



National
Bureau of
Economic
Research

THE ECONOMICS OF ARTIFICIAL INTELLIGENCE

An Agenda

Edited by Ajay Agrawal,
Joshua Gans, and Avi Goldfarb




Artificial Intelligence and the Modern Product...

The Economics of Artificial Intelligence



**National Bureau of
Economic Research
Conