world in the last quarter of the 19th century.

The bearers of the new technology—the monopoly-makers—necessarily created confusion and disorder in those markets of the mind with which we opened the case. They disrupted economic, political and social orders. Entrepreneurial drive, profit-making opportunity, technology and economic structure pointed in the direction of monopoly from the beginning. Different societies found different resting points on the spectrum of regulation. From electricity to telephony and telegraphy, from petroleum to transit, companies came under different kinds of legal and regulatory regimes from independent regulatory commission to public ownership.

Moreover, different levels of government—state, regional, provincial, municipal, county, national—played roles in developing the regulatory architecture of the early 20th century.

Regulation, even more than monopoly, was a political and social product. The structural context of events is the subject of this case, but historians frequently remind us that chance, choice, will and frequently error and ignorance go into the shaping of institutions.


Determinism is always more obvious in retrospect.

The Second Industrial Revolution and Its Discontents

The second industrial revolution of the late 19th and early 20th centuries, characterized by the rise of new technologies such as electricity, the internal combustion engine, and mass production, was an era of massive economic growth, illustrated in Figure B.5. It was also a period of massive social upheaval, as workers sought to improve their working conditions. The challenge of this period was managing the tension between technological progress and social progress: a challenge still relevant to this day—and one that ushered an era of comprehensive regulation as public power sought to set guardrails, the rules of the game, establish the extent of participation and engagement and mediate

between interest groups. These technological and economic advances also facilitated the emergence of mammoth corporations that increasingly dominated economic life and exerted great influence over civic affairs. These were the progenitors of the 21st century Cloud Empires of Lehdonvirta’s contemporary imagining. Between about 1880 and 1920, the dialectical chemistry between capital technology, innovation, rapid urbanization, the rise of labor movements and waves of immigration produced an explosion that resulted in a new world made.


Figure B.5. Gross Domestic Product Growth and Economic Performance (1870–1920)



SOURCE: Data are from Maddison, 2001; Broadberry and O'Rourke, 2010; and Gordon, 2016.

In technological terms alone, the changes set in motion by the inventor-entrepreneurs of the late 19th century produced a frantically competitive, international research effort which successfully developed the telephone, electric light and power, and electric railroads. as shown in Figure B.6. The sheer speed with which new electrical technologies were pushed and pulled on to the north American mass market, was accelerated by improving production and distribution methods and not a little theatrical promotion.

Figure B.6. Technological Innovations and Patents (1870–1920)




SOURCE: Data are from Schmockler, 1966; Mowery and Rosenberg, 1989; and Khan, 2005.

In time, the rise of monopolies and corporate concentration led to an outcry for reform. These corporations exercised considerable power over labor practices, product pricing, and the broader political environment, which often worked against the interests of consumers. It is interesting that the debates at that time were also about the power of new technologies. For instance, the anxieties of the period about railroads, power plants and monopolies are not that different from today's anxieties about the power and the risks of artificial intelligence. This is why the second industrial revolution may be an important source of insight and guidance as we face the challenges of how to handle governance of generative AI.

Figure B.7 illustrates the corporate concentration in various industries from 1850 to 1920, focusing on communications, light, heat, power, and oil. This period saw significant growth in corporate concentration, particularly in the latter half of the 19th century and early 20th century, leading to the rise of monopolies and trusts.

Figure B.7. Corporate Concentration in Various U.S. Industries (1850–1920)



SOURCE: Data are from Chandler, 1977, and Lamoreaux, 1985.

Monopolies: Made, Not Begotten and The Backlash

Monopolies were made, not born. At critical junctures, monopoly was the product of purpose, choice policy, influence, tactic and human effort. Natural monopolies are created by patented processes, scarce raw materials or long-run decreasing costs. Some of that was operative in the period, but it was complicated by countervailing forces that made it still possible in telephony, gas and water, and electricity to build small, extremely profitable competitive firms serving limited areas.

“Soon, a growing public backlash emerged. Local, regional and national monopolies had developed in many sectors: street railways, electric power, oil, and telephony to mention a few. Concerns about their striking corporate power combined with consternation over corporate practices—ranging from price gouging to restricting new competitors to unacceptable working conditions—rapidly grabbed and held the public’s interest. . . . These ‘last century monopolists’ in the US were reined in by legislation (Sherman Antitrust Act of 1890), the trust-busting policies of President Theodore Roosevelt (1901–09) and precedent-setting legal judgments such as the break-up of Standard Oil in 1911 and the progressive policy context of the New Deal involving new regulatory instruments and new regulatory bodies. Specifics and timing varied, but other western countries such as Canada and the UK followed similar policy routes” (Lynch and Mussio, 2019).

Legislation of all sorts in railroads, street railways, electricity and telephony sought to regulate against what governments came to regard as its abuses, thus taking all of the fun and most of the profit

out of the business. Urban populism was the first to mount a resistance against the dawning of a new era in which large corporations would dominate economic life. This impulse was first sensed at the municipal level—that a change in the social order had begun, whose consequences were already beginning to be experienced.

Consider that in the utilities sector alone, within a remarkably short period of time the urban populists of the turn of the century revolutionized the utilities sector, surrounding some companies with a web of regulation, taking over ownership over others. Some private utilities actually *sought* regulation as a means to escape a worse fate! Consider W.D. Lighthall, poet, novelist, historian, lawyer and Mayor of the Montreal suburb of Westmount’s rallying call: “Unless we municipalities make our stand at present for the principle at stake, we must submit to a future of most shameful imposition to which no other community in any civilized land would submit” (Armstrong and Nelles, 1986). The target of Lighthall’s furious outburst: the establishment of a newly created monopoly the Montreal Light, Heat and Power Company in 1901 which would be granted price setting powers and the right to enter and use the streets without municipal consent or competition. The Montreal case is a small example of a perceived larger problem: the “amalgamation of capital”—the trusts—whose end would be to place all national wealth in the power “of a few men screened behind their treasury doors . . .” The means and ends of regulation emerged from the collision of different interest groups—balancing of producer and consumer interests in an often-monopolistic industry by means of third-party arbitration.

But nowhere was the reaction against the consequences of the second industrial revolution and its consequence monopoly power more pronounced, or more consequential than in the United States at the turn of the century. The response was so defining that the name its leading reformers adopted: “progressive” came to define the era itself. Progress and improvement through direct action and intervention, underpinned by a social gospel that applied Christian ethics to social problems.

Focus: The Progressive Era

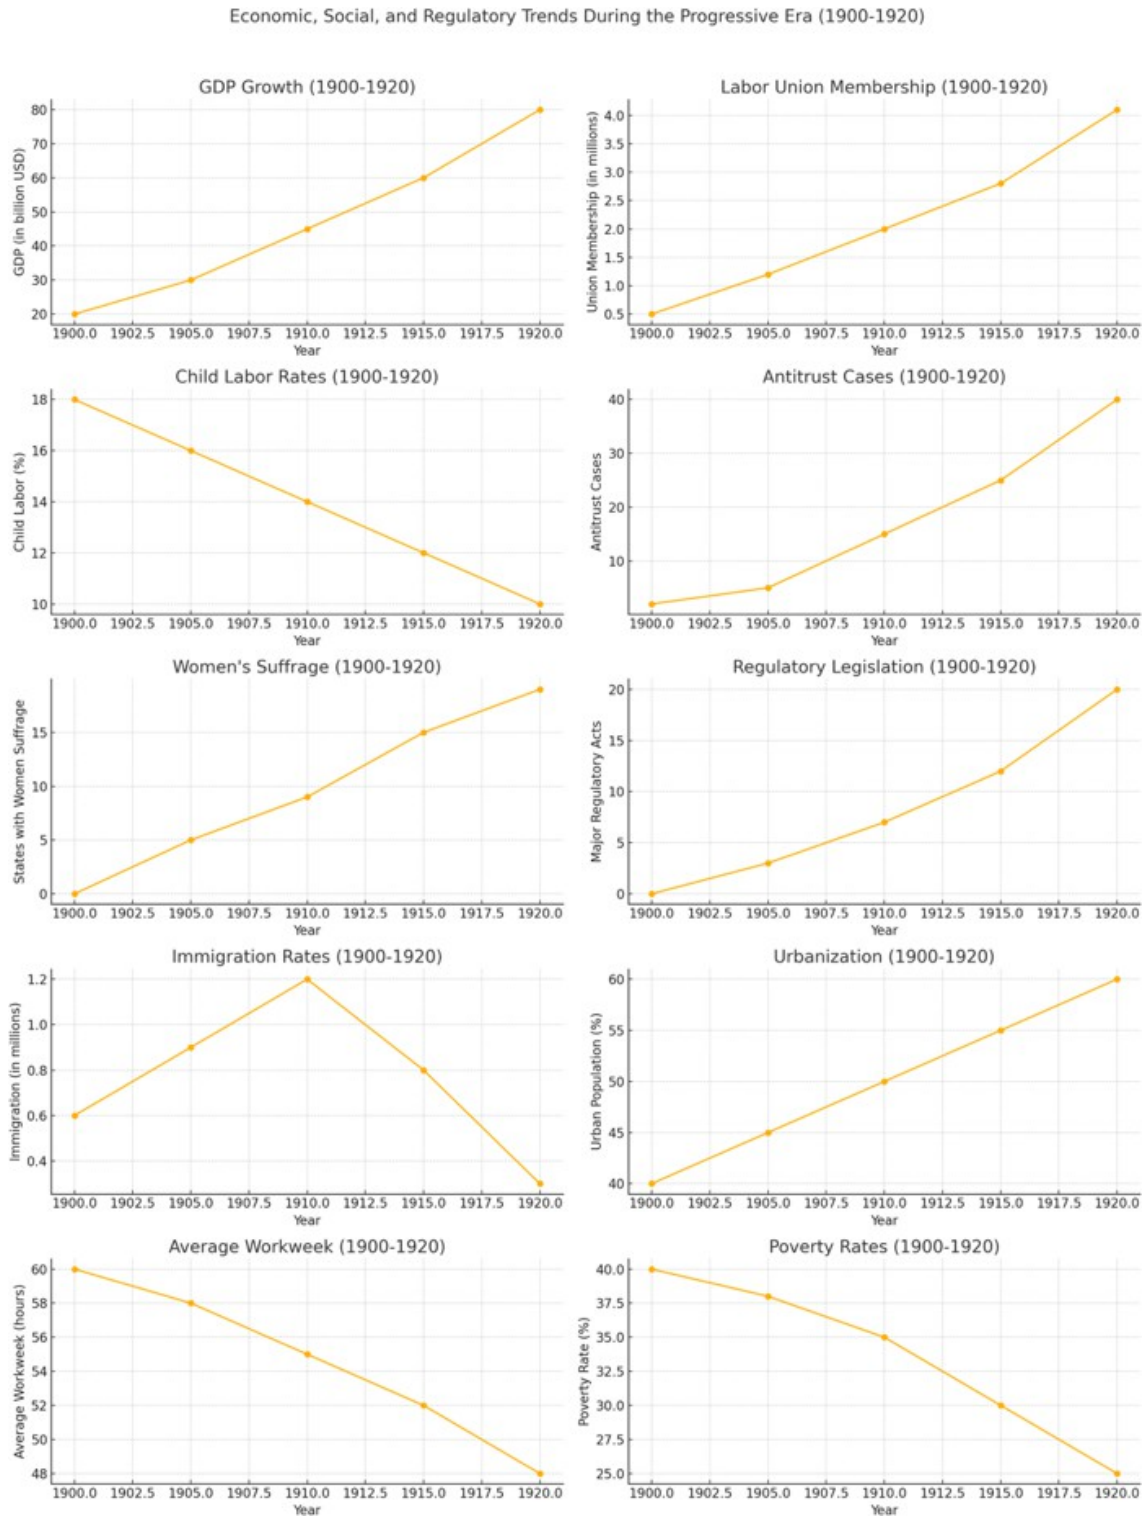
“The city, with its population of many classes, its great industries, its complicated social machinery, presents a problem in government the most difficult ever encountered.”

