



National
Bureau of
Economic
Research

THE ECONOMICS OF ARTIFICIAL INTELLIGENCE

An Agenda

Edited by Ajay Agrawal,
Joshua Gans, and Avi Goldfarb



Artificial Intelligence: An Agenda

Edited by **Ajay Agrawal, Joshua Gans,
and Avi Goldfarb**

The University of Chicago Press

Chicago and London

The University of Chicago Press, Chicago 60637
The University of Chicago Press, Ltd., London
© 2019 by the National Bureau of Economic Research, Inc.
All rights reserved. No part of this book may be used or reproduced
in any manner whatsoever without written permission, except in the
case of brief quotations in critical articles and reviews. For more
information, contact the University of Chicago Press, 1427 E. 60th St.,
Chicago, IL 60637.
Published 2019
Printed in the United States of America

28 27 26 25 24 23 22 21 20 19 1 2 3 4 5

ISBN-13: 978-0-226-61333-8 (cloth)
ISBN-13: 978-0-226-61347-5 (e-book)
DOI: <https://doi.org/10.7208/chicago/9780226613475.001.0001>

Library of Congress Cataloging-in-Publication Data

Names: Agrawal, Ajay, editor. | Gans, Joshua, 1968– editor. | Goldfarb, Avi, editor.

Title: The economics of artificial intelligence : an agenda / Ajay Agrawal, Joshua Gans, and Avi Goldfarb, editors.

Other titles: National Bureau of Economic Research conference report.

Description: Chicago ; London : The University of Chicago Press, 2019. | Series: National Bureau of Economic Research conference report | Includes bibliographical references and index.

Identifiers: LCCN 2018037552 | ISBN 9780226613338 (cloth : alk. paper) | ISBN 9780226613475 (ebook)

Subjects: LCSH: Artificial intelligence—Economic aspects.

Classification: LCC TA347.A78 E365 2019 | DDC 338.4/70063—dc23

LC record available at <https://lcn.loc.gov/2018037552>

© This paper meets the requirements of ANSI/NISO Z39.48-1992 (Permanence of Paper).

Contents

	Acknowledgments	xi
	Introduction	1
	Ajay Agrawal, Joshua Gans, and Avi Goldfarb	
I.	AI AS A GPT	
	1. Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics	23
	Erik Brynjolfsson, Daniel Rock, and Chad Syverson	
	<i>Comment:</i> Rebecca Henderson	
	2. The Technological Elements of Artificial Intelligence	61
	Matt Taddy	
	3. Prediction, Judgment, and Complexity: A Theory of Decision-Making and Artificial Intelligence	89
	Ajay Agrawal, Joshua Gans, and Avi Goldfarb	
	<i>Comment:</i> Andrea Prat	
	4. The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis	115
	Iain M. Cockburn, Rebecca Henderson, and Scott Stern	
	<i>Comment:</i> Matthew Mitchell	

	5. Finding Needles in Haystacks: Artificial Intelligence and Recombinant Growth	149
	Ajay Agrawal, John McHale, and Alexander Oettl	
	6. Artificial Intelligence as the Next GPT: A Political-Economy Perspective	175
	Manuel Trajtenberg	
II.	GROWTH, JOBS, AND INEQUALITY	
	7. Artificial Intelligence, Income, Employment, and Meaning	189
	Betsy Stevenson	
	8. Artificial Intelligence, Automation, and Work	197
	Daron Acemoglu and Pascual Restrepo	
	9. Artificial Intelligence and Economic Growth	237
	Philippe Aghion, Benjamin F. Jones, and Charles I. Jones	
	<i>Comment:</i> Patrick Francois	
	10. Artificial Intelligence and Jobs: The Role of Demand	291
	James Bessen	
	11. Public Policy in an AI Economy	309
	Austan Goolsbee	
	12. Should We Be Reassured If Automation in the Future Looks Like Automation in the Past?	317
	Jason Furman	
	13. R&D, Structural Transformation, and the Distribution of Income	329
	Jeffrey D. Sachs	
	14. Artificial Intelligence and Its Implications for Income Distribution and Unemployment	349
	Anton Korinek and Joseph E. Stiglitz	
	15. Neglected Open Questions in the Economics of Artificial Intelligence	391
	Tyler Cowen	

III. MACHINE LEARNING AND REGULATION

- | | |
|--|-----|
| 16. Artificial Intelligence, Economics, and Industrial Organization | 399 |
| Hal Varian | |
| <i>Comment:</i> Judith Chevalier | |
| 17. Privacy, Algorithms, and Artificial Intelligence | 423 |
| Catherine Tucker | |
| 18. Artificial Intelligence and Consumer Privacy | 439 |
| Ginger Zhe Jin | |
| 19. Artificial Intelligence and International Trade | 463 |
| Avi Goldfarb and Daniel Trefler | |
| 20. Punishing Robots: Issues in the Economics of Tort Liability and Innovation in Artificial Intelligence | 493 |
| Alberto Galasso and Hong Luo | |

IV. MACHINE LEARNING AND ECONOMICS

- | | |
|---|-----|
| 21. The Impact of Machine Learning on Economics | 507 |
| Susan Athey | |
| <i>Comment:</i> Mara Lederman | |
| 22. Artificial Intelligence, Labor, Productivity, and the Need for Firm-Level Data | 553 |
| Manav Raj and Robert Seamans | |
| 23. How Artificial Intelligence and Machine Learning Can Impact Market Design | 567 |
| Paul R. Milgrom and Steven Tadelis | |
| 24. Artificial Intelligence and Behavioral Economics | 587 |
| Colin F. Camerer | |
| <i>Comment:</i> Daniel Kahneman | |
| Contributors | 611 |
| Author Index | 615 |
| Subject Index | 625 |

Acknowledgments

This volume contains chapters and ideas discussed at the first NBER Conference on the Economics of Artificial Intelligence, held in September 2017 in Toronto. We thank all the authors and discussants for their contributions. Funds for the conference and book project were provided by the Sloan Foundation, the Canadian Institute for Advanced Research, and the Creative Destruction Lab at the University of Toronto. At the Sloan Foundation, Danny Goroff provided guidance that improved the overall agenda. The NBER digitization initiative, under the leadership of Shane Greenstein, was a key early supporter. We thank our dean, Tiff Macklem. In addition, Jim Poterba at the NBER has been generous, giving us the flexibility needed to bring this project together. Special thanks are due to Rob Shannon, Denis Healy, Carl Beck, and Dawn Bloomfield for managing the conference and logistics and to Helena Fitz-Patrick for guiding the book through the editorial process. Finally we thank our families, Gina, Natalie, Rachel, Amelia, Andreas, Belanna, Ariel, Annika, Anna, Sam, and Ben.

Ajay Agrawal, Joshua Gans, and Avi Goldfarb

Artificial intelligence (AI) technologies have advanced rapidly over the last several years. As the technology continues to improve, it may have a substantial impact on the economy with respect to productivity, growth, inequality, market power, innovation, and employment. In 2016, the White House put out several reports emphasizing this potential impact. Despite its importance, there is little economics research on the topic. The research that exists is derived from past technologies (such as factory robots) that capture only part of the economic reach of AI. Without a better understanding of how AI might impact the economy, we cannot design policy to prepare for these changes.

To address these challenges, the National Bureau of Economic Research held its first conference on the Economics of Artificial Intelligence in September 2017 in Toronto, with support from the NBER Economics Digitization Initiative, the Sloan Foundation, the Canadian Institute for Advanced Research, and the University of Toronto's Creative Destruction Lab. The purpose of the conference was to set the research agenda for economists working on AI. The invitation emphasized these points as follows:

Ajay Agrawal is the Peter Munk Professor of Entrepreneurship at the Rotman School of Management, University of Toronto, and a research associate of the National Bureau of Economic Research. Joshua Gans is professor of strategic management and holder of the Jeffrey S. Skoll Chair of Technical Innovation and Entrepreneurship at the Rotman School of Management, University of Toronto (with a cross appointment in the Department of Economics), and a research associate of the National Bureau of Economic Research. Avi Goldfarb holds the Rotman Chair in Artificial Intelligence and Healthcare and is professor of marketing at the Rotman School of Management, University of Toronto, and a research associate of the National Bureau of Economic Research.

For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c14005.ack>.

The context is this: imagine back to 1995 when the internet was about to begin transforming industries. What would have happened to economic research into that revolution had the leading economists gathered to scope out a research agenda at that time? Today, we are facing the same opportunity with regard to AI. This time around we are convening a group of 30 leading economists to scope out the research agenda for the next 20 years into the economics of AI.

Scholars who accepted the invitation were asked to write up and present ideas around a specific topic related to their expertise. For each paper, a discussant was assigned. Throughout the conference, in presentations, discussions, and debates, participants weighed in with their ideas for what the key questions will be, what research has already shown, and where the challenges will lie. Pioneering AI researchers Geoffrey Hinton, Yann LeCun, and Russ Salakhutdinov attended, providing useful context and detail about the current and expected future capabilities of the technology. The conference was unique because it emphasized the work that still needs to be done, rather than the presentation of standard research papers. Participants had the freedom to engage in informed speculation and healthy debate about the most important areas of inquiry.

This volume contains a summary of the proceedings of the conference. We provided authors with few constraints. This meant diversity in topics and chapter style. Many of the chapters contained herein are updated versions of the original papers and presentations at the conference. Some discussants commented directly on the chapters while others went further afield, emphasizing concepts that did not make it into the formal presentations but instead arose as part of debate and discussion. The volume also contains a small number of chapters that were not presented at the conference, but nevertheless represent ideas that came up in the general discussion and that warranted inclusion in a volume describing the proceedings of the conference.

We categorize the chapters into four broad themes. First, several chapters emphasize the role of AI as a general purpose technology (GPT), building on the existing literature on general purpose technologies from the steam engine to the internet. Second, many chapters highlight the impact of AI on growth, jobs, and inequality, focusing on research and tools from macro and labor economics. Third, five chapters discuss machine learning and economic regulation, with an emphasis on microeconomic consequences and industrial organization. The final set of chapters explores how AI will affect research in economics.

Of course, these themes are not mutually exclusive. Discussion of AI as a GPT naturally leads to discussions of economic growth. Regulation can enhance or reduce inequality. And AI's impact on economics is a consequence of it being a general purpose technology for scientific discovery (as emphasized in chapter 4 by Cockburn, Henderson, and Stern). Furthermore, a handful of concepts cut across the various parts, most notably the

role of humans as AI improves and the interaction between technological advance and political economy.

Below, we summarize these four broad themes in detail. Before doing so, we provide a definition of the technology that brings together the various themes.

What Is Artificial Intelligence?

The Oxford English Dictionary defines artificial intelligence as “the theory and development of computer systems able to perform tasks normally requiring human intelligence.” This definition is both broad and fluid. There is an old joke among computer scientists that artificial intelligence defines what machines cannot yet do. Before a machine could beat a human expert at chess, such a win would mean artificial intelligence. After the famed match between IBM’s Deep Blue and Gary Kasparov, playing chess was called computer science and other challenges became artificial intelligence.

The chapters in this volume discuss three related, but distinct, concepts of artificial intelligence. First, there is the technology that has driven the recent excitement around artificial intelligence: machine learning. Machine learning is a branch of computational statistics. It is a tool of prediction in the statistical sense, taking information you have and using it to fill in information you do not have. Since 2012, the uses of machine learning as a prediction technology have grown substantially. One set of machine-learning algorithms, in particular, called “deep learning,” has been shown to be useful and commercially viable for a variety of prediction tasks from search engine design to image recognition to language translation. The chapter in the book authored by us—Agrawal, Gans, and Goldfarb—emphasizes that rapid improvements in prediction technology can have a profound impact on organizations and policy (chapter 3). The chapter by Taddy (chapter 2) defines prediction with machine learning as one component of a true artificial intelligence and provides detail on the various machine-learning technologies.

While the recent interest in AI is driven by machine learning, computer scientists and philosophers have emphasized the feasibility of a true artificial general intelligence that equals or exceeds human intelligence (Bostrom 2014; Kaplan 2016). The closing sentence of this volume summarizes this possibility bluntly. Daniel Kahneman writes, “I do not think that there is very much that we can do that computers will not eventually be programmed to do.” The economic and societal impact of machines that surpass human intelligence would be extraordinary. Therefore—whether such an event occurs imminently, in a few decades, in a millennium, or never—it is worth exploring the economic consequences of such an event. While not a focal aspect of any chapter, several of the chapters in this volume touch on the economic consequences of such superintelligent machines.

A third type of technology that is often labeled “artificial intelligence” is

better seen as a process: automation. Much of the existing empirical work on the impact of artificial intelligence uses data on factory automation through robotics. Daron Acemoglu and Pascual Restrepo use data on factory robots to explore the impact of AI and automation on work (chapter 8). Automation is a potential consequence of artificial intelligence, rather than artificial intelligence per se. Nevertheless, discussions of the consequences of artificial intelligence and automation are tightly connected.

While most chapters in the book focus on the first definition—artificial intelligence as machine learning—a prediction technology, the economic implications of artificial general intelligence and automation receive serious attention.

AI as a GPT

A GPT is characterized by pervasive use in a wide range of sectors combined with technological dynamism (Bresnahan and Trajtenberg 1995). General purpose technologies are enabling technologies that open up new opportunities. While electric motors did reduce energy costs, the productivity impact was largely driven by increased flexibility in the design and location of factories (David 1990). Much of the interest in artificial intelligence and its impact on the economy stems from its potential as a GPT. Human intelligence is a general purpose tool. Artificial intelligence, whether defined as prediction technology, general intelligence, or automation, similarly has potential to apply across a broad range of sectors.

Brynjolfsson, Rock, and Syverson (chapter 1) argue the case for AI as a GPT. They focus on machine learning and identify a variety of sectors in which machine learning is likely to have a broad impact. They note expected continual technological progress in machine learning and a number of complementary innovations that have appeared along with machine learning. By establishing AI as a GPT, they can turn to the general lessons of the productivity literature on GPTs with respect to initially low rates of productivity growth, organizational challenges, and adjustment costs. They propose four potential explanations for the surprisingly low measured productivity growth given rapid innovation in AI and related technologies—false hopes, mismeasurement, redistribution, and implementation lags—and conclude that lags due to missing complementary innovations are most likely the primary source of missing productivity growth: “an underrated area of research involves the complements to the new AI technologies, not only in areas of human capital and skills, but also new processes and business models. The intangible assets associated with the last wave of computerization were about ten times as large as the direct investments in computer hardware itself.”

Henderson’s comment emphasizes the impact of a GPT on employment and the distribution of income, directly linking the discussion of AI as a

GPT to questions addressed in the section on Growth, Jobs, and Inequality. She agrees with the central thesis “One of the reasons I like the paper so much is that it takes seriously an idea that economists long resisted—namely, that things as nebulous as ‘culture’ and ‘organizational capabilities’ might be (a) very important, (b) expensive, and (c) hard to change.” At the same time, she adds emphasis on additional implications: “I think that the authors may be underestimating the implications of this dynamic in important ways. . . . I’m worried about the transition problem at the societal level quite as much as I’m worried about it at the organizational level.”

The next chapters provide micro-level detail on the nature of AI as a technology. Taddy (chapter 2) provides a broad overview of the meaning of intelligence in computer science. He then provides some technical detail on two key machine-learning techniques, deep learning and reinforcement learning. He explains the technology in a manner intuitive to economists: “Machine learning is a field that thinks about how to automatically build robust predictions from complex data. It is closely related to modern statistics, and indeed many of the best ideas in ML have come from statisticians (the lasso, trees, forests, etc.). But whereas statisticians have often focused on *model inference*—on understanding the parameters of their models (e.g., testing on individual coefficients in a regression)—the ML community has been more focused on the single goal of maximizing predictive performance. The entire field of ML is calibrated against ‘out-of-sample’ experiments that evaluate how well a model trained on one data set will predict new data.”

Building on ideas in Agrawal, Gans, and Goldfarb (2018), we argue in chapter 3 that the current excitement around AI is driven by advances in prediction technology. We then show that modeling AI as a drop in the cost of prediction provides useful insight into the microeconomic impact of AI on organizations. We emphasize that AI is likely to substitute for human prediction, but complement other skills such as human judgment—defined as knowing the utility or valuation function: “a key departure from the usual assumptions of rational decision-making is that the decision-maker does not know the payoff from the risky action in each state and must apply *judgment* to determine the payoff. . . . Judgment does not come for free.”