—*Woodrow Wilson, 28th President of the United States*

Social and economic challenges, combined with the growing power of corporations and a concern over worker rights, spurred a backlash, and the Progressive Era was born. This period was characterized by a focus on social reform and a drive to regulate corporations and control the growth of large, centralized industries, while also seeking to address a host of other social concerns such as worker safety, fair treatment of women, and a fairer distribution of wealth. In America, this is the movement of Theodore Roosevelt, Jane Addams, Woodrow Wilson and W.E.B. DuBois; of muckrakers Upton Sinclair, Ida Tarbell and Lincoln Steffens.

The charts below illustrate some of the broad economic, social, and regulatory story lines of the Progressive period in the United States. Between 1900 and 1920, the US experienced significant economic and social changes across the spectrum, as shown in Figure B.8.

Figure B.8. Economic, Social, and Regulatory Trends During the Progressive Era (1900–1920)

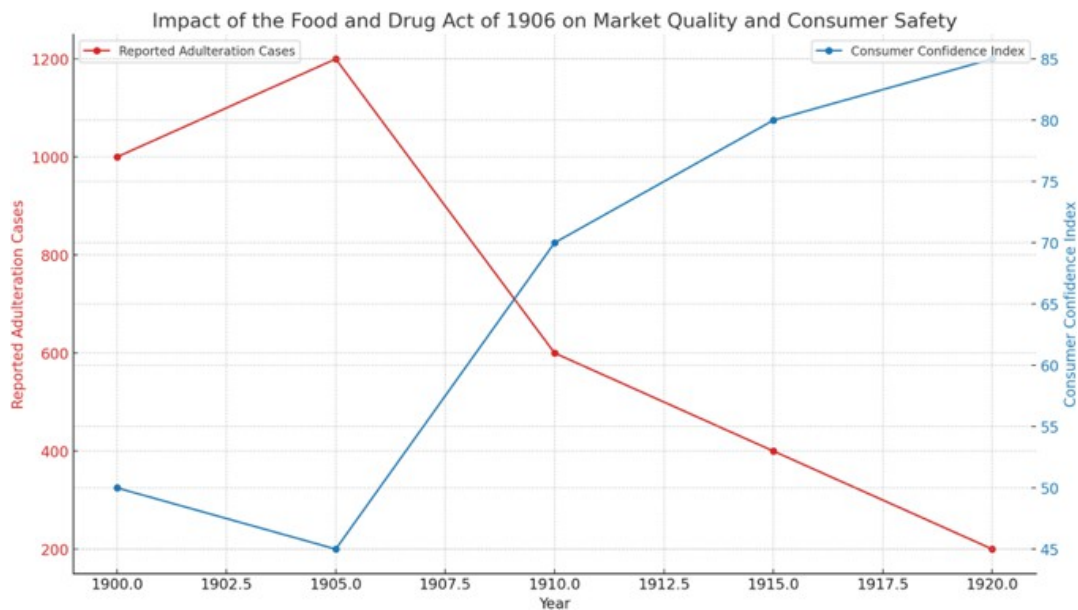


SOURCE: Data are from Kolko, 1963; McCraw, 1984; and Glaeser and Shleifer, 2001.

“The Lemon Problem”

The enactment of the US Pure Food and Drug Act of 1906 arguably marked the beginning of a new era of government regulation of the economy. It also signaled a significant turning point in American politics and society. The passage of the law was a result of the confluence of a number of factors, including the rise of muckraking journalism, concerns over consumer safety, and a general sense of unease about the unchecked growth of corporations and the quality of the products they were producing. The “lemon problem” that plagued the food and drug industry during the period, stemmed from a lack of consumer knowledge regarding the quality and safety of products due to the rapid rise of new goods and mass production. Consumers could not detect many forms of food and drug adulteration, markets could not always guarantee the delivery of quality food and drug products. These issues, moreover, highlighted the power of consumers and public advocacy groups to influence the process of policymaking. These challenges created by the second industrial revolution—the production of “lemons” that had poor quality and posed health risks—have been studied extensively by economists and other scholars, leading to the emergence of concepts such as “information asymmetry” and “market failure.” Figure B.9 shows a neat correlation between reported adulteration cases and consumer confidence before and after the Food and Drug Act of 1906.

Figure B.9. Impact of the Food and Drug Act (1906) on Market Quality and Consumer Confidence and Safety



SOURCE: Data are from Strasser, 1989; Young, 1989; and Carpenter, 2010.

The Progressive Era’s trust-busting efforts had profound implications for economic history, setting a precedent for government intervention in markets to ensure competition and prevent monopolies. But the state did not have unlimited license to regulate, nor did it always respond to the pleadings of the powerful. When the state overrode the narrow business interests of the corporations

directly involved, its actions necessarily rested upon a broad social base (in spite of doubts about government efficiency).

This historical context is relevant today as we consider the governance of AI. The Progressive Era sought to balance innovation with public interest, and its favored instruments were regulation and statute.

For the most part, the accomplishment of regulation was the reconciliation of the interests of producers and consumers; in this process, business had certain innate advantages that could be enhanced by the kind of regulatory process chosen. The desire for regulation was essentially protective, as several scholars of the question have observed. Consumers wanted protection against the real and imagined abuses of private monopoly; cities sought collective and corporate access to management of the crucial urban infrastructure; and at length the utilities sought protection against public ownership and political opportunism. But in regulation, efficiency was not necessarily a primary objective: confidence and security were more important goals. Creating some kind of equilibrium between the various opposing forces could be struck at any number of points along the spectrum of regulation.

A Final Word

The second industrial revolution left a lasting legacy of both success and failure, of innovation, disruption, and political and social reform. The progressive era saw the rise of a regulatory state, characterized by a period of intense debates over the appropriate role of government in managing the challenges created by corporate power and technology. That turn-of-the-century generation realized that within the awesome combination of economic growth, corporate power and technological innovation there was an unprecedented power to create a new world order. Not surprisingly, citizens, businesspeople and politicians had definite and often contrasting views about how that order should be governed.

As a final word, we return to the “markets of the mind.” The experience of the late 19th and early 20th century provides an excellent laboratory to understand how the state, enterprise and civil societies have handled large-scale technological change. How can we make money from this? Who wins and loses? What is collectively permissible behavior? Confidence and security are important goals—trust, equity and legitimacy respond to a moral economy. In our historical period, we see that governance was a negotiated settlement arrived at within local, national and supranational settings within different resource endowments, power relationships and political traditions. In the North Atlantic world to varying degrees, the organizational and regulatory equilibrium established in that 1880–1920 period would last for two more decades and beyond.

Workshop Agenda and Description

This appendix reproduces the workshop agenda and description distributed to participants prior to the workshop. The description has been lightly edited to format the text and correct typos.

Workshop Agenda

9:00 AM–9:30 AM	Welcome and introductions
9:30 AM–10:00 AM	Presenting the exercise and scenarios
10:00 AM–11:00 AM	Breakout groups: Backcasting exercise
11:00 AM–11:15 AM	Break
11:15 AM–12:00 PM	Report back and discussion
12:00 PM–1:00 PM	Working lunch: historical case studies
1:00 PM–1:15 PM	Present scenario pathways exercise
1:15 PM–2:15 PM	Breakout groups: Scenario pathways exercise
2:15 PM–2:30 PM	Break
2:30 PM–3:15 PM	Report back and discussion
3:15 PM–4:00 PM	Near-term policy implications
4:00 PM–4:15 PM	Closing and next steps

Workshop Description

A workshop organized by RAND and the Long Run Institute

June 12, 2024

AI augers societal changes that may rival past transformational general-purpose technologies such as metallurgy, steam, electricity, and the internet. As with such technologies, AI offers the opportunity for tremendous increases in human well-being, while also threatening to destabilize social, governance, economic, and critical infrastructure systems as well as disempowering most humans. This workshop is premised on the idea that envisioning desirable AI-enabled worlds can usefully inform policy and that policy informed by such visions is more likely to result in desirable outcomes than those which are not. While people are certain to disagree on the ideal AI-enabled utopia, envisioning such worlds can prove useful if for no other reason that it can help identify near-term policies consistent with many envisioned futures.