Prat’s comment emphasizes that economists typically assume that the valuation function is given, and that loosening that assumption will lead to a deeper understanding of the impact of AI on organizations. He offers an example to illustrate: “Admissions offices of many universities are turning to AI to choose which applicants to make offers to. Algorithms can be trained on past admissions data. We observe the characteristics of applicants and the grades of past and present students. . . . The obvious problem is that we do not know how admitting someone who is likely to get high grades is going to affect the long-term payoff of our university. . . . Progress in AI should induce our university leaders to ask deeper questions about the relationship between student quality and the long-term goals of our higher-learning

institutions. These questions cannot be answered with AI, but rather with more theory-driven retrospective approaches or perhaps more qualitative methodologies.”

The next chapters explore AI as a GPT that will enhance science and innovation. After reviewing the history of artificial intelligence, Cockburn, Henderson, and Stern (chapter 4) provide empirical support for the widespread application of machine learning in general, and deep learning in particular, in scientific fields outside of computer science: “we develop what we believe is the first systematic database that captures the corpus of scientific paper and patenting activity in artificial intelligence . . . we find striking evidence for a rapid and meaningful shift in the application orientation of learning-oriented publications, particularly after 2009.” The authors make a compelling case for AI as a general purpose tool in the method of invention. The chapter concludes by discussing the implications for innovation policy and innovation management: “the potential commercial reward from mastering this mode of research is likely to usher in a period of racing, driven by powerful incentives for individual companies to acquire and control critical large data sets and application-specific algorithms.”

Mitchell’s comment emphasizes the regulatory effects of AI as a GPT for science and innovation—in terms of intellectual property, privacy, and competition policy: “It is not obvious whether AI is a general purpose technology for innovation or a very efficient method of imitation. The answer has a direct relevance for policy. A technology that made innovation cheaper would often (but not always) imply less need for strong IP protection, since the balance would swing toward limiting monopoly power and away from compensating innovation costs. To the extent that a technology reduces the cost of imitation, however, it typically necessitates greater protection.” Several later chapters detail these and other regulatory issues.

Agrawal, McHale, and Oettl (chapter 5) provide a recombinant growth model that explores how a general purpose technology for innovation could affect the rate of scientific discovery: “instead of emphasising the potential substitution of machines for workers in existing tasks, we emphasise the importance of AI in overcoming a specific problem that impedes human researchers—finding useful combinations in complex discovery spaces . . . we develop a relatively simple combinatorial-based knowledge production function that converges in the limit to the Romer/Jones function. . . . If the curse of dimensionality is both the blessing and curse of discovery, then advances in AI offer renewed hope of breaking the curse while helping to deliver on the blessing.” This idea of AI as an input into innovation is a key component of Cockburn, Henderson, and Stern (chapter 4), as well as in several later chapters. It is an important element of Aghion, Jones, and Jones’s model of the impact of AI on economic growth (chapter 9), emphasizing endogenous growth through AI (self-)improvements. It also underlies

the chapters focused on how AI will impact the way economics research is conducted (chapters 21 through 24).

The section on AI as a general purpose technology concludes with Manuel Trajtenberg's discussion of political and societal consequences (chapter 6). At the conference, Trajtenberg discussed Joel Mokyr's paper "The Past and Future of Innovation: Some Lessons from Economic History," which will be published elsewhere. The chapter therefore sits between a stand-alone chapter and a discussion. Trajtenberg's chapter does not comment directly on Mokyr, but uses Mokyr's paper as a jumping-off point to discuss how technology creates winners and losers, and the policy challenges associated with the political consequences of the diffusion of a GPT. "The sharp split between winners and losers, if left to its own, may have serious consequences far beyond the costs for the individuals involved: when it coincides with the political divide, it may threaten the very fabric of democracy, as we have seen recently both in America and in Europe. Thus, if AI bursts onto the scene and triggers mass displacement of workers, and demography plays out its fateful hand, the economy will be faced with a formidable dual challenge, that may require a serious reassessment of policy options . . . we need to anticipate the required institutional changes, to experiment in the design of new policies, particularly in education and skills development, in the professionalization of service occupations, and in affecting the direction of technical advance. Furthermore, economists possess a vast methodological arsenal that may prove very useful for that purpose—we should not shy away from stepping into this area, since its importance for the economy cannot be overstated." The next set of chapters also emphasize the distributional challenges of economic growth driven by rapid technological change.

Growth, Jobs, and Inequality

Much of the popular discussion around AI focuses on the impact on jobs. If machines can do what humans do, then will there still be work for humans in the future? The chapters in this section dig into the consequences of AI for jobs, economic growth, and inequality. Almost all chapters emphasize that technological change means an increase in wealth for society. As Jason Furman puts it in chapter 12, "We need more artificial intelligence." At the same time, it is clear that the impact of AI on society will depend on how the increased income from AI is distributed. The most recent GPTs to diffuse, computers and the internet, likely led to increased inequality due to skill-bias (e.g., Autor, Katz, and Krueger 1998; Akerman, Gaarder, and Mogstad 2015) and to an increased capital share (e.g., Autor et al. 2017). This section brings together those chapters that emphasize (largely macro-economic) ideas related to growth, inequality, and jobs. If the impact of AI will be like these other technologies, then what will the consequences

look like for inequality, political economy, economic growth, jobs, and the meaning of work?

Stevenson (chapter 7) outlines many of the key issues. She emphasizes that economists generally agree that in the long run society will be wealthier. She highlights issues with respect to the short run and income distribution. Summarizing both the tension in the public debate and the key themes in several other chapters, she notes, “In the end, there’s really two separate questions: there’s an employment question, in which the fundamental question is can we find fulfilling ways to spend our time if robots take our jobs? And there’s an income question, can we find a stable and fair distribution of income?”

Acemoglu and Restrepo (chapter 8) examine how AI and automation might change the nature of work. They suggest a task-based approach to understanding automation, emphasizing the relative roles of labor and capital in the economy. “At the heart of our framework is the idea that automation and thus AI and robotics replace workers in tasks that they previously performed, and via this channel, create a powerful *displacement effect*.” This will lead to a lower labor share of economic output. At the same time, productivity will increase and capital will accumulate, thereby increasing the demand for labor. More importantly, “we argue that there is a more powerful countervailing force that increases the demand for labor as well as the share of labor in the national income: the *creation of new tasks*, functions, and activities in which labor has a comparative advantage relative to machines. The creation of new tasks generates a *reinstatement effect* directly counterbalancing the *displacement effect*.” Like Stevenson, the long-run message is optimistic; however, a key point is that adjustment costs may be high. New skills are a necessary condition of the long-run optimistic forecast, and there is likely to be a short- and medium-term mismatch between skills and technologies. They conclude with a discussion of open questions about which skills are needed, the political economy of technological change (reinforcing ideas highlighted in the earlier chapter by Trajtenberg), and the interaction between inequality and the type of innovation enabled by automation going forward.

Aghion, Jones, and Jones (chapter 9) build on the task-based model, focusing on the impact on economic growth. They emphasize Baumol’s cost disease: “Baumol (1967) observed that sectors with rapid productivity growth, such as agriculture and even manufacturing today, often see their share of GDP decline while those sectors with relatively slow productivity growth—perhaps including many services—experience increases. As a consequence, economic growth may be constrained not by what we do well, but rather by what is essential and yet hard to improve. We suggest that combining this feature of growth with automation can yield a rich description of the growth process, including consequences for future growth and income distribution.” Thus, even in the limit where there is an artificial general intelligence that creates a singularity or intelligence explosion with a self-

improving AI, cost disease forces may constrain growth. This link between technological advance and Baumol's cost disease provides a fundamental limit to the most optimistic and the most pessimistic views. Scarcity limits both growth and the downside risk. The chapter also explores how AI might reduce economic growth if it makes it easier to imitate a rival's innovations, returning to issues of intellectual property highlighted in Mitchell's comment. Finally, they discuss inequality within and across firms. They note that AI will increase wages of the least skilled employees of technologically advanced firms, but also increasingly outsource the tasks undertaken by such employees.

Francois's comment takes this emphasis on cost disease as a starting point, asking what those tasks will be that humans are left to do. "But it is when we turn to thinking about what are the products or services where humans will remain essential in production that we start to run into problems. What if humans can't do anything better than machines? Many discussions at the conference centered around this very possibility. And I must admit that I found the scientists' views compelling on this. . . . The point I wish to make is that even in such a world where machines are better at all tasks, there will still be an important role for human 'work.' And that work will become the almost political task of managing the machines." He argues that humans must tell the machines what to optimize. Bostrom (2014) describes this as the value-loading problem. Francois emphasizes that this is largely a political problem, and links the challenges in identifying values with Arrow's ([1951] 1963) impossibility theorem. He identifies key questions around ownership of the machines, length of time that rents should accrue to those owners, and the political structure of decision-making. In raising these questions, he provides a different perspective on issues highlighted by Stevenson on the meaning of work and Trajtenberg on the political economy of technological change.

The discussion of the meaning of work is a direct consequence of concerns about the impact of AI on jobs. Jobs have been the key focus of public discussion on AI and the economy. If human tasks get automated, what is left for humans to do? Bessen (chapter 10) explores this question, using data about other technological advances to support his arguments. He emphasizes that technological change can lead to an increase in demand and so the impact of automation on jobs is ambiguous, even within a sector. "The reason automation in textiles, steel, and automotive manufacturing led to strong job growth has to do with the effect of technology on demand. . . . New technologies do not just replace labor with machines, but in a competitive market, automation will reduce prices. In addition, technology may improve product quality, customization, or speed of delivery. All of these things can increase demand. If demand increases sufficiently, employment will grow even though the labor required per unit of output declines."

Like Bessen, Goolsbee (chapter 11) notes that much of the popular dis-

cussion around AI relates to labor market consequences. Recognizing that those consequences matter, his chapter mostly emphasizes the positive: growth and productivity are good. Artificial Intelligence has potential to increase our standard of living. Like Acemoglu and Restrepo, he notes that the short-term displacement effects could be substantial. One frequently cited solution to the displacement effects of AI is a universal basic income, in which all members of society receive a cash transfer from the government. He then discusses the economics of such a policy and the numerous challenges to making it work. “First . . . in a world where AI-induced unemployment is already high, separating work and income is an advantage. In a world like the one we are in now, offering a basic income will likely cause a sizable drop in the labor market participation by low-wage groups. . . . Second, for a given amount of money to be used on redistribution, UBI likely shifts money away from the very poor. . . . Third, . . . converting things to a UBI and getting rid of the in-kind safety net will lead to a situation in which, even if among a small share of UBI recipients, SOME people will blow their money in unsympathetic ways—gambling, drugs, junk food, Ponzi schemes, whatever. And now those people will come to the emergency room or their kids will be hungry and by the rules, they will be out of luck. That’s what they were supposed to have used their UBI for.” Before concluding, he touches on a variety of regulatory issues that receive more detailed discussion in chapters 16 through 20. His conclusion mirrors that of Francois, emphasizing the importance of humans in determining policy direction, even if AI improves to the point where it surpasses human intelligence.

Furman (chapter 12) is similarly optimistic, emphasizing that we need more, not less AI. “AI is a critical area of innovation in the U.S. economy right now. At least to date, AI has not had a large impact on the aggregate performance of the macroeconomy or the labor market. But it will likely become more important in the years to come, bringing substantial opportunities – and our first impulse should be to embrace it fully.” Referencing data on productivity growth and on the diffusion of industrial robots, he then discusses potential negative effects on the economy as AI diffuses, particularly with respect to inequality and reduced labor force participation. The issues around labor force participation highlight the importance of Stevenson’s questions on the meaning of work. Like Goolsbee, Furman notes several challenges to implementing a universal basic income as a solution to these negative effects. He concludes that policy has an important role to play in enabling society to fully reap the benefits of technological change while minimizing the disruptive effects.

Returning to the question of labor share highlighted by Acemoglu and Restrepo, Sachs (chapter 13) emphasizes that the income share going to capital grows with automation: “Rather than Solow-era stylized facts, I would therefore propose the following alternative stylized facts: (a) the share of national income accruing to capital rises over time in sectors expe-

riencing automation, especially when capital is measured to include human capital; (b) the share of national income accruing to low-skill labor drops while the share accruing to high-skill labor rises; (c) the dynamics across sectors vary according to the differential timing of automation, with automation spreading from low-skilled and predictable tasks toward high-skilled and less predictable tasks; (d) automation reflects the rising intensity of science and technology throughout the economy . . . , and (e) future technological changes associated with AI are likely to shift national income from medium-skilled and high-skilled toward owners of business capital.” The chapter concludes with a list of key open questions about the dynamics of automation, the role of monopoly rents, and the consequences for income distribution and labor force participation.

Korinek and Stiglitz (chapter 14) also emphasize income distribution, discussing the implications of AI-related innovation for inequality. They show that, in a first-best economy, contracts can be specified in advance that make innovation Pareto improving. However, imperfect markets and costly redistribution can imply a move away from the first-best. Innovation may then drive inequality directly by giving innovators a surplus, or indirectly by changing the demand for different types of labor and capital. They discuss policies that could help reduce the increase in inequality, emphasizing different taxation tools. Related to the ideas introduced in Mitchell’s comment, they also explore IP policies: “If outright redistribution is infeasible, there may be other institutional changes that result in market distributions that are more favorable to workers. For example, intervention to steer technological progress may act as a second-best device . . . we provide an example in which a change in intellectual property rights—a shortening of the term of patent protection—effectively redistributes some of the innovators’ surplus to workers (consumers) to mitigate the pecuniary externalities on wages that they experience, with the ultimate goal that the benefits of the innovation are more widely shared.” Stiglitz and Korinek conclude with a more speculative discussion of artificial general intelligence (superhuman artificial intelligence), emphasizing that such a technological development will likely further increase inequality.

The final chapter in the section on growth, jobs, and inequality calls for a different emphasis. Cowen (chapter 15) emphasizes consumer surplus, international effects, and political economy. With respect to consumer surplus, he writes, “Imagine education and manufactured goods being much cheaper because we produced them using a greater dose of smart software. The upshot is that even if a robot puts you out of a job or lowers your pay, there will be some recompense on the consumer side.” Cowen also speculates that AI might hurt developing countries much more than developed, as automation means that labor cost reasons to offshore decline. Finally, like Trajtenberg and Francois, he emphasizes the political economy of AI, highlighting questions related to income distribution.

Taken together, the chapters in this section highlight several key issues and provide models that identify challenges related to growth, jobs, inequality, and politics. These models set up a number of theoretical and empirical questions about how AI will impact economic outcomes within and across countries.

The discussions are necessarily speculative because AI has not yet diffused widely, so research must either be entirely theoretical or it must use related technologies (such as factory robots) as a proxy for AI. The discussions are also speculative because of the challenges in measuring the relevant variables. In order to determine the impact of AI on the economy, we need consistent measures of AI, productivity, intangible capital, and growth across sectors, regions, and contexts. Going forward, to the extent that progress occurs against the proposed research agenda, it will depend on advances in measurement.

Machine Learning and Regulation

Industry will be a key innovator and adopter of artificial intelligence. A number of regulatory issues arise. The regulatory issues related to truly intelligent machines are touched on by Trajtenberg, Francois, Goolsbee, and Cowen. Mitchell's comment of Cockburn, Henderson, and Stern emphasizes intellectual property regulation. This section focuses on other regulatory challenges with respect to advances in machine learning.

Varian (chapter 16) sets up the issues by describing the key models from industrial organization that are relevant to understanding the impact of machine learning on firms. He highlights the importance of data as a scarce resource, and discusses the economics of data as an input: it is nonrival and it exhibits decreasing returns to scale in a technical sense (because prediction accuracy increases in the square root of N). He discusses the structure of ML-using industries including vertical integration, economies of scale, and the potential for price discrimination. He emphasizes the difference between learning by doing and data network effects: "There is a concept that is circulating among lawyers and regulators called 'data network effects.' The model is that a firm with more customers can collect more data and use this data to improve its product. This is often true—the prospect of improving operations is what makes ML attractive—but it is hardly novel. And it is certainly not a network effect! This is essentially a supply-side effect known as 'learning by doing.' . . . A company can have huge amounts of data, but if it does nothing with the data, it produces no value. In my experience, the problem is not lack of resources, but is lack of skills. A company that has data but no one to analyze it is in a poor position to take advantage of that data." He concludes by highlighting policy questions related to algorithmic collusion (which was discussed at the conference as "economist catnip,"

interesting and fun but unlikely to be of first-order importance), security, privacy, and transparency.