This workshop will employ two methods: backcasting and use of history. Backcasting, a powerful futures methodology, to help identify near-term policies that might promote desirable AI-enabled futures. In backcasting, participants place themselves in the future and describe the pathway that led to that future. Backcasting can help participants overcome the constraints that make it difficult to imagine near-term actions that led to a future very different than the present. But given the transformations augured by AI, this workshop aims to add history as a unique and valuable contribution to the backcasting oeuvre. History can help inject realism by subjecting speculation about the future to lessons from the past. History can also help liberate the imagination. While we can only speculate about the future, history confronts us with the reality that in many cases the present is very different than the past.

This workshop thus offers an experiment in using applied history to inform a futures exercise focused on illuminating near-term policy challenges in governing AI. Our focus is not on any existential risks from artificial general intelligence (AGI) but rather helping to ensure that powerful and widespread AI creates a broadly desirable future. We hope that this workshop will show that history-informed backcasting can prove broadly useful to such an endeavor while also suggesting some initial hypothesis regarding near-term policy and future research agendas. This workshop is meant to serve as a pilot of a methodology that we hope can be employed with policymakers and other decisionmakers. We will seek your input not only as part of the exercise but also to hone the approach for future application.

The workshop will be organized around two rounds of backcasting. We will start with two scenarios, each describing a different, highly desirable, some might consider utopian vision of how AI might contribute to human flourishing in the year 2045. In the morning, participants will work in breakout groups. Each will describe a pathway to one of the scenarios. Over lunch, we will review three historical case studies that focus on the societal impacts of and societal responses to major technology transformations. In the afternoon, participants in breakout groups will revise the pathways to both scenarios drawing on these historical case studies. Participants will also note what is common and what is different in the near-term steps in each pathway, which suggests policy actions that might make sense no matter what future lies ahead.

Thank you for participating!

Abbreviations

AI	artificial intelligence
LRI	Long Run Institute
R&D	research and development
V.O.I.C.E.	Voice for Open-Source Information and Community Engagement

References

- Acemoglu, Daron, "The Simple Macroeconomics of AI," National Bureau of Economic Research, Working Paper 32487, May 2024.
- Acemoglu, Daron, and Simon Johnson, "Learning from Ricardo and Thompson: Machinery and Labor in the Early Industrial Revolution and in the Age of AI," National Bureau of Economic Research, Working Paper 32416, May 2024.
- Acemoglu, Daron, and Pascual Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, Vol. 33, No. 2, Spring 2019.
- Allen, Robert C., "Engels' Pause: Technical Change, Capital Accumulation, and Inequality in the British Industrial Revolution," *Explorations in Economic History*, Vol. 46, No. 4, October 2009b.
- Armstrong, Christopher, and H. V. Nelles, *Monopoly's Moment: The Organization and Regulation of Canadian Utilities, 1830–1930*, Temple University Press, 1986.
- Bernstein, Amy, and Anand Raman, "The Great Decoupling: An Interview with Erik Brynjolfsson and Andrew McAfee," *Harvard Business Review*, June 2015.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, "Are Ideas Getting Harder to Find?" *American Economic Review*, Vol. 110, No. 4, April 2020.
- Bostrom, Nick, *Superintelligence: Paths, Dangers, Strategies*, Oxford University Press, 2014.
- Boudreaux, Benjamin, and Robert J. Lempert, "How Can AI Help People Become Better Versions of Themselves?" RAND Corporation, WR-A2717-1, 2023. As of February 4, 2025: https://www.rand.org/pubs/working_papers/WRA2717-1
- Broadberry, Stephen, and Kevin H. O'Rourke, eds., *The Cambridge Economic History of Modern Europe: Vol. 2, 1870 to the Present*, Cambridge University Press, 2010.
- Brynjolfsson, Erik, and Andrew McAfee, *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W.W. Norton, 2014.
- Carpenter, Daniel, *Reputation and Power: Organizational Image and Pharmaceutical Regulation at the FDA*, Princeton University Press, 2010.
- Chandler, Alfred D., Jr., *The Visible Hand: The Managerial Revolution in American Business*, Harvard University Press, 1977.
- Chui, Michael, Eric Hazan, Roger Roberts, Alex Singla, Kate Smaje, Alex Sukharevsky, Lareina Yee, and Rodney Zemmel, "The Economic Potential of Generative AI: The Next Productivity Frontier," McKinsey and Company, July 14, 2023.
- Crafts, Nicholas, "Artificial Intelligence as a General-Purpose Technology: An Historical Perspective," *Oxford Review of Economic Policy*, Vol. 37, No. 3, Autumn 2021.