Chevalier's comment builds on Varian's emphasis on the importance of data, exploring the potential of antitrust policy aimed at companies that use machine learning. Legal scholars and policymakers have asked whether antitrust essential facilities doctrine should be applied to data ownership. She emphasizes the trade-off between static and dynamic considerations for such a policy: "In evaluating antitrust policies in innovative industries, it is important to recognize that consumer benefits from new technologies arise not just from obtaining goods and services at competitive prices, but also from the flow of new and improved products and services that arise from innovation. Thus, antitrust policy should be evaluated not just in terms of its effect on prices and outputs, but also on its effect on the speed of innovation. Indeed, in the high technology industries, it seems likely that these dynamic efficiency considerations dwarf the static efficiency considerations." She also explores several practical challenges.

Another regulatory issue that arises from the importance of data is privacy. Tucker (chapter 17) notes that machine learning uses data to make predictions about what individuals may desire, be influenced by, or do. She emphasizes that privacy is challenging for three reasons: cheap storage means that data may persist longer than the person who generated the data intended, nonrivalry means that data may be repurposed for uses other than originally intended, and externalities caused by data created by one individual that contains information about others: "For example, in the case of genetics, the decision to create genetic data has immediate consequences for family members, since one individual's genetic data is significantly similar to the genetic data of their family members. . . . There may also be spillovers across a person's decision to keep some information secret, if such secrecy predicts other aspects of that individual's behavior that AI might be able to project from." She discusses potential negative impacts of these three challenges, concluding with some key open questions.

Jin (chapter 18) also focuses on the importance of data as an input into machine learning. She emphasizes that reduced privacy creates security challenges, such as identity theft, ransomware, and misleading algorithms (such as Russian-sponsored posts in the 2016 US election): "In my opinion, the leading concern is that firms are not fully accountable for the risk they bring to consumer privacy and data security. To restore full accountability, one needs to overcome three obstacles, namely (a) the difficulty to observe firms' actual action in data collection, data storage, and data use; (b) the difficulty to quantify the consequence of data practice, especially before low-probability adverse events realize themselves; and (c) the difficulty to draw a causal link between a firm's data practice and its consequence." Combined, Tucker and Jin's chapters emphasize that any discussion of growth and

impact of AI requires an understanding of the privacy framework. Access to data drives innovation, underlies the potential for economic growth, and frames the antitrust debate.

The economics of data also create challenges with respect to the rules governing international trade. Goldfarb and Trefler (chapter 19) argue that economies of scale in data through feedback loops, along with economies of scope and knowledge externalities in AI innovation, could create the opportunity for country-level rents and strategic trade policy. At the same time, they emphasize that the geographic constraints on data and knowledge would have to be high for such a policy to be optimal at the country level. They highlight the rise of China: “China has become the focal point for much of the international discussion. The US narrative has it that Chinese protection has reduced the ability of dynamic US firms such as Google and Amazon to penetrate Chinese markets. This protection has allowed China to develop significant commercial AI capabilities, as evidenced by companies such as Baidu (a search engine like Google), Alibaba (an e-commerce web portal like Amazon), and Tencent (the developer of WeChat, which can be seen as combining the functions of Skype, Facebook, and Apple Pay) . . . we collected time-series data on the institutional affiliation of all authors of papers presented at a major AI research conference . . . we compare the 2012 and 2017 conferences. . . . While these countries all increased their absolute number of participants, in relative terms they all lost ground to China, which leapt from 10 percent in 2012 to 23 percent in 2017.” The authors discuss the international dimensions of domestic regulation related to privacy, access to government data, and industrial standards.

The final regulatory issue highlighted in this section is tort liability. Galasso and Luo (chapter 20) review prior literature on the relationship between liability and innovation. They emphasize the importance of getting the balance right between consumer protection and innovation incentives: “A central question in designing a liability system for AI technologies is how liability risk should be allocated between producers and consumers, and how this allocation might affect innovation. . . . A key promise of AI technologies is to achieve autonomy. With less room for consumers to take precautions, the relative liability burden is likely to shift toward producers, especially in situations in which producers are in a better position than individual users to control risk. . . . On the other hand, during the transitional period of an AI technology, substantial human supervision may still be required. . . . In many of these situations, it may be impractical or too costly for producers to monitor individual users and to intervene. Therefore, it would be important to maintain consumer liability to the extent that users of AI technologies have sufficient incentives to take precautions and invest in training, thus internalizing potential harm to others.”

Broadly, regulation will affect the speed at which AI diffuses. Too much regulation, and industry will not have incentives to invest. Too little regu-

lation, and consumers will not trust the products that result. In this way, getting the regulatory balance right is key to understanding when and how any impact of AI on economic growth and inequality will arise.

Impact on the Practice of Economics

Cockburn, Henderson, and Stern emphasize that machine learning is a general purpose technology for science and innovation. As such, it is likely to have an impact on research in a variety of disciplines, including economics. Athey (chapter 21) provides an overview of the various ways in which machine learning is likely to affect the practice of economics. For example: “I believe that machine learning (ML) will have a dramatic impact on the field of economics within a short time frame. . . . ML does not add much to questions about identification, which concern when the object of interest, for example, a causal effect, can be estimated with infinite data, but rather yields great improvements when the goal is semiparametric estimation or when there are a large number of covariates relative to the number of observations . . . a key advantage of ML is that ML views empirical analysis as “algorithms” that estimate and compare many alternative models . . . ‘outsourcing’ model selection to algorithms works very well when the problem is ‘simple’—for example, prediction and classification tasks, where performance of a model can be evaluated by looking at goodness of fit in a held-out test set.” She emphasizes the usefulness of machine-learning techniques for policy problems related to prediction (as in Kleinberg et al. 2015). The chapter then details recent advances in using machine-learning techniques in causal inference, which she views as a fundamental new tool kit for empirical economists. She concludes with a list of sixteen predictions of how machine learning will impact economics, emphasizing new econometric tools, new data sets and measurement techniques, increased engagement of economists as engineers (and plumbers), and, of course, increased study of the economic impact of machine learning on the economy as a whole.

Lederman’s comment emphasizes the usefulness of machine learning to create new variables for economic analysis, and how the use of machine learning by organizations creates a new kind of endogeneity problem: “We develop theoretical models to help us understand the data-generation process which, in turn, informs both our concerns about causality as well as the identification strategies we develop. . . . Overall, as applied researchers working with real-world data sets, we need to recognize that increasingly the data we are analyzing is going to be the result of decisions that are made by algorithms in which the decision-making process may or may not resemble the decision-making processes we model as social scientists.”

If the study of AI is going to be a key question for economists going forward, Raj and Seamans (chapter 22) emphasize that we need better data: “While there is generally a paucity of data examining the adoption, use, and

effects of both AI and robotics, there is currently less information available regarding AI. There are no public data sets on the utilization or adoption of AI at either the macro or micro level. The most complete source of information, the McKinsey Global Institute study, is proprietary and inaccessible to the general public or the academic community. The most comprehensive and widely used data set examining the diffusion of robotics is the International Federation of Robotics (IFR) Robot Shipment Data . . . the IFR does not collect any information on dedicated industrial robots that serve one purpose. Furthermore, some of the robots are not classified by industry, detailed data is only available for industrial robots (and not robots in service, transportation, warehousing, or other sectors), and geographical information is often aggregated” They provide a detailed discussion of data-collection opportunities by government and by academic researchers. If the agenda set up in the other chapters is to be answered, it is important to have a reliable data set that defines AI, measures its quality, and tracks its diffusion.

Related to Athey’s emphasis of increased engagement of economists as engineering, Milgrom and Tadelis (chapter 23) describe how machine learning is already affecting market-design decisions. Using specific examples from online marketplaces and telecommunications auctions, they emphasize the potential of AI to improve efficiency by predicting demand and supply, overcoming computational barriers, and reducing search frictions: “AI and machine learning are emerging as important tools for market design. Retailers and marketplaces such as eBay, Taobao, Amazon, Uber, and many others are mining their vast amounts of data to identify patterns that help them create better experiences for their customers and increase the efficiency of their markets . . . two-sided markets such as Google, which match advertisers with consumers, are not only using AI to set reserve prices and segment consumers into finer categories for ad targeting, but they also develop AI-based tools to help advertisers bid on ads. . . . Another important application of AI’s strength in improving forecasting to help markets operate more efficiently is in electricity markets. To operate efficiently, electricity market makers . . . must engage in demand and supply forecasting.” The authors argue that AI will play a substantial role in the design and implementation of markets over a wide range of applications.

Camerer (chapter 24) also emphasizes the role of AI as a tool for predicting choice: “Behavioral economics can be defined as the study of natural limits on computation, willpower, and self-interest, and the implications of those limits for economic analysis (market equilibrium, IO, public finance, etc.). A different approach is to define behavioral economics more generally, as simply being open-minded about what variables are likely to influence economic choices. . . . In a general ML approach, predictive features could be—and *should* be—any variables that predict. . . . If behavioral economics is recast as open-mindedness about what variables might predict, then ML is an ideal way to do behavioral economics because it can make use of

a wide set of variables and select which ones predict.” He argues that firms, policymakers, and market designers can implement AI as either a “bionic patch” that improves human decision-making or “malware” that exploits human weaknesses. In this way, AI could reduce or exacerbate the political economy and inequality issues highlighted in earlier chapters. In addition, Camerer explores two other ways in which AI and behavioral economics will interact. He hypothesizes that machine learning could help predict human behavior in a variety of settings including bargaining, risky choice, and games, helping to verify or reject theory. He also emphasizes that (poor) implementation of AI might provide insight into new ways to model biases in human decision-making.

The book concludes with Kahneman’s brief and insightful comment. Kahneman begins with a discussion of Camerer’s idea of using prediction to verify theory, but continues with a broader discussion of a variety of themes that arose over the course of the conference. With an optimistic tone, he emphasizes that there are no obvious limits to what artificial intelligence may be able to do: “Wisdom is breadth. Wisdom is not having too narrow a view. That is the essence of wisdom; it is broad framing. A robot will be endowed with broad framing. When it has learned enough, it will be wiser than we people because we do not have broad framing. We are narrow thinkers, we are noisy thinkers, and it is very easy to improve upon us. I do not think that there is very much that we can do that computers will not eventually be programmed to do.”

The Future of Research on the Economics of Artificial Intelligence

The chapters in this book are the beginning. They highlight key questions, recognize the usefulness of several economic models, and identify areas for further development. We can leverage what we know about GPTs to anticipate the impact of AI as it diffuses, recognizing that no two GPTs are identical. If AI is a general purpose technology, it is likely to lead to increased economic growth. A common theme in these chapters is that slowing down scientific progress—even if it were possible—would come at a significant cost. At the same time, many attendees emphasized that the distribution of the benefits of AI might not be even. It depends on who owns the AI, the effect on jobs, and the speed of diffusion.

The task given to the conference presenters was to scope out the research agenda. Perhaps more than anything, this volume highlights all that we do not know. It emphasizes questions around growth, inequality, privacy, trade, innovation, political economy, and so forth. We do not have answers yet. Of course, the lack of answers is a consequence of the early stage of AI’s diffusion. We cannot measure the impact until AI is widespread.

With the current state of measurement, however, we may never get answers. As highlighted in the chapter by Raj and Seamans, we do not have

good measures of AI. We also do not have a good measure of improvement to AI. What is the AI equivalent to the computational speed of a microchip or the horsepower of an internal combustion engine that will allow for quality-adjusted prices and measurement? We also do not have good measures of productivity growth when that growth is primarily driven by intangible capital. To answer these questions, the gross domestic product (GDP) measurement apparatus needs to focus on adjusting for intangible capital, software, and changes to the innovation process (Haskel and Westlake 2017). Furthermore, to the extent that the benefits of AI generate heterogeneous benefits to people as consumers and as workers, measurement of the benefit of AI will be tricky. For example, if AI enables more leisure and people choose to take more leisure, should that be accounted for in measures of inequality? If so, how?

While each chapter has its own take on the agenda, several themes cut across the volume as key aspects of the research agenda going forward. To the extent there is consensus on the questions, the consensus focuses on the potential of AI as a GPT, and the associated potential consequences on growth and inequality. A second consistent theme is the role of regulation in accelerating or constraining the diffusion of the technology. A third theme is that AI will change the way we do our work as economists. Finally, a number of issues appear in many chapters that are somewhat outside the standard economic models of technology's impact. How do people find meaning if AI replaces work with leisure? How can economists inform the policy debate on solutions proposed by technologists in the popular press such as taxing robots or a universal basic income? How does a technology's diffusion affect the political environment, and vice versa?

This book highlights the questions and provides direction. We hope readers of this book take it as a starting point for their own research into this new and exciting area of study.

References

- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston, MA: Harvard Business Review Press.
- Akerman, Anders, Ingvil Gaarder, and Magne Mogstad. 2015. "The Skill Complementarity of Broadband Internet." *Quarterly Journal of Economics* 130 (4): 1781–824.
- Arrow, Kenneth. (1951) 1963. *Social Choice and Individual Values*, 2nd ed. New York: John Wiley and Sons.
- Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2017. "The Fall of the Labor Share and the Rise of Superstar Firms." Working paper, Massachusetts Institute of Technology.
- Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. "Computing

-
- Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113 (4): 1169–213.
- Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.
- Bresnahan, Timothy F., and M. Trajtenberg. 1995. "General Purpose Technologies 'Engines of Growth?'" *Journal of Econometrics* 65:83–108.
- David, Paul A. 1990. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review Papers and Proceedings* 80 (2): 355–61.
- Haskel, Jonathan, and Stian Westlake. 2017. *Capitalism without Capital: The Rise of the Intangible Economy*. Princeton, NJ: Princeton University Press.
- Kaplan, Jerry. 2016. *Artificial Intelligence: What Everyone Needs to Know*. Oxford: Oxford University Press.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. 2015. "Prediction Policy Problems." *American Economic Review* 105 (5): 491–95.

Copyrighted image

AI as a GPT

Artificial Intelligence and the Modern Productivity Paradox A Clash of Expectations and Statistics

Erik Brynjolfsson, Daniel Rock, and Chad Syverson

The discussion around the recent patterns in aggregate productivity growth highlights a seeming contradiction. On the one hand, there are astonishing examples of potentially transformative new technologies that could greatly increase productivity and economic welfare (see Brynjolfsson and McAfee 2014). There are some early concrete signs of these technologies' promise, recent leaps in artificial intelligence (AI) performance being the most prominent example. However, at the same time, measured productivity growth over the past decade has slowed significantly. This deceleration is large, cutting productivity growth by half or more in the decade preceding the slowdown. It is also widespread, having occurred throughout the Organisation for Economic Co-operation and Development (OECD) and, more recently, among many large emerging economies as well (Syverson 2017).¹

Erik Brynjolfsson is director of the MIT Initiative on the Digital Economy, the Schussel Family Professor of Management Science and professor of information technology at the MIT Sloan School of Management, and a research associate of the National Bureau of Economic Research. Daniel Rock is a PhD candidate at the MIT Sloan School of Management and a researcher at the MIT Initiative on the Digital Economy. Chad Syverson is the Eli B. and Harriet B. Williams Professor of Economics at the University of Chicago Booth School of Business and a research associate of the National Bureau of Economic Research.

We thank Eliot Abrams, Ajay Agrawal, David Autor, Seth Benzell, Joshua Gans, Avi Goldfarb, Austan Goolsbee, Andrea Meyer, Guillaume Saint-Jacques, Manuel Trajtenberg, and numerous participants at the NBER Workshop on AI and Economics in September 2017. In particular, Rebecca Henderson provided detailed and very helpful comments on an earlier draft and Larry Summers suggested the analogy to the J-curve. Generous funding for this research was provided in part by the MIT Initiative on the Digital Economy. This is a minor revision of NBER Working Paper no. 24001. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c14007.ack>.

1. A parallel, yet more pessimistically oriented debate about potential technological progress is the active discussion about robots taking jobs from more and more workers (e.g., Brynjolfsson and McAfee 2011; Acemoglu and Restrepo 2017; Bessen 2017; Autor and Salomons 2017).

We thus appear to be facing a redux of the Solow (1987) paradox: we see transformative new technologies everywhere but in the productivity statistics.

In this chapter, we review the evidence and explanations for the modern productivity paradox and propose a resolution. Namely, there is no inherent inconsistency between forward-looking technological optimism and backward-looking disappointment. Both can simultaneously exist. Indeed, there are good conceptual reasons to *expect* them to simultaneously exist when the economy undergoes the kind of restructuring associated with transformative technologies. In essence, the forecasters of future company wealth and the measurers of historical economic performance show the greatest disagreement during times of technological change. In this chapter, we argue and present some evidence that the economy is in such a period now.

1.1 Sources of Technological Optimism

Paul Polman, Unilever’s CEO, recently claimed that “The speed of innovation has never been faster.” Similarly, Bill Gates, Microsoft’s cofounder, observes that “Innovation is moving at a scarily fast pace.” Vinod Khosla of Khosla Ventures sees “the beginnings of . . . [a] rapid acceleration in the next 10, 15, 20 years.” Eric Schmidt of Alphabet Inc., believes “we’re entering . . . the age of abundance [and] during the age of abundance, we’re going to see a new age . . . the age of intelligence.”² Assertions like these are especially common among technology leaders and venture capitalists.

In part, these assertions reflect the continuing progress of information technology (IT) in many areas, from core technology advances like further doublings of basic computer power (but from ever larger bases) to successful investment in the essential complementary innovations like cloud infrastructure and new service-based business models. But the bigger source of optimism is the wave of recent improvements in AI, especially machine learning (ML). Machine learning represents a fundamental change from the first wave of computerization. Historically, most computer programs were created by meticulously codifying human knowledge, mapping inputs to outputs as prescribed by the programmers. In contrast, machine-learning systems use categories of general algorithms (e.g., neural networks) to figure out relevant mappings on their own, typically by being fed very large sample data sets. By using these machine-learning methods that leverage the growth in total data and data-processing resources, machines have made impressive gains in perception and cognition, two essential skills for most

2. <http://www.khoslaventures.com/fireside-chat-with-google-co-founders-larry-page-and-sergey-brin>; https://en.wikipedia.org/wiki/Predictions_made_by_Ray_Kurzweil#2045:_The_Singularity; <https://www.theguardian.com/small-business-network/2017/jun/22/alphabets-eric-schmidt-google-artificial-intelligence-viva-technology-mckinsey>.

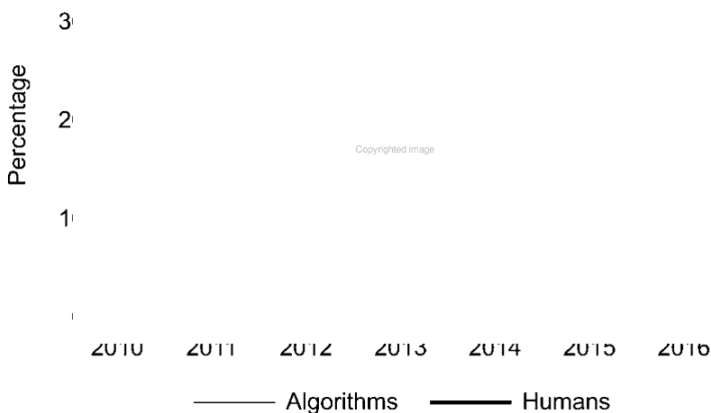


Fig. 1.1 AI versus human image recognition error rates

types of human work. For instance, error rates in labeling the content of photos on ImageNet, a data set of over ten million images, have fallen from over 30 percent in 2010 to less than 5 percent in 2016, and most recently as low as 2.2 percent with SE-ResNet152 in the ILSVRC2017 competition (see figure 1.1).³ Error rates in voice recognition on the Switchboard speech recording corpus, often used to measure progress in speech recognition, have decreased to 5.5 percent from 8.5 percent over the past year (Saon et al. 2017). The 5 percent threshold is important because that is roughly the performance of humans on each of these tasks on the same test data.

Although not at the level of professional human performance yet, Facebook's AI research team recently improved upon the best machine language translation algorithms available using convolutional neural net sequence prediction techniques (Gehring et al. 2017). Deep learning techniques have also been combined with reinforcement learning, a powerful set of techniques used to generate control and action systems whereby autonomous agents are trained to take actions given an environment state to maximize future rewards. Though nascent, advances in this field are impressive. In addition to its victories in the game of Go, Google DeepMind has achieved superhuman performance in many Atari games (Fortunato et al. 2017).

These are notable technological milestones. But they can also change the economic landscape, creating new opportunities for business value creation and cost reduction. For example, a system using deep neural networks was tested against twenty-one board-certified dermatologists and matched their

3. <http://image-net.org/challenges/LSVRC/2017/results>. ImageNet includes labels for each image, originally provided by humans. For instance, there are 339,000 labeled as flowers, 1,001,000 as food, 188,000 as fruit, 137,000 as fungus, and so on.

performance in diagnosing skin cancer (Esteva et al. 2017). Facebook uses neural networks for over 4.5 billion translations each day.⁴

An increasing number of companies have responded to these opportunities. Google now describes its focus as “AI first,” while Microsoft’s CEO Satya Nadella says AI is the “ultimate breakthrough” in technology. Their optimism about AI is not just cheap talk. They are making heavy investments in AI, as are Apple, Facebook, and Amazon. As of September 2017, these companies comprise the five most valuable companies in the world. Meanwhile, the tech-heavy NASDAQ composite index more than doubled between 2012 and 2017. According to CBInsights, global investment in private companies focused on AI has grown even faster, increasing from \$589 million in 2012 to over \$5 billion in 2016.⁵

1.2 The Disappointing Recent Reality

Although the technologies discussed above hold great potential, there is little sign that they have yet affected aggregate productivity statistics. Labor productivity growth rates in a broad swath of developed economies fell in the middle of the first decade of the twenty-first century and have stayed low since then. For example, aggregate labor productivity growth in the United States averaged only 1.3 percent per year from 2005 to 2016, less than half of the 2.8 percent annual growth rate sustained from 1995 to 2004. Fully twenty-eight of the twenty-nine other countries for which the OECD has compiled productivity growth data saw similar decelerations. The unweighted average annual labor productivity growth rate across these countries was 2.3 percent from 1995 to 2004, but only 1.1 percent from 2005 to 2015.⁶ What’s more, real median income has stagnated since the late 1990s and noneconomic measures of well-being, like life expectancy, have fallen for some groups (Case and Deaton 2017).

Figure 1.2 replicates the Conference Board’s analysis of its country-level Total Economy Database (Conference Board 2016). It plots highly smoothed annual productivity growth rate series for the United States, other mature economies (which combined match much of the OECD sample cited above), emerging and developing economies, and the world overall. The aforementioned slowdowns in the United States and other mature economies are clear in the figure. The figure also reveals that the productivity growth acceleration in emerging and developing economies during the first decade of the twenty-

4. <https://code.facebook.com/posts/289921871474277/transitioning-entirely-to-neural-machine-translation/>.

5. And the number of deals increased from 160 to 658. See <https://www.cbinsights.com/research/artificial-intelligence-startup-funding/>.

6. These slowdowns are statistically significant. For the United States, where the slowdown is measured using quarterly data, equality of the two periods’ growth rates is rejected with a t -statistic of 2.9. The OECD numbers come from annual data across the thirty countries. Here, the null hypothesis of equality is rejected with a t -statistic of 7.2.



Fig. 1.2 Smoothed average annual labor productivity growth (percent) by region

Source: The Conference Board Total Economy Database™ (adjusted version), November 2016.

Note: Trend growth rates are obtained using HP filter, assuming a $\lambda = 100$.

first century ended around the time of the Great Recession, causing a recent decline in productivity growth rates in these countries too.

These slowdowns do not appear to simply reflect the effects of the Great Recession. In the OECD data, twenty-eight of the thirty countries still exhibit productivity decelerations if 2008–2009 growth rates are excluded from the totals. Cetto, Fernald, and Mojon (2016), using other data, also find substantial evidence that the slowdowns began before the Great Recession.

Both capital deepening and total factor productivity (TFP) growth lead to labor productivity growth, and both seem to be playing a role in the slowdown (Fernald 2014; OECD 2015). Disappointing technological progress can be tied to each of these components. Total factor productivity directly reflects such progress. Capital deepening is indirectly influenced by technological change because firms’ investment decisions respond to improvements in capital’s current or expected marginal product.

These facts have been read by some as reasons for pessimism about the ability of new technologies like AI to greatly affect productivity and income. Gordon (2014, 2015) argues that productivity growth has been in long-run decline, with the IT-driven acceleration of 1995 to 2004 being a one-off aberration. While not claiming technological progress will be nil in the coming decades, Gordon essentially argues that we have been experiencing the new, low-growth normal and should expect to continue to do so going forward. Cowen (2011) similarly offers multiple reasons why innovation may be slow, at least for the foreseeable future. Bloom et al. (2017) document

that in many fields of technological progress research productivity has been falling, while Nordhaus (2015) finds that the hypothesis of an acceleration of technology-driven growth fails a variety of tests.

This pessimistic view of future technological progress has entered into long-range policy planning. The Congressional Budget Office, for instance, reduced its ten-year forecast for average US annual labor productivity growth from 1.8 percent in 2016 (CBO 2016) to 1.5 percent in 2017 (CBO 2017). Although perhaps modest on its surface, that drop implies US gross domestic product (GDP) will be considerably smaller ten years from now than it would in the more optimistic scenario—a difference equivalent to almost \$600 billion in 2017.

1.3 Potential Explanations for the Paradox

There are four principal candidate explanations for the current confluence of technological optimism and poor productivity performance: (a) false hopes, (b) mismeasurement, (c) concentrated distribution and rent dissipation, and (d) implementation and restructuring lags.⁷

1.3.1 False Hopes

The simplest possibility is that the optimism about the potential technologies is misplaced and unfounded. Perhaps these technologies won't be as transformative as many expect, and although they might have modest and noteworthy effects on specific sectors, their aggregate impact might be small. In this case, the paradox will be resolved in the future because realized productivity growth never escapes its current doldrums, which will force the optimists to mark their beliefs to market.

History and some current examples offer a quantum of credence to this possibility. Certainly one can point to many prior exciting technologies that did not live up to initially optimistic expectations. Nuclear power never became too cheap to meter, and fusion energy has been twenty years away for sixty years. Mars may still beckon, but it has been more than forty years since Eugene Cernan was the last person to walk on the moon. Flying cars never got off the ground,⁸ and passenger jets no longer fly at supersonic speeds. Even AI, perhaps the most promising technology of our era, is well behind Marvin Minsky's 1967 prediction that "Within a generation the problem of creating 'artificial intelligence' will be substantially solved" (Minsky 1967, 2).

On the other hand, there remains a compelling case for optimism. As we outline below, it is not difficult to construct back-of-the-envelope scenarios

7. To some extent, these explanations parallel the explanations for the Solow paradox (Brynjolfsson 1993).

8. But coming soon? <https://kittyhawk.aero/about/>.

in which even a modest number of currently existing technologies could combine to substantially raise productivity growth and societal welfare. Indeed, knowledgeable investors and researchers are betting their money and time on exactly such outcomes. Thus, while we recognize the potential for overoptimism—and the experience with early predictions for AI makes an especially relevant reminder for us to be somewhat circumspect in this chapter—we judge that it would be highly preliminary to dismiss optimism at this point.

1.3.2 Mismeasurement

Another potential explanation for the paradox is mismeasurement of output and productivity. In this case, it is the pessimistic reading of the empirical past, not the optimism about the future, that is mistaken. Indeed, this explanation implies that the productivity benefits of the new wave of technologies are already being enjoyed, but have yet to be accurately measured. Under this explanation, the slowdown of the past decade is illusory. This “mismeasurement hypothesis” has been put forth in several works (e.g., Mokyr 2014; Alloway 2015; Feldstein 2015; Hatzius and Dawsey 2015; Smith 2015).

There is a *prima facie* case for the mismeasurement hypothesis. Many new technologies, like smartphones, online social networks, and downloadable media involve little monetary cost, yet consumers spend large amounts of time with these technologies. Thus, the technologies might deliver substantial utility even if they account for a small share of GDP due to their low relative price. Guvenen et al. (2017) also show how growing offshore profit shifting can be another source of mismeasurement.

However, a set of recent studies provide good reason to think that mismeasurement is not the entire, or even a substantial, explanation for the slowdown. Cardarelli and Lusinyan (2015), Byrne, Fernald, and Reinsdorf (2016), Nakamura and Soloveichik (2015), and Syverson (2017), each using different methodologies and data, present evidence that mismeasurement is not the primary explanation for the productivity slowdown. After all, while there is convincing evidence that many of the benefits of today’s technologies are not reflected in GDP and therefore productivity statistics, the same was undoubtedly true in earlier eras as well.

1.3.3 Concentrated Distribution and Rent Dissipation

A third possibility is that the gains of the new technologies are already attainable, but that through a combination of concentrated distribution of those gains and dissipative efforts to attain or preserve them (assuming the technologies are at least partially rivalrous), their effect on average productivity growth is modest overall, and is virtually nil for the median worker. For instance, two of the most profitable uses of AI to date have been for targeting and pricing online ads, and for automated trading of financial instruments, both applications with many zero-sum aspects.

One version of this story asserts that the benefits of the new technologies are being enjoyed by a relatively small fraction of the economy, but the technologies' narrowly scoped and rivalrous nature creates wasteful "gold rush"-type activities. Both those seeking to be one of the few beneficiaries, as well as those who have attained some gains and seek to block access to others, engage in these dissipative efforts, destroying many of the benefits of the new technologies.⁹

Recent research offers some indirect support for elements of this story. Productivity differences between frontier firms and average firms in the same industry have been increasing in recent years (Andrews, Criscuolo, and Gal 2016; Furman and Orszag 2015). Differences in profit margins between the top and bottom performers in most industries have also grown (McAfee and Brynjolfsson 2008). A smaller number of superstar firms are gaining market share (Autor et al. 2017; Brynjolfsson et al. 2008), while workers' earnings are increasingly tied to firm-level productivity differences (Song et al. 2015). There are concerns that industry concentration is leading to substantial aggregate welfare losses due to the distortions of market power (e.g., De Loecker and Eeckhout 2017; Gutiérrez and Philippon 2017). Furthermore, growing inequality can lead to stagnating median incomes and associated socioeconomic costs, even when total income continues to grow.

Although this evidence is important, it is not dispositive. The aggregate effects of industry concentration are still under debate, and the mere fact that a technology's gains are not evenly distributed is no guarantee that resources will be dissipated in trying to capture them—especially that there would be enough waste to erase noticeable aggregate benefits.

1.3.4 Implementation and Restructuring Lags

Each of the first three possibilities, especially the first two, relies on explaining away the discordance between high hopes and disappointing statistical realities. One of the two elements is presumed to be somehow "wrong." In the misplaced optimism scenario, the expectations for technology by technologists and investors are off base. In the mismeasurement explanation, the tools we use to gauge empirical reality are not up to the task of accurately doing so. And in the concentrated distribution stories, the private gains for the few may be very real, but they do not translate into broader gains for the many.

But there is a fourth explanation that allows both halves of the seeming paradox to be correct. It asserts that there really is good reason to be optimistic about the future productivity growth potential of new technologies, while at the same time recognizing that recent productivity growth has been low. The core of this story is that it takes a considerable time—often more than

9. Stiglitz (2014) offers a different mechanism where technological progress with concentrated benefits in the presence of restructuring costs can lead to increased inequality and even, in the short run, economic downturns.

is commonly appreciated—to be able to sufficiently harness new technologies. Ironically, this is especially true for those major new technologies that ultimately have an important effect on aggregate statistics and welfare. That is, those with such broad potential application that they qualify as general purpose technologies (GPTs). Indeed, the more profound and far-reaching the potential restructuring, the longer the time lag between the initial invention of the technology and its full impact on the economy and society.