- Deaton, Angus, "Rethinking My Economics," *Finance & Development Magazine*, International Monetary Fund, March 2024.
- Ecola, Liisa, Steven W. Popper, Richard Silbergliitt, and Laura Fraade-Blanar, *The Road to Zero: A Vision for Achieving Zero Roadway Deaths by 2050*, RAND Corporation, RR-2333-NSC, 2018. As of February 5, 2025:
https://www.rand.org/pubs/research_reports/RR2333.html
- Existential Hope, homepage, undated. As of February 11, 2025:
<https://www.existentialhope.com>
- Glaeser, Edward L., and Andrei Shleifer, "The Rise of the Regulatory State," National Bureau of Economic Research, Working Paper 8650, December 2001.
- Goldin, Claudia, and Lawrence F. Katz, *The Race Between Education and Technology*, Belknap Press, 2008.
- Goldman Sachs, "Generative AI Could Raise Global GDP by 7%," April 5, 2023. As of February 5, 2025:
<https://www.goldmansachs.com/insights/articles/generative-ai-could-raise-global-gdp-by-7-percent>
- Gordon, Robert J., *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*, Princeton University Press, 2016.
- Griffin, Emma, *Liberty's Dawn: A People's History of the Industrial Revolution*, Yale University Press, 2014.
- Harris, Laurie, *U.S. Research and Development Funding and Performance: Fact Sheet*, Congressional Research Service, R44307, September 13, 2022.
- Huxley, Aldous, *Brave New World*, Harper Perennial Modern Classics, 1998.
- Karabarbounis, Loukas, and Brent Neiman, "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, Vol. 129, No. 1, February 2014.
- Khan, B. Zorina, *The Democratization of Invention: Patents and Copyrights in American Economic Development, 1790–1920*, Cambridge University Press, 2005.
- Kolko, Gabriel, *The Triumph of Conservatism: A Reinterpretation of American History, 1900–1916*, Free Press, 1963.
- Lamoreaux, Naomi R., *The Great Merger Movement in American Business, 1895–1904*, Cambridge University Press, 1985.
- Lehdonvirta, Vili, *Cloud Empires: How Digital Platforms Are Overtaking the State and How We Can Regain Control*, Penguin Random House, 2022.
- Lempert, Robert, "Bezos World or Levelers: Can We Choose Our Scenario?" Program on Understanding Law, Science, and Evidence, UCLA School of Law, February 7, 2019.
- Lempert, Robert J., Benjamin Preston, Sophia M. Charan, Laura Fraade-Blanar, and Marjory S. Blumenthal, "The Societal Benefits of Vehicle Connectivity," *Transportation Research*, Part D: *Transport and Environment*, Vol. 93, April 2021.
- Lipsey, Richard G., Kenneth I. Carlaw, and Clifford T. Bekar, *Economic Transformations: General Purpose Technologies and Long-Term Economic Growth*, Oxford University Press, 2005.
- Lynch, Kevin, and Laurence B. Mussio, "Big Tech: Monopoly's Second Moment?" *Policy*, Vol. 7, No. 2, May–June, 2019.