This explanation implies there will be a period in which the technologies are developed enough that investors, commentators, researchers, and policy-makers can imagine their potentially transformative effects, even though they have had no discernable effect on recent productivity growth. It isn't until a sufficient stock of the new technology is built and the necessary invention of complementary processes and assets occurs that the promise of the technology actually blossoms in aggregate economic data. Investors are forward looking and economic statistics are backward looking. In times of technological stability or steady change (constant velocity), the disjoint measurements will seem to track each other. But in periods of rapid change, the two measurements can become uncorrelated.

There are two main sources of the delay between recognition of a new technology's potential and its measurable effects. One is that it takes time to build the stock of the new technology to a size sufficient enough to have an aggregate effect. The other is that complementary investments are necessary to obtain the full benefit of the new technology, and it takes time to discover and develop these complements and to implement them. While the fundamental importance of the core invention and its potential for society might be clearly recognizable at the outset, the myriad necessary coinventions, obstacles, and adjustments needed along the way await discovery over time, and the required path may be lengthy and arduous. Never mistake a clear view for a short distance.

This explanation resolves the paradox by acknowledging that its two seemingly contradictory parts are not actually in conflict. Rather, both parts are in some sense natural manifestations of the same underlying phenomenon of building and implementing a new technology.

While each of the first three explanations for the paradox might have a role in describing its source, the explanations also face serious questions in their ability to describe key parts of the data. We find the fourth—the implementation and restructuring lags story—to be the most compelling in light of the evidence we discuss below. Thus it is the focus of our explorations in the remainder of this chapter.

1.4 The Argument in Favor of the Implementation and Restructuring Lags Explanation

Implicit or explicit in the pessimistic view of the future is that the recent slowdown in productivity growth portends slower productivity growth in the future.

We begin by establishing one of the most basic elements of the story: that slow productivity growth today does not rule out faster productivity growth in the future. In fact, the evidence is clear that it is barely predictive at all.

Total factor productivity growth is the component of overall output growth that cannot be explained by accounting for changes in observable labor and capital inputs. It has been called a “measure of our ignorance” (Abramovitz 1956). It is a residual, so an econometrician should not be surprised if it is not very predictable from past levels. Labor productivity is a similar measure, but instead of accounting for capital accumulation, simply divides total output by the labor hours used to produce that output.

Figures 1.3 and 1.4 plot, respectively, US productivity indices since 1948 and productivity growth by decade. The data include average labor productivity (LP), average total factor productivity (TFP), and Fernald’s (2014) utilization-adjusted TFP (TFPua).¹⁰

Productivity has consistently grown in the postwar era, albeit at different rates at different times. Despite the consistent growth, however, past productivity growth rates have historically been poor predictors of future productivity growth. In other words, the productivity growth of the past decade tells us little about productivity growth for the coming decade. Looking only at productivity data, it would have been hard to predict the decrease in productivity growth in the early 1970s or foresee the beneficial impact of IT in the 1990s.

As it turns out, while there is some correlation in productivity growth rates over short intervals, the correlation between adjacent ten-year periods is not statistically significant. We present below the results from a regression of different measures of average productivity growth on the previous period’s average productivity growth for ten-year intervals as well as scatterplots of productivity for each ten-year interval against the productivity in the subsequent period. The regressions in table 1.1 allow for autocorrelation in error terms across years (1 lag). Table 1.2 clusters the standard errors by decade. Similar results allowing for autocorrelation at longer time scales are presented in the appendix.

In all cases, the R^2 of these regressions is low, and the previous decade’s productivity growth does not have statistically discernable predictive power over the next decade’s growth. For labor productivity, the R^2 is 0.009. Although the intercept in the regression is significantly different from zero (productivity growth is positive, on average), the coefficient on the previous period’s growth is not statistically significant. The point estimate is economically small, too. Taking the estimate at face value, 1 percent higher annual labor productivity growth in the prior decade (around an unconditional mean of about 2 percent per year) corresponds to less than 0.1 percent

10. Available at <http://www.frbsf.org/economic-research/indicators-data/total-factor-productivity-tfp/>.




Fig. 1.3 US TFP and labor productivity indices, 1948–2016
Note: 1990 = 100.



Fig. 1.4 US TFP and labor productivity growth (percent) by decade

Table 1.1 Regressions with Newey-West standard errors

	(1) Labor productivity growth	(2) Total factor productivity growth	(3) Utilization-adjusted productivity growth
Newey-West regressions (1 lag allowed) ten-year average productivity growth			
Previous ten-year average LP growth	0.0857 (0.177)		
Previous ten-year average TFP growth		0.136 (0.158)	
Previous ten-year average TFP _{ua} growth			0.158 (0.187)
Constant	1.949*** (0.398)	0.911*** (0.188)	0.910*** (0.259)
Observations	50	50	50
R-squared	0.009	0.023	0.030

Note: Standard errors in parentheses.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Table 1.2 Regressions with standard errors clustered by decade

	(1) Labor productivity growth	(2) Total factor productivity growth	(3) Utilization-adjusted productivity growth
Ten-year average productivity growth (SEs clustered by decade)			
Previous ten-year average LP growth	0.0857 (0.284)		
Previous ten-year average TFP growth		0.136 (0.241)	
Previous ten-year average TFP _{ua} growth			0.158 (0.362)
Constant	1.949** (0.682)	0.911** (0.310)	0.910 (0.524)
Observations	50	50	50
R-squared	0.009	0.023	0.030

Note: Robust standard errors in parentheses.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.




Fig 1.5 Labor productivity growth scatterplot




Fig. 1.6 Total factor productivity growth scatterplot

faster growth in the following decade. In the TFP growth regression, the R^2 is 0.023, and again the coefficient on the previous period’s growth is insignificant. Similar patterns hold in the utilization-adjusted TFP regression (R^2 of 0.03). The lack of explanatory power of past productivity growth is also apparent in the scatterplots (see figures 1.5, 1.6, and 1.7).

The old adage that “past performance is not predictive of future results” applies well to trying to predict productivity growth in the years to come,



Fig. 1.7 Utilization-adjusted total factor productivity growth scatterplot

especially in periods of a decade or longer. Historical stagnation does not justify forward-looking pessimism.

1.5 A Technology-Driven Case for Productivity Optimism

Simply extrapolating recent productivity growth rates forward is not a good way to estimate the next decade’s productivity growth. Does that imply we have no hope at all of predicting productivity growth? We don’t think so.

Instead of relying only on past productivity statistics, we can consider the technological and innovation environment we expect to see in the near future. In particular, we need to study and understand the specific technologies that actually exist and make an assessment of their potential.

One does not have to dig too deeply into the pool of existing technologies or assume incredibly large benefits from any one of them to make a case that existing but still nascent technologies can potentially combine to create noticeable accelerations in aggregate productivity growth. We begin by looking at a few specific examples. We will then make the case that AI is a GPT, with broader implications.

First, let’s consider the productivity potential of autonomous vehicles. According to the US Bureau of Labor Statistics (BLS), in 2016 there were 3.5 million people working in private industry as “motor vehicle operators” of one sort or another (this includes truck drivers, taxi drivers, bus drivers, and other similar occupations). Suppose autonomous vehicles were to reduce, over some period, the number of drivers necessary to do the current workload to 1.5 million. We do not think this is a far-fetched scenario given the potential of the technology. Total nonfarm private employment in mid-

2016 was 122 million. Therefore, autonomous vehicles would reduce the number of workers necessary to achieve the same output to 120 million. This would result in aggregate labor productivity (calculated using the standard BLS nonfarm private series) increasing by 1.7 percent ($122/120 = 1.017$). Supposing this transition occurred over ten years, this single technology would provide a direct boost of 0.17 percent to annual productivity growth over that decade.

This gain is significant, and it does not include many potential productivity gains from complementary changes that could accompany the diffusion of autonomous vehicles. For instance, self-driving cars are a natural complement to transportation-as-a-service rather than individual car ownership. The typical car is currently parked 95 percent of the time, making it readily available for its owner or primary user (Morris 2016). However, in locations with sufficient density, a self-driving car could be summoned on demand. This would make it possible for cars to provide useful transportation services for a larger fraction of the time, reducing capital costs per passenger-mile, even after accounting for increased wear and tear. Thus, in addition to the obvious improvements in labor productivity from replacing drivers, capital productivity would also be significantly improved. Of course, the speed of adoption is important for estimation of the impact of these technologies. Levy (2018) is more pessimistic, suggesting in the near term that long distance truck driver jobs will grow about 2 percent between 2014 and 2024. This is 3 percent less (about 55,000 jobs in that category) than they would have grown without autonomous vehicle technology and about 3 percent of total employment of long distance truck drivers. A second example is call centers. As of 2015, there were about 2.2 million people working in more than 6,800 call centers in the United States, and hundreds of thousands more work as home-based call center agents or in smaller sites.¹¹ Improved voice-recognition systems coupled with intelligence question-answering tools like IBM's Watson might plausibly be able to handle 60–70 percent or more of the calls, especially since, in accordance with the Pareto principle, a large fraction of call volume is due to variants on a small number of basic queries. If AI reduced the number of workers by 60 percent, it would increase US labor productivity by 1 percent, perhaps again spread over ten years. Again, this would likely spur complementary innovations, from shopping recommendation and travel services to legal advice, consulting, and real-time personal coaching. Relatedly, citing advances in AI-assisted customer service, Levy (2018) projects zero growth in customer service representatives from 2014 to 2024 (a difference of 260,000 jobs from BLS projections).

Beyond labor savings, advances in AI have the potential to boost total factor productivity. In particular, energy efficiency and materials usage could be improved in many large-scale industrial plants. For instance, a

11. <https://info.siteselectiongroup.com/blog/how-big-is-the-us-call-center-industry-compared-to-india-and-philippines>.

team from Google DeepMind recently trained an ensemble of neural networks to optimize power consumption in a data center. By carefully tracking the data already collected from thousands of sensors tracking temperatures, electricity usage, and pump speeds, the system learned how to make adjustments in the operating parameters. As a result, the AI was able to reduce the amount of energy used for cooling by 40 percent compared to the levels achieved by human experts. The algorithm was a general-purpose framework designed to account complex dynamics, so it is easy to see how such a system could be applied to other data centers at Google, or indeed, around the world. Overall, data center electricity costs in the United States are about \$6 billion per year, including about \$2 billion just for cooling.¹²

What's more, similar applications of machine learning could be implemented in a variety of commercial and industrial activities. For instance, manufacturing accounts for about \$2.2 trillion of value added each year. Manufacturing companies like GE are already using AI to forecast product demand, future customer maintenance needs, and analyze performance data coming from sensors on their capital equipment. Recent work on training deep neural network models to perceive objects and achieve sensorimotor control have at the same time yielded robots that can perform a variety of hand-eye coordination tasks (e.g., unscrewing bottle caps and hanging coat hangers; Levine et al., [2016]). Liu et al. (2017) trained robots to perform a number of household chores, like sweeping and pouring almonds into a pan, using a technique called imitation learning.¹³ In this approach, the robot learns to perform a task using a raw video demonstration of what it needs to do. These techniques will surely be important for automating manufacturing processes in the future. The results suggest that artificial intelligence may soon improve productivity in household production tasks as well, which in 2010 were worth as much as \$2.5 trillion in nonmarket value added (Bridgman et al. 2012).¹⁴

Although these examples are each suggestive of nontrivial productivity gains, they are only a fraction of the set of applications for AI and machine learning that have been identified so far. James Manyika et al. (2017) analyzed 2,000 tasks and estimated that about 45 percent of the activities that people are paid to perform in the US economy could be automated using existing levels of AI and other technologies. They stress that the pace of

12. According to personal communication, August 24, 2017, with Jon Koomey, Arman Shehabi, and Sarah Smith of Lawrence Berkeley Lab.

13. Videos of these efforts available here: <https://sites.google.com/site/imitationfromobservation/>.

14. One factor that might temper the aggregate impact of AI-driven productivity gains is if product demand for the sectors with the largest productivity AI gains is sufficiently inelastic. In this case, these sectors' shares of total expenditure will shrink, shifting activity toward slower-growing sectors and muting aggregate productivity growth à la Baumol and Bowen (1966). It is unclear what the elasticities of demand are for the product classes most likely to be affected by AI.

automation will depend on factors other than technical feasibility, including the costs of automation, regulatory barriers, and social acceptance.

1.6 Artificial Intelligence Is a General Purpose Technology

As important as specific applications of AI may be, we argue that the more important economic effects of AI, machine learning, and associated new technologies stem from the fact that they embody the characteristics of general purpose technologies (GPTs). Bresnahan and Trajtenberg (1996) argue that a GPT should be pervasive, able to be improved upon over time, and be able to spawn complementary innovations.

The steam engine, electricity, the internal combustion engine, and computers are each examples of important general purpose technologies. Each of them increased productivity not only directly, but also by spurring important complementary innovations. For instance, the steam engine not only helped to pump water from coal mines, its most important initial application, but also spurred the invention of more effective factory machinery and new forms of transportation like steamships and railroads. In turn, these coinventions helped give rise to innovations in supply chains and mass marketing, to new organizations with hundreds of thousands of employees, and even to seemingly unrelated innovations like standard time, which was needed to manage railroad schedules.

Artificial intelligence, and in particular machine learning, certainly has the potential to be pervasive, to be improved upon over time, and to spawn complementary innovations, making it a candidate for an important GPT.

As noted by Agrawal, Gans, and Goldfarb (2017), the current generation of machine-learning systems is particularly suited for augmenting or automating tasks that involve at least some prediction aspect, broadly defined. These cover a wide range of tasks, occupations, and industries, from driving a car (predicting the correct direction to turn the steering wheel) and diagnosing a disease (predicting its cause) to recommending a product (predicting what the customer will like) and writing a song (predicting which note sequence will be most popular). The core capabilities of perception and cognition addressed by current systems are pervasive, if not indispensable, for many tasks done by humans.

Machine-learning systems are also designed to improve over time. Indeed, what sets them apart from earlier technologies is that they are designed to improve *themselves* over time. Instead of requiring an inventor or developer to codify, or code, each step of a process to be automated, a machine-learning algorithm can discover on its own a function that connects a set of inputs X to a set of outputs Y as long as it is given a sufficiently large set of labeled examples mapping some of the inputs to outputs (Brynjolfsson and Mitchell 2017). The improvements reflect not only the discovery of new algorithms and techniques, particularly for deep neural networks, but

also their complementarities with vastly more powerful computer hardware and the availability of much larger digital data sets that can be used to train the systems (Brynjolfsson and McAfee 2017). More and more digital data is collected as a byproduct of digitizing operations, customer interactions, communications, and other aspects of our lives, providing fodder for more and better machine-learning applications.¹⁵

Most important, machine-learning systems can spur a variety of complementary innovations. For instance, machine learning has transformed the abilities of machines to perform a number of basic types of perception that enable a broader set of applications. Consider machine vision—the ability to see and recognize objects, to label them in photos, and to interpret video streams. As error rates in identifying pedestrians improve from one per 30 frames to about one per 30 million frames, self-driving cars become increasingly feasible (Brynjolfsson and McAfee 2017).

Improved machine vision also makes practical a variety of factory automation tasks and medical diagnoses. Gill Pratt has made an analogy to the development of vision in animals 500 million years ago, which helped ignite the Cambrian explosion and a burst of new species on earth (Pratt 2015). He also noted that machines have a new capability that no biological species has: the ability to share knowledge and skills almost instantaneously with others. Specifically, the rise of cloud computing has made it significantly easier to scale up new ideas at much lower cost than before. This is an especially important development for advancing the economic impact of machine learning because it enables cloud robotics: the sharing of knowledge among robots. Once a new skill is learned by a machine in one location, it can be replicated to other machines via digital networks. Data as well as skills can be shared, increasing the amount of data that any given machine learner can use.

This in turn increases the rate of improvement. For instance, self-driving cars that encounter an unusual situation can upload that information with a shared platform where enough examples can be aggregated to infer a pattern. Only one self-driving vehicle needs to experience an anomaly for many vehicles to learn from it. Waymo, a subsidiary of Google, has cars driving 25,000 “real” autonomous and about 19 million simulated miles each week.¹⁶ All of the Waymo cars learn from the joint experience of the others. Similarly, a robot struggling with a task can benefit from sharing data and learnings with other robots that use a compatible knowledge-representation framework.¹⁷

When one thinks of AI as a GPT, the implications for output and welfare gains are much larger than in our earlier analysis. For example, self-driving cars could substantially transform many nontransport industries.

15. For example, through enterprise resource planning systems in factories, internet commerce, mobile phones, and the “Internet of Things.”

16. <http://ben-evans.com/benedictevans/2017/8/20/winner-takes-all>.

17. Rethink Robotics is developing exactly such a platform.

Retail could shift much further toward home delivery on demand, creating consumer welfare gains and further freeing up valuable high-density land now used for parking. Traffic and safety could be optimized, and insurance risks could fall. With over 30,000 deaths due to automobile crashes in the United States each year, and nearly a million worldwide, there is an opportunity to save many lives.¹⁸

1.7 Why Future Technological Progress Is Consistent with Low Current Productivity Growth

Having made a case for technological optimism, we now turn to explaining why it is not inconsistent with—and in fact may even be naturally related to—low current productivity growth.

Like other GPTs, AI has the potential to be an important driver of productivity. However, as Jovanovic and Rousseau (2005) point out (with additional reference to David's [1991] historical example), "a GPT does not deliver productivity gains immediately upon arrival" (1184). The technology can be present and developed enough to allow some notion of its transformative effects even though it is not affecting current productivity levels in any noticeable way. This is precisely the state that we argue the economy may be in now.

We discussed earlier that a GPT can at one moment both be present and yet not affect current productivity growth if there is a need to build a sufficiently large stock of the new capital, or if complementary types of capital, both tangible and intangible, need to be identified, produced, and put in place to fully harness the GPT's productivity benefits.

The time necessary to build a sufficient capital stock can be extensive. For example, it was not until the late 1980s, more than twenty-five years after the invention of the integrated circuit, that the computer capital stock reached its long-run plateau at about 5 percent (at historical cost) of total nonresidential equipment capital. It was at only half that level ten years prior. Thus, when Solow pointed out his now eponymous paradox, the computers were *finally just then* getting to the point where they really could be seen everywhere.

David (1991) notes a similar phenomenon in the diffusion of electrification. At least half of US manufacturing establishments remained unelectrified until 1919, about thirty years after the shift to polyphase alternating current began. Initially, adoption was driven by simple cost savings in pro-

18. These latter two consequences of autonomous vehicles, while certainly reflecting welfare improvements, would need to be capitalized in prices of goods or services to be measured in standard GDP and productivity measures. We will discuss AI-related measurement issues in greater depth later. Of course, it is worth remembering that autonomous vehicles also hold the potential to create new economic costs if, say, the congestion from lower marginal costs of operating a vehicle is not counteracted by sufficiently large improvements in traffic management technology or certain infrastructure investments.

viding motive power. The biggest benefits came later, when complementary innovations were made. Managers began to fundamentally reorganize work by replacing factories' centralized power source and giving every individual machine its own electric motor. This enabled much more flexibility in the location of equipment and made possible effective assembly lines of materials flow.

This approach to organizing factories is obvious in retrospect, yet it took as many as thirty years for it to become widely adopted. Why? As noted by Henderson (1993, 2006), it is exactly *because* incumbents are designed around the current ways of doing things and so proficient at them that they are blind to or unable to absorb the new approaches and get trapped in the status quo—they suffer the “curse of knowledge.”¹⁹

The factory electrification example demonstrates the other contributor to the time gap between a technology's emergence and its measured productivity effects: the need for installation (and often invention) of complementary capital. This includes both tangible and intangible investments. The timeline necessary to invent, acquire, and install these complements is typically more extensive than the time-to-build considerations just discussed. Consider the measured lag between large investments in IT and productivity benefits within firms. Brynjolfsson and Hitt (2003) found that while small productivity benefits were associated with firms' IT investments when one-year differences were considered, the benefits grew substantially as longer differences were examined, peaking after about seven years. They attributed this pattern to the need for complementary changes in business processes. For instance, when implementing large enterprise-planning systems, firms almost always spend several times more money on business process redesign and training than on the direct costs of hardware and software. Hiring and other human-resources practices often need considerable adjustment to match the firm's human capital to the new structure of production. In fact, Bresnahan, Brynjolfsson, and Hitt (2002) find evidence of three-way complementarities between IT, human capital, and organizational changes in the investment decisions and productivity levels. Furthermore, Brynjolfsson, Hitt, and Yang (2002) show each dollar of IT capital stock is correlated with about \$10 of market value. They interpret this as evidence of substantial IT-related intangible assets and show that firms that combine IT investments with a specific set of organizational practices are not just more productive, they also have disproportionately higher market values than firms that invest in only one or the other. This pattern in the data is consistent with a long stream of research on the importance of organizational and even

19. Atkeson and Kehoe (2007) note manufacturers' reluctance to abandon their large knowledge stock at the beginning of the transition to electric power to adopt what was, initially, only a marginally superior technology. David and Wright (2006) are more specific, focusing on “the need for organizational and above all for *conceptual* changes in the ways tasks and products are defined and structured” (147, emphasis in original).

cultural change when making IT investments and technology investments more generally (e.g., Aral, Brynjolfsson, and Wu 2012; Brynjolfsson and Hitt 2000; Orlikowski 1996; Henderson 2006).

But such changes take substantial time and resources, contributing to organizational inertia. Firms are complex systems that require an extensive web of complementary assets to allow the GPT to fully transform the system. Firms that are attempting transformation often must reevaluate and reconfigure not only their internal processes but often their supply and distribution chains as well. These changes can take time, but managers and entrepreneurs will direct invention in ways that economize on the most expensive inputs (Acemoglu and Restrepo 2017). According to LeChatelier's principle (Milgrom and Roberts 1996), elasticities will therefore tend to be greater in the long run than in the short run as quasi-fixed factors adjust.

There is no assurance that the adjustments will be successful. Indeed, there is evidence that the modal transformation of GPT-level magnitude fails. Alon et al. (2017) find that cohorts of firms over five years old contribute little to aggregate productivity growth on net—that is, among established firms, productivity improvements in one firm are offset by productivity declines in other firms. It is hard to teach the proverbial old dog new tricks. Moreover, the old dogs (companies) often have internal incentives to not learn them (Arrow 1962; Holmes, Levine, and Schmitz 2012). In some ways, technology advances in industry one company death at a time.

Transforming industries and sectors requires still more adjustment and reconfiguration. Retail offers a vivid example. Despite being one of the biggest innovations to come out of the 1990s dot-com boom, the largest change in retail in the two decades that followed was not e-commerce, but instead the expansion of warehouse stores and supercenters (Hortaçsu and Syverson 2015). Only very recently did e-commerce become a force for general retailers to reckon with. Why did it take so long? Brynjolfsson and Smith (2000) document the difficulties incumbent retailers had in adapting their business processes to take full advantage of the internet and electronic commerce. Many complementary investments were required. The sector as a whole required the build out of an entire distribution infrastructure. Customers had to be “retrained.” None of this could happen quickly. The potential of e-commerce to revolutionize retailing was widely recognized, and even hyped in the late 1990s, but its actual share of retail commerce was miniscule, 0.2 percent of all retail sales in 1999. Only after two decades of widely predicted yet time-consuming change in the industry, is e-commerce starting to approach 10 percent of total retail sales and companies like Amazon are having a first-order effect on more traditional retailers' sales and stock market valuations.

The case of self-driving cars discussed earlier provides a more prospective example of how productivity might lag technology. Consider what happens to the current pools of vehicle production and vehicle operation workers

when autonomous vehicles are introduced. Employment on production side will initially increase to handle research and development (R&D), AI development, and new vehicle engineering. Furthermore, learning curve issues could well imply lower productivity in manufacturing these vehicles during the early years (Levitt, List, and Syverson 2013). Thus labor input in the short run can actually increase, rather than decrease, for the same amount of vehicle production. In the early years of autonomous vehicle development and production, the marginal labor added by producers exceeds the marginal labor displaced among the motor vehicle operators. It is only later when the fleet of deployed autonomous vehicles gets closer to a steady state that measured productivity reflects the full benefits of the technology.

1.8 Viewing Today's Paradox through the Lens of Previous General Purpose Technologies

We have indicated in the previous discussion that we see parallels between the current paradox and those that have happened in the past. It is closely related to the Solow paradox era circa 1990, certainly, but it is also tied closely to the experience during the diffusion of portable power (combining the contemporaneous growth and transformative effects of electrification and the internal combustion engine).

Comparing the productivity growth patterns of the two eras is instructive. Figure 1.8 is an updated version of an analysis from Syverson (2013). It overlays US labor productivity since 1970 with that from 1890 to 1940, the period after portable power technologies had been invented and were starting to be placed into production. (The historical series values are from Kendrick [1961].) The modern series timeline is indexed to a value of 100 in 1995 and

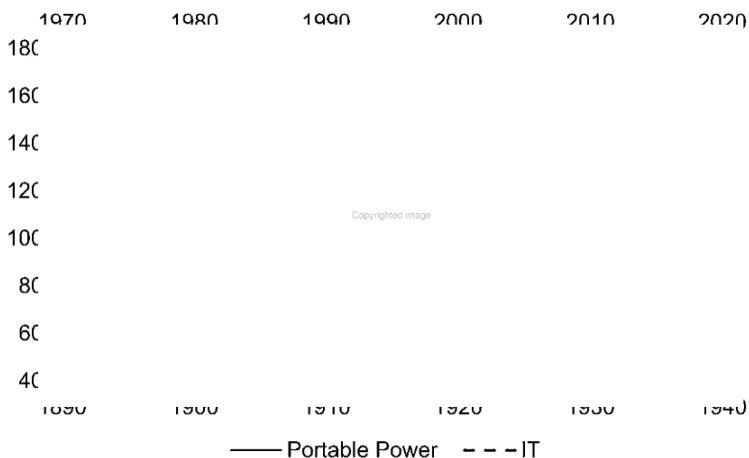


Fig. 1.8 Labor productivity growth in the portable power and IT eras

is labeled on the upper horizontal axis. The portable power era index has a value of 100 in 1915, and its years are shown on the lower horizontal axis.

Labor productivity during the portable power era shared remarkably similar patterns with the current series. In both eras, there was an initial period of roughly a quarter century of relatively slow productivity growth. Then both eras saw decade-long accelerations in productivity growth, spanning 1915 to 1924 in the portable power era and 1995 to 2004 more recently.

The late-1990s acceleration was the (at least partial) resolution of the Solow paradox. We imagine that the late 1910s acceleration could have similarly answered some economist's query in 1910 as to why one sees electric motors and internal combustion engines everywhere but in the productivity statistics.²⁰

Very interesting, and quite relevant to the current situation, the productivity growth slowdown we have experienced after 2004 also has a parallel in the historical data, a slowdown from 1924 to 1932. As can be seen in the figure, and instructive to the point of whether a new wave of AI and associated technologies (or if one prefers, a second wave of IT-based technology) could reaccelerate productivity growth, labor productivity growth at the end of the portable power era rose again, averaging 2.7 percent per year between 1933 and 1940.

Of course this past breakout growth is no guarantee that productivity must speed up again today. However, it does raise two relevant points. First, it is another example of a period of sluggish productivity growth followed by an acceleration. Second, it demonstrates that productivity growth driven by a core GPT can arrive in multiple waves.

1.9 Expected Productivity Effects of an AI-Driven Acceleration

To understand the likely productivity effects of AI, it is useful to think of AI as a type of capital, specifically a type of intangible capital. It can be accumulated through investment, it is a durable factor of production, and its value can depreciate. Treating AI as a type of capital clarifies how its development and installation as a productive factor will affect productivity.

As with any capital deepening, increasing AI will raise labor productivity. This would be true regardless of how well AI capital is measured (which we might expect it won't be for several reasons discussed below) though there may be lags.

The effects of AI on TFP are more complex and the impact *will* depend on its measurement. If AI (and its output elasticity) were to be measured perfectly and included in both the input bundle in the denominator of TFP

20. We are not aware of anyone who actually said this, and of course today's system of national economic statistics did not exist at that time, but we find the scenario amusing, instructive, and in some ways plausible.

and the output bundle in the numerator, then measured TFP will accurately reflect true TFP. In this case, AI could be treated just like any other measurable capital input. Its effect on output could be properly accounted for and “removed” by the TFP input measure, leading to no change in TFP. This isn’t to say that there would not be productive benefits from diffusion of AI; it is just that it could be valued like other types of capital input.

There are reasons why economists and national statistical agencies might face measurement problems when dealing with AI. Some are instances of more general capital measurement issues, but others are likely to be idiosyncratic to AI. We discuss this next.

1.10 Measuring AI Capital

Regardless of the effects of AI and AI-related technologies on actual output and productivity, it is clear from the productivity outlook that the ways AI’s effects will be *measured* are dependent on how well countries’ statistics programs measure AI capital.

The primary difficulty in AI capital measurement is, as mentioned earlier, that many of its outputs will be intangible. This issue is exacerbated by the extensive use of AI as an input in making other capital, including new types of software, as well as human and organizational capital, rather than final consumption goods. Much of this other capital, including human capital, will, like AI itself, be mostly intangible (Jones and Romer 2010).

To be more specific, effective use of AI requires developing data sets, building firm-specific human capital, and implementing new business processes. These all require substantial capital outlays and maintenance. The tangible counterparts to these intangible expenditures, including purchases of computing resources, servers, and real estate, are easily measured in the standard neoclassical growth accounting model (Solow 1957). On the other hand, the value of capital goods production for complementary intangible investments is difficult to quantify. Both tangible and intangible capital stocks generate a capital service flow yield that accrues over time. Realizing these yields requires more than simply renting capital stock. After purchasing capital assets, firms incur additional adjustment costs (e.g., business process redesigns and installation costs). These adjustment costs make capital less flexible than frictionless rental markets would imply. Much of the market value of AI capital specifically, and IT capital more generally, may be derived from the capitalized short-term quasi-rents earned by firms that have already reorganized to extract service flows from new investment.

Yet while the stock of tangible assets is booked on corporate balance sheets, expenditures on the intangible complements and adjustment costs to AI investment commonly are not. Without including the production and use of intangible AI capital, the usual growth accounting decompositions of changes in value added can misattribute AI intangible capital deepening

to changes in TFP. As discussed in Hall (2000) and Yang and Brynjolfsson (2001), this constitutes an omission of a potentially important component of capital goods production in the calculation of final output. Estimates of TFP will therefore be inaccurate, though possibly in either direction. In the case where the intangible AI capital stock is growing faster than output, then TFP growth will be underestimated, while TFP will be overestimated if capital stock is growing more slowly than output.

The intuition for this effect is that in any given period t , the output of (unmeasured) AI capital stock in period $t + 1$ is a function the input (unmeasured) existing AI capital stock in period t . When AI stock is growing rapidly, the unmeasured outputs (AI capital stock created) will be greater than the unmeasured inputs (AI capital stock used).

Furthermore, suppose the relevant costs in terms of labor and other resources needed to create intangible assets are measured, but the resulting increases in intangible assets are not measured as contributions to output. In this case, not only will total GDP be undercounted but so will productivity, which uses GDP as its numerator. Thus periods of rapid intangible capital accumulation may be associated with *lower* measured productivity growth, even if true productivity is increasing.

With missing capital goods production, measured productivity will only reflect the fact that more capital and labor inputs are used up in producing measured output. The inputs used to produce unmeasured capital goods will instead resemble lost potential output. For example, a recent report from the Brookings Institution estimates that investments in autonomous vehicles have topped \$80 billion from 2014 to 2017 with little consumer adoption of the technology so far.²¹ This is roughly 0.44 percent of 2016 GDP (spread over three years). If all of the capital formation in autonomous vehicles was generated by equally costly labor inputs, this would lower estimated labor productivity by 0.1 percent per year over the last three years since autonomous vehicles have not yet led to any significant increase in measured final output. Similarly, according to the AI Index, enrollment in AI and ML courses at leading universities has roughly tripled over the past ten years, and the number of venture-back AI-related start-ups has more than quadrupled. To the extent that they create intangible assets beyond the costs of production, GDP will be underestimated.