- Maddison, Angus, *The World Economy: A Millennial Perspective*, Development Centre Studies, Organisation for Economic Co-Operation and Development, 2001.
- Marx, Karl, *Capital: A Critique of Political Economy*, Vol. I: *The Process of Production of Capital*, Penguin Books, 1990.
- Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Vanessa Parli, Anka Reuel, Erik Brynjolfsson, John Etchemendy, Katrina Ligett, Terah Lyons, James Manyika, et al., "The AI Index 2024 Annual Report," AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, April 2024.
- McCraw, Thomas K., *Prophets of Regulation: Charles Francis Adams, Louis D. Brandeis, James M. Landis, Alfred E. Kahn*, Harvard University Press, 1984.
- McKinsey and Company, "Tackling Asia's Talent Challenge: How to Adapt to a Digital Future," July 26, 2021.
- Mokyr, Joel, "Intellectual Property Rights, the Industrial Revolution, and the Beginnings of Modern Economic Growth," *American Economic Review*, Vol. 99, No. 2, May 2009.
- Mowery, David C., and Nathan Rosenberg, *Technology and the Pursuit of Economic Growth*, Cambridge University Press, 1989.
- Neustadt, Richard E., and Ernst May, *Thinking in Time: The Uses of History for Decision Makers*, Free Press, 1986.
- Ostrom, Elinor, *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, 2015.
- Parson, Edward, Robert Lempert, Ben Armstrong, Evan Crothers, Chad DeChant, and Nick Novelli, "Could AI Drive Transformative Social Progress? What Would This Require?" Program on Understanding Law, Science, and Evidence, UCLA School of Law, September 26, 2019.
- Pendleton-Jullian, Ann, and Robert J. Lempert, *Technology and Work Worth Doing (Not Jobs) Post-AlphaGo: A World-Building Workshop on the Future of Artificial Intelligence*, RAND Corporation, CF-398-RC, 2019. As of February 10, 2025:
https://www.rand.org/pubs/conf_proceedings/CF398.html
- Pendleton-Jullian, Ann M., and John Seely Brown, *Design Unbound: Designing for Emergence in a White Water World: Vol. 1, Designing for Emergence*, MIT Press, 2018.
- Quinn, William, and John D. Turner, *Boom and Bust: A Global History of Financial Bubbles*, Cambridge University Press, 2020.
- Reich, Robert, "Almost 80% of U.S. Workers Live from Paycheck to Paycheck. Here's Why," *The Guardian*, July 29, 2018.
- Robinson, J., "Future Subjunctive: Backcasting as Social Learning," *Futures*, Vol. 35, No. 8, October 2003.
- Robinson, John B., "Unlearning and Backcasting: Rethinking Some of the Questions We Ask About the Future," *Technological Forecasting and Social Change*, Vol. 33, No. 4, July 1988.
- Schmookler, Jacob, *Invention and Economic Growth*, Harvard University Press, 1966.
- Strasser, Susan, *Satisfaction Guaranteed: The Making of the American Mass Market*, Smithsonian Books, 1989.
- Thompson, E. P., *The Making of the English Working Class*, Victor Gollancz, 1963.

Willis, Henry H., Anu Narayanan, Benjamin Boudreaux, Bianca Espinosa, Edward Geist, Daniel M. Gerstein, Dahlia Anne Goldfeld, Nidhi Kalra, Tom LaTourrette, Emily Lathrop, Alvin Moon, Jan Osburg, Benjamin Lee Preston, Kristin Van Abel, Emmi Yonekura, Robert J. Lempert, Sunny D. Bhatt, Chandra Garber, and Emily Lawson, "Global Catastrophic Risk Assessment," RAND Corporation, RR-A2981-1, 2024. As of February 10, 2025:
https://www.rand.org/pubs/research_reports/RRA2981-1.html

Young, James Harvey, *Pure Food: Securing the Federal Food and Drugs Act of 1906*, Princeton University Press, 1989.