Eventually the mismeasured intangible capital goods investments are expected to yield a return (i.e., output) by their investors. If and when measurable output is produced by these hidden assets, another mismeasurement effect leading to overestimation of productivity will kick in. When the output share and stock of mismeasured or omitted capital grows, the measured output increases produced by that capital will be incorrectly attributed to total factor productivity improvements. As the growth rate of investment in unmeasured capital goods decreases, the capital service flow from

21. <https://www.brookings.edu/research/gauging-investment-in-self-driving-cars/>.

unmeasured goods effect on TFP can exceed the underestimation error from unmeasured capital goods.

Combining these two effects produces a “J-curve” wherein early production of intangible capital leads to underestimation of productivity growth, but later returns from the stock of unmeasured capital creates measured output growth that might be incorrectly attributed to TFP.

Formally:

$$(1) \quad Y + zI_2 = f(A, K_1, K_2, L)$$

$$(2) \quad dY + z dI_2 = F_A dA + F_{K_1} dK_1 + F_L dL + F_{K_2} dK_2.$$

Output Y and unmeasured capital goods with price $z(zI_2)$ are produced with production function f . The inputs of $f(\cdot)$ are the total factor productivity A , ordinary capital K_1 , unmeasured capital K_2 , and labor L . Equation (2) describes the total differential of output as a function of the inputs to the production function. If the rental price of ordinary capital is r_1 , the rental price of unmeasured capital is r_2 , and the wage rate is w , we have

$$(3) \quad \hat{S} = \frac{dY}{Y} - \left(\frac{r_1 K_1}{Y} \right) \left(\frac{dK_1}{K_1} \right) - \left(\frac{wL}{Y} \right) \left(\frac{dL}{L} \right)$$

and

$$(4) \quad S^* = \frac{dY}{Y} - \left(\frac{r_1 K_1}{Y} \right) \left(\frac{dK_1}{K_1} \right) - \left(\frac{wL}{Y} \right) \left(\frac{dL}{L} \right) - \left(\frac{r_2 K_2}{Y} \right) \left(\frac{dK_2}{K_2} \right) + \left(\frac{zI_2}{Y} \right) \left(\frac{dI_2}{I_2} \right),$$

where \hat{S} is the familiar Solow residual as measured and S^* is the correct Solow residual accounting for mismeasured capital investments and stock.

The mismeasurement is then

$$(5) \quad \hat{S} - S^* = \left(\frac{r_2 K_2}{Y} \right) \left(\frac{dK_2}{K_2} \right) - \left(\frac{zI_2}{Y} \right) \left(\frac{dI_2}{I_2} \right) = \left(\frac{r_2 K_2}{Y} \right) g_{K_2} - \left(\frac{zI_2}{Y} \right) g_{I_2}$$

The right side of the equation describes a hidden capital effect and a hidden investment effect. When the growth rate of new investment in unmeasured capital multiplied by its share of output is larger (smaller) than the growth rate of the stock of unmeasured capital multiplied by its share of output, the estimated Solow residual will underestimate (overestimate) the rate of productivity growth. Initially, new types of capital will have a high marginal product. Firms will accumulate that capital until its marginal rate of return is equal to the rate of return of other capital. As capital accumulates, the growth rate of net investment in the unmeasured capital will turn negative, causing a greater overestimate TFP. In steady state, neither net investment's share of output nor the net stock of unmeasured capital grows and the productivity mismeasurement is zero. Figure 1.9 provides an illustration.²²

22. The price of new investment (z) and rental price of capital (r) are 0.3 and 0.12, respectively, in this toy economy. Other values used to create the figure are included in the appendix.

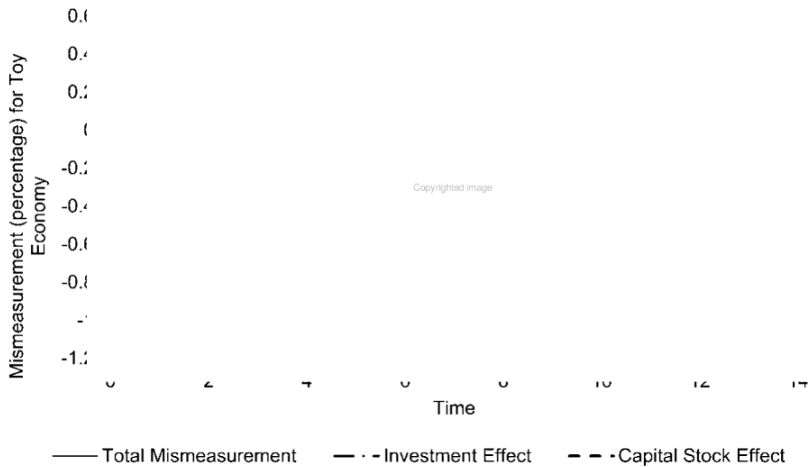


Fig. 1.9 The mismeasurement J-curve for an economy accumulating a new kind of capital

Looking forward, these problems may be particularly stark for AI capital, because its accumulation will almost surely outstrip the pace of ordinary capital accumulation in the short run. AI capital is a new category of capital—new in economic statistics, certainly, but we would argue practically so as well.

This also means that capital quantity indexes that are computed from within-type capital growth might have problems benchmarking size and effect of AI early on. National statistics agencies do not really focus on measuring capital types that are not already ubiquitous. New capital categories will tend to either be rolled into existing types, possibly with lower inferred marginal products (leading to an understatement of the productive effect of the new capital), or missed altogether. This problem is akin to the new goods problem in price indexes.

A related issue—once AI is measured separately—is how closely its units of measurement will capture AI’s marginal product relative to other capital stock. That is, if a dollar of AI stock has a marginal product that is twice as high as the modal unit of non-AI capital in the economy, will the quantity indexes of AI reflect this? This requires measured relative prices of AI and non-AI capital to capture differences in marginal product. Measuring levels correctly is less important than measuring accurate proportional differences (whether intertemporally or in the cross section) correctly. What is needed in the end is that a unit of AI capital twice as productive as another should be twice as large in the capital stock.

It is worth noting that these are all classic problems in capital measurement and not new to AI. Perhaps these problems will be systematically worse for AI, but this is not obvious *ex ante*. What it does mean is that econo-

mists and national statistical agencies at least have experience in, if not quite a full solution for, dealing with these sorts of limitations. That said, some measurement issues are likely to be especially prevalent for AI. For instance, a substantial part of the value of AI output may be firm-specific. Imagine a program that figures out individual consumers' product preferences or price elasticities and matches products and pricing to predictions. This has different value to different companies depending on their customer bases and product selection, and knowledge may not be transferrable across firms. The value also depends on companies' abilities to implement price discrimination. Such limits could come from characteristics of a company's market, like resale opportunities, which are not always under firms' control, or from the existence in the firm of complementary implementation assets and/or abilities. Likewise, each firm will likely have a different skill mix that it seeks in its employees, unique needs in its production process, and a particular set of supply constraints. In such cases, firm-specific data sets and applications of those data will differentiate the machine-learning capabilities of one firm from another (Brynjolfsson and McAfee 2017).

1.11 Conclusion

There are plenty of both optimists and pessimists about technology and growth. The optimists tend to be technologists and venture capitalists, and many are clustered in technology hubs. The pessimists tend to be economists, sociologists, statisticians, and government officials. Many of them are clustered in major state and national capitals. There is much less interaction between the two groups than within them, and it often seems as though they are talking past each other. In this chapter, we argue that in an important sense, they are.

When we talk with the optimists, we are convinced that the recent breakthroughs in AI and machine learning are real and significant. We also would argue that they form the core of a new, economically important potential GPT. When we speak with the pessimists, we are convinced that productivity growth has slowed down recently and what gains there have been are unevenly distributed, leaving many people with stagnating incomes, declining metrics of health and well-being, and good cause for concern. People are uncertain about the future, and many of the industrial titans that once dominated the employment and market value leaderboard have fallen on harder times.

These two stories are not contradictory. In fact, in many ways they are consistent and symptomatic of an economy in transition. Our analysis suggests that while the recent past has been difficult, it is not destiny. Although it is always dangerous to make predictions, and we are humble about our ability to foretell the future, our reading of the evidence does provide some cause for optimism. The breakthroughs of AI technologies already demon-

strated are not yet affecting much of the economy, but they portend bigger effects as they diffuse. More important, they enable complementary innovations that could multiply their impact. Both the AI investments and the complementary changes are costly, hard to measure, and take time to implement, and this can, at least initially, depress productivity as it is currently measured. Entrepreneurs, managers, and end-users will find powerful new applications for machines that can now learn how to recognize objects, understand human language, speak, make accurate predictions, solve problems, and interact with the world with increasing dexterity and mobility.

Further advances in the core technologies of machine learning would likely yield substantial benefits. However, our perspective suggests that an underrated area of research involves the complements to the new AI technologies, not only in areas of human capital and skills, but also new processes and business models. The intangible assets associated with the last wave of computerization were about ten times as large as the direct investments in computer hardware itself. We think it is plausible that AI-associated intangibles could be of a comparable or greater magnitude. Given the big changes in coordination and production possibilities made possible by AI, the ways that we organized work and education in the past are unlikely to remain optimal in the future.

Relatedly, we need to update our economic measurement tool kits. As AI and its complements more rapidly add to our (intangible) capital stock, traditional metrics like GDP and productivity can become more difficult to measure and interpret. Successful companies do not need large investments in factories or even computer hardware, but they do have intangible assets that are costly to replicate. The large market values associated with companies developing and/or implementing AI suggest that investors believe there is real value in those companies. In the case that claims on the assets of the firm are publicly traded and markets are efficient, the financial market will properly value the firm as the present value of its risk-adjusted discounted cash flows. This can provide an estimate of the value of both the tangible and intangible assets owned by the firm. What's more, the effects on living standards may be even larger than the benefits that investors hope to capture. It is also possible, even likely, that many people will not share in those benefits. Economists are well positioned to contribute to a research agenda of documenting and understanding the often intangible changes associated with AI and its broader economic implications.

Realizing the benefits of AI is far from automatic. It will require effort and entrepreneurship to develop the needed complements, and adaptability at the individual, organizational, and societal levels to undertake the associated restructuring. Theory predicts that the winners will be those with the lowest adjustment costs and that put as many of the right complements in place as possible. This is partly a matter of good fortune, but with the right road map, it is also something for which they, and all of us, can prepare.

Appendix

Table 1A.1 Regressions with Newey-West standard errors with longer time dependence

	(1) 1 lag allowed	(2) 2 lags allowed	(3) 3 lags allowed	(4) 4 lags allowed	(5) 10 lags allowed
Newey-West regressions, ten-year average, labor productivity growth					
Previous ten-year average productivity growth	0.0857 (0.177)	0.0857 (0.207)	0.0857 (0.227)	0.0857 (0.242)	0.0857 (0.278)
Constant	1.949*** (0.398)	1.949*** (0.465)	1.949*** (0.511)	1.949*** (0.545)	1.949*** (0.624)
Observations	50	50	50	50	50
R-squared	0.009	0.009	0.009	0.009	0.009
Newey-West regressions, ten-year average, TFP growth					
Previous ten-year average TFP growth	0.136 (0.158)	0.136 (0.181)	0.136 (0.197)	0.136 (0.208)	0.136 (0.233)
Constant	0.911*** (0.188)	0.911*** (0.216)	0.911*** (0.233)	0.911*** (0.244)	0.911*** (0.257)
Observations	50	50	50	50	50
R-squared	0.023	0.023	0.023	0.023	0.023
Newey-West regressions, ten-year average, TFP (util. adj.) growth					
Previous ten-year average TFPua growth	0.158 (0.187)	0.158 (0.221)	0.158 (0.246)	0.158 (0.266)	0.158 (0.311)
Constant	0.910*** (0.259)	0.910*** (0.306)	0.910*** (0.341)	0.910*** (0.368)	0.910*** (0.412)
Observations	50	50	50	50	50
R-squared	0.030	0.030	0.030	0.030	0.030

Note: Standard errors in parentheses.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Table 1A.2 Parameters for the toy economy J-curve

Time	Net investment	Net capital stock	Investment growth rate	Capital stock growth rate	Output
0.0	1.0	10.0			10,000.0
1.0	15.0	25.0	14.0	1.5	10,500.0
2.0	80.0	105.0	4.3	3.2	11,025.0
3.0	160.0	265.0	1.0	1.5	11,576.3
4.0	220.0	485.0	0.4	0.8	12,155.1
5.0	250.0	735.0	0.1	0.5	12,762.8
6.0	220.0	955.0	-0.1	0.3	13,401.0
7.0	140.0	1,095.0	-0.4	0.1	14,071.0
8.0	100.0	1,195.0	-0.3	0.1	14,774.6
9.0	50.0	1,245.0	-0.5	0.0	15,513.3
10.0	20.0	1,265.0	-0.6	0.0	16,288.9
11.0	10.0	1,275.0	-0.5	0.0	17,103.4
12.0	0.0	1,275.0	-1.0	0.0	17,958.6

References

- Abramovitz, Moses. 1956. "Resource and Output Trends in the U.S. Since 1870." *American Economic Review, Papers and Proceedings* 46 (2): 5–23.
- Acemoglu, D., and P. Restrepo. 2017. "The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment." NBER Working Paper no. 22252, Cambridge, MA.
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2017. "What to Expect from Artificial Intelligence." *Sloan Management Review*, Feb. 7. <https://sloanreview.mit.edu/article/what-to-expect-from-artificial-intelligence/>.
- Alloway, Tracy. 2015. "Goldman: How 'Grand Theft Auto' Explains One of the Biggest Mysteries of the U.S. Economy." *Bloomberg Business*, May 26. <http://www.bloomberg.com/news/articles/2015-05-26/goldman-how-grand-theft-auto-explains-one-of-the-biggest-mysteries-of-the-u-s-economy>.
- Alon, Titan, David Berger, Robert Dent, and Benjamin Pugsley. 2017. "Older and Slower: The Startup Deficit's Lasting Effects on Aggregate Productivity Growth." NBER Working Paper no. 23875, Cambridge, MA.
- Andrews, Dan, Chiara Criscuolo, and Peter Gal. 2016. "The Best *versus* the Rest: The Global Productivity Slowdown, Divergence across Firms and the Role of Public Policy." OECD Productivity Working Papers no. 5, Paris, OECD Publishing, Dec. 2.
- Aral, Sinan, Erik Brynjolfsson, and Lynn Wu. 2012. "Three-Way Complementarities: Performance Pay, HR Analytics and Information Technology." *Management Science* 58 (5): 913–31.
- Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, edited by Richard R. Nelson, 609–26. Princeton, NJ: Princeton University Press.
- Atkeson, Andrew, and Patrick J. Kehoe. 2007. "Modeling the Transition to a New

- Economy: Lessons from Two Technological Revolutions.” *American Economic Review* 97 (1): 64–88.
- Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2017. “Concentrating on the Fall of the Labor Share.” *American Economic Review, Papers and Proceedings* 107 (5): 180–85.
- Autor, David, and Anna Salomons. 2017. “Robocalypse Now—Does Productivity Growth Threaten Employment?” European Central Bank Conference Proceedings, June.
- Baumol, William, and William Bowen. 1966. *Performing Arts, The Economic Dilemma: A Study of Problems Common to Theater, Opera, Music, and Dance*. New York: Twentieth Century Fund.
- Bessen, James E. 2017. “AI and Jobs: The Role of Demand.” Law and Economics Research Paper no. 17-46, Boston University School of Law.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2017. “Are Ideas Getting Harder to Find?” NBER Working Paper no. 23782, Cambridge, MA.
- Bresnahan, Timothy, Erik Brynjolfsson, and Lorin Hitt. 2002. “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence.” *Quarterly Journal of Economics* 117 (1): 339–76.
- Bresnahan, Timothy F., and Manuel Trajtenberg. 1996. “General Purpose Technologies: ‘Engines of Growth?’” *Journal of Econometrics, Annals of Econometrics* 65 (1): 83–108.
- Bridgman, B., A. Dugan, M. Lal, M. Osborne, and S. Villones. 2012. “Accounting for Household Production in the National Accounts, 1965–2010.” *Survey of Current Business* 92 (5): 23–36.
- Brynjolfsson, Erik. 1993. “The Productivity Paradox of Information Technology.” *Communications of the ACM* 36 (12): 66–77.
- Brynjolfsson, Erik, and Lorin Hitt. 2000. “Beyond Computation: Information Technology, Organizational Transformation and Business Performance.” *Journal of Economic Perspectives* 14 (4): 23–48.
- . 2003. “Computing Productivity: Firm-Level Evidence.” *Review of Economics and Statistics* 85 (4): 793–808.
- Brynjolfsson, Erik, Lorin Hitt, and Shinkyu Yang. 2002. “Intangible Assets: Computers and Organizational Capital.” *Brookings Papers on Economic Activity* 2002 (1). Brookings Institution. <https://www.brookings.edu/bpea-articles/intangible-assets-computers-and-organizational-capital/>.
- Brynjolfsson, Erik, and Andrew McAfee. 2011. *Race against the Machine*. Lexington, MA: Digital Frontier.
- . 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. New York: W. W. Norton & Company.
- . 2017. “What’s Driving the Machine Learning Explosion?” *Harvard Business Review* 18:3–11. July.
- Brynjolfsson, Erik, Andrew McAfee, Michael Sorell, and Feng Zhu. 2008. “Scale without Mass: Business Process Replication and Industry Dynamics.” HBS Working Paper no. 07–016, Harvard Business School. <https://hbswk.hbs.edu/item/scale-without-mass-business-process-replication-and-industry-dynamics>.
- Brynjolfsson, Erik, and Tom Mitchell. 2017. “What Can Machine Learning Do? Workforce Implications.” *Science* 358 (6370): 1530–34.
- Brynjolfsson, Erik, and Michael D. Smith. 2000. “Frictionless Commerce? A Comparison of Internet and Conventional Retailers.” *Management Science* 46 (4): 563–85.
- Byrne, David M., John G. Fernald, and Marshall B. Reinsdorf. 2016. “Does the

- United States Have a Productivity Slowdown or a Measurement Problem?" *Brookings Papers on Economic Activity* Spring:109–82.
- Cardarelli, Roberto, and Lusine Lusinyan. 2015. "U.S. Total Factor Productivity Slowdown: Evidence from the U.S. States." IMF Working Paper no. WP/15/116, International Monetary Fund.
- Case, Anne, and Angus Deaton. 2017. "Mortality and Morbidity in the 21st Century." *Brookings Papers on Economic Activity* Spring: 397–476.
- Cette, Gilbert, John G. Fernald, and Benoit Mojon. 2016. "The Pre-Great Recession Slowdown in Productivity." *European Economic Review* 88:3–20.
- Conference Board. 2016. "The Conference Board Total Economy Database: Summary Tables (November 2016)." New York, The Conference Board.
- Congressional Budget Office (CBO). 2016. *The Budget and Economic Outlook: 2016 to 2026*. <https://www.cbo.gov/publication/51129>.
- . 2017. *The 2017 Long-Term Budget Outlook*. <https://www.cbo.gov/publication/152480>.
- Cowen, Tyler. 2011. *The Great Stagnation: How America Ate All the Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better*. New York: Dutton.
- David, Paul. 1991. "Computer and Dynamo: The Modern Productivity Paradox in a Not-Too-Distant Mirror." In *Technology and Productivity: The Challenge for Economic Policy*, 315–47. Paris: OECD Publishing.
- David, Paul A., and Gavin Wright. 2006. "General Purpose Technologies and Surges in Productivity: Historical Reflections on the Future of the ICT Revolution." In *The Economic Future in Historical Perspective*, vol. 13, edited by Paul A. David and Mark Thomas. Oxford: Oxford University Press.
- De Loecker, Jan, and Jan Eeckhout. 2017. "The Rise of Market Power and the Macroeconomic Implications." NBER Working Paper no. 23687, Cambridge, MA.
- Esteva, A., B. Kuprel, R. A. Novoa, J. Ko, S. M. Swetter, H. M. Blau, and S. Thrun. 2017. "Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks." *Nature* 542 (7639): 115–18.
- Feldstein, Martin. 2015. "The U.S. Underestimates Growth." *Wall Street Journal*, May 18.
- Fernald, John G. 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." FRBSF Working Paper no. 2012–19, Federal Reserve Bank of San Francisco. Updated March 2014. <https://www.frbsf.org/economic-research/files/wp12-19bk.pdf>.
- Fortunato, Meire, Mohammad Gheshlaghi Azar, Bilal Piot, Jacob Menick, Ian Osband, Alex Graves, Vlad Mnih, et al. 2017. "Noisy Networks for Exploration." arXiv preprint arXiv:1706.10295. <https://arxiv.org/abs/1706.10295>.
- Furman, Jason, and Peter Orszag. 2015. "A Firm-Level Perspective on the Role of Rents in the Rise in Inequality." Presentation at A Just Society Centennial Event in Honor of Joseph Stiglitz at Columbia University, Oct. 16.
- Gehring, J., M. Auli, D. Grangier, D. Yarats, and Y. N. Dauphin. 2017. "Convolutional Sequence to Sequence Learning." arXiv preprint arXiv:1705.03122. <https://arxiv.org/abs/1705.03122>.
- Gordon, Robert J. 2014. "The Demise of US Economic Growth: Restatement, Rebuttal, and Reflections." NBER Working Paper no. 19895, Cambridge, MA.
- . 2015. *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*. Princeton, NJ: Princeton University Press.
- Gutiérrez, Germán, and Thomas Philippon. 2017. "Declining Competition and Investment in the U.S." NBER Working Paper no. 23583, Cambridge, MA.

- Guvenen, Fatih, Raymond J. Mataloni Jr., Dylan G. Rassier, and Kim J. Ruhl. 2017. "Offshore Profit Shifting and Domestic Productivity Measurement." NBER Working Paper no. 23324, Cambridge, MA.
- Hall, Robert E. 2000. "E-Capital: The Link between the Stock Market and the Labor Market in the 1990s." *Brookings Papers on Economic Activity* Fall:73–118.
- Hatzius, Jan, and Kris Dawsey. 2015. "Doing the Sums on Productivity Paradox v2.0." *Goldman Sachs U.S. Economics Analyst*, no. 15/30.
- Henderson, Rebecca. 1993. "Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Industry." *RAND Journal of Economics* 24 (2): 248–70.
- . 2006. "The Innovator's Dilemma as a Problem of Organizational Competence." *Journal of Product Innovation Management* 23:5–11.
- Holmes, Thomas J., David K. Levine, and James A. Schmitz. 2012. "Monopoly and the Incentive to Innovate When Adoption Involves Switchover Disruptions." *American Economic Journal: Microeconomics* 4 (3): 1–33.
- Hortaçsu, Ali, and Chad Syverson. 2015. "The Ongoing Evolution of US Retail: A Format Tug-of-War." *Journal of Economic Perspectives* 29 (4): 89–112.
- Jones, C. I., and P. M. Romer. 2010. "The New Kaldor Facts: Ideas, Institutions, Population, and Human Capital." *American Economic Journal: Macroeconomics* 2 (1): 224–45.
- Jovanovic, Boyan, and Peter L. Rousseau. 2005. "General Purpose Technologies." In *Handbook of Economic Growth*, vol. 1B, edited by Philippe Aghion and Steven N. Durlauf, 1181–224. Amsterdam: Elsevier B.V.
- Kendrick, John W. 1961. *Productivity Trends in the United States*. National Bureau of Economic Research. Princeton, NJ: Princeton University Press.
- Levine, S., C. Finn, T. Darrell, and P. Abbeel. 2016. "End-to-End Training of Deep Visuomotor Policies." *Journal of Machine Learning Research* 17 (39): 1–40.
- Levitt, Steven D., John A. List, and Chad Syverson. 2013. "Toward an Understanding of Learning by Doing: Evidence from an Automobile Plant." *Journal of Political Economy* 121 (4): 643–81.
- Levy, Frank. 2018. "Computers and Populism: Artificial Intelligence, Jobs, and Politics in the Near Term." *Oxford Review of Economic Policy* 34 (3): 393–417.
- Liu, Y., A. Gupta, P. Abbeel, and S. Levine. 2017. "Imitation from Observation: Learning to Imitate Behaviors from Raw Video via Context Translation." arXiv preprint arXiv:1707.03374. <https://arxiv.org/abs/1707.03374>.
- Manyika, James, Michael Chui, Mehdi Miremadi, Jacques Bughin, Katy George, Paul Willmott, and Martin Dewhurst. 2017. "Harnessing Automation for a Future That Works." *McKinsey Global Institute*, January. <https://www.mckinsey.com/global-themes/digital-disruption/harnessing-automation-for-a-future-that-works>.
- McAfee, Andrew, and Erik Brynjolfsson. 2008. "Investing in the IT that Makes a Competitive Difference." *Harvard Business Review* July:98.
- Milgrom, P., and J. Roberts. 1996. "The LeChatelier Principle." *American Economic Review* 86 (1): 173–79.
- Minsky, Marvin. 1967. *Computation: Finite and Infinite Machines*. Upper Saddle River, NJ: Prentice-Hall.
- Mokyr, J. 2014. "Secular Stagnation? Not in Your Life." *Geneva Reports on the World Economy* August:83–89.
- Morris, David Z. 2016. "Today's Cars Are Parked 95 Percent of the Time." *Fortune*, Mar. 13.
- Nakamura, Leonard, and Rachel Soloveichik. 2015. "Capturing the Productivity

- Impact of the 'Free' Apps and Other Online Media." FRBP Working Paper no. 15–25, Federal Reserve Bank of Philadelphia.
- Nordhaus, W. D. 2015. "Are We Approaching an Economic Singularity? Information Technology and the Future of Economic Growth." NBER Working Paper no. 21547, Cambridge, MA.
- Organisation for Economic Co-operation and Development (OECD). 2015. *The Future of Productivity*. <https://www.oecd.org/eeco/growth/OECD-2015-The-future-of-productivity-book.pdf>.
- Orlikowski, W. J. 1996. "Improvising Organizational Transformation over Time: A Situated Change Perspective." *Information Systems Research* 7 (1): 63–92.
- Pratt, Gill A. 2015. "Is a Cambrian Explosion Coming for Robotics?" *Journal of Economic Perspectives* 29 (3): 51–60.
- Saon, G., G. Kurata, T. Sercu, K. Audhkhasi, S. Thomas, D. Dimitriadis, X. Cui, et al. 2017. "English Conversational Telephone Speech Recognition by Humans and Machines." arXiv preprint arXiv:1703.02136. <https://arxiv.org/abs/1703.02136>.
- Smith, Noah. 2015. "The Internet's Hidden Wealth." *Bloomberg View*, June 6. <http://www.bloombergview.com/articles/2015-06-10/wealth-created-by-the-internet-may-not-appear-in-gdp>.
- Solow, Robert M. 1957. "Technical Change and the Aggregate Production Function." *Review of Economics and Statistics* 39 (3): 312–20.
- . 1987. "We'd Better Watch Out." *New York Times Book Review*, July 12, 36.
- Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom, and Till von Wachter. 2015. "Firming Up Inequality." NBER Working Paper no. 21199, Cambridge, MA.
- Stiglitz, Joseph E. 2014. "Unemployment and Innovation." NBER Working Paper no. 20670, Cambridge, MA.
- Syverson, Chad. 2013. "Will History Repeat Itself? Comments on 'Is the Information Technology Revolution Over?'" *International Productivity Monitor* 25:37–40.
- . 2017. "Challenges to Mismeasurement Explanations for the US Productivity Slowdown." *Journal of Economic Perspectives* 31 (2): 165–86.
- Yang, Shinkyu, and Erik Brynjolfsson. 2001. "Intangible Assets and Growth Accounting: Evidence from Computer Investments." Unpublished manuscript, Massachusetts Institute of Technology.

Comment Rebecca Henderson

"Artificial Intelligence and the Modern Productivity Paradox" is a fabulous chapter. It is beautifully written, extremely interesting, and goes right to the heart of a centrally important question, namely, what effects will AI have on economic growth? The authors make two central claims. The first is that AI

Rebecca Henderson is the John and Natty McArthur University Professor at Harvard University, where she has a joint appointment at the Harvard Business School in the General Management and Strategy units, and a research associate of the National Bureau of Economic Research.

For acknowledgments, sources of research support, and disclosure of the author's material financial relationships, if any, please see <http://www.nber.org/chapters/c14020.ack>.

is a general purpose technology, or GPT, and as such is likely to have a dramatic impact on productivity and economic growth. The second is that the reason we do not yet see it in the productivity statistics is because—like all GPTs—this is a technology that will take time to diffuse across the economy.

More specifically, the authors argue that AI will take time to diffuse because its adoption will require mastering “adjustment costs, organizational changes, and new skills.” They suggest that just as we did not see IT in the productivity statistics until firms had made the organizational changes and hired the human capital necessary to master it, so the adoption of AI will require not only the diffusion of the technology itself but also the development of the organizational and human assets that will be required to exploit its full potential.

This is a fascinating idea. One of the reasons I like the chapter so much is that takes seriously an idea that economists long resisted—namely, that things as nebulous as “culture” and “organizational capabilities” might be (a) very important, (b) expensive, and (c) hard to change. Twenty-five years ago, when I submitted a paper to the *RAND Journal of Economics* that suggested that incumbents were fundamentally disadvantaged compared to entrants because they were constrained by old ways of acting and perceiving, I got a letter from the editor that began “Dear Rebecca, you have written a paper suggesting that the moon is made of green cheese, and that economists have too little considered the motions of cheesy planetoids”

I like to think that few editors would respond that way today. Thanks to a wave of new work in organizational economics and the pioneering empirical research of scholars like Nick Bloom, John van Reenen, Raffaella Sadun, and the authors themselves, we now have good reason to believe that managerial processes and organizational structures have very real effects on performance and that they take a significant time to change. One of the most exciting things about this chapter is that it takes these ideas sufficiently seriously to suggest that the current slowdown in productivity is largely a function of organizational inertia—that a central macroeconomic outcome is a function of a phenomenon that thirty years ago was barely on the radar.

That’s exciting. Is it true? And if it is, what are its implications?

My guess is that the deployment of AI will indeed be gated by the need to change organizational structures and processes. But I think that the authors may be underestimating the implications of this dynamic in important ways.

Take the case of accounting. A few months ago, I happened to meet the chief strategy officer for one of the world’s largest accounting firms. He told me that his firm is the largest hirer of college graduates in the world—which may or may not be true, but which he certainly believed—and that his firm was planning to reduce the number of college graduates they hire by 75 percent over the next four to five years—largely because it is increasingly clear that AI is going to be able to take over much of the auditing work currently performed by humans. This shift will certainly be mediated by

every accounting firm's ability to integrate AI into their procedures and to persuade their customers that it is worth paying for—examples of exactly the kinds of barriers that this chapter suggests are so important—but in principle it should dramatically increase the productivity of accounting services, exactly the effects that Erik and his coauthors are hoping for.

But I am worried about all the college graduates the accounting firms are not going to hire. More broadly, as AI begins to diffuse across the economy it seems likely that a lot of people will get pushed into new positions and a lot of people will be laid off. And just as changing organizational processes takes time, so it's going to take time to remake the social context in ways that will make it possible to handle these dislocations. Without these kinds of investments—one can imagine they might be in education, in relocation assistance, and the like—there is a real risk of a public backlash against AI that could dramatically reduce its diffusion rate.

For example, the authors are excited about the benefits that the widespread diffusion of autonomous vehicles are likely to bring. Productivity seems likely to skyrocket, while with luck tens of thousands of people will no longer perish in car crashes every year. But “driving” is one of the largest occupations there is. What will happen when millions of people begin to be laid off? I'm with the authors in believing that the diffusion of AI could be an enormous source of innovation and growth. But I can see challenges in the transition at the societal level, as well as at the organizational level. And there will also be challenges if too large a share of the economic gains from the initial deployment of the technology goes to the owners of capital rather than to the rest of society.

Which is to say that I am a little more pessimistic than Erik and his coauthors as to the speed at which AI will diffuse—and this is even before I start talking about the issues that Scott, Iain, and I touch on in our own chapter, namely, that we are likely to have significant underinvestment in AI relative to the social option, coupled with a fair amount of dissipative racing.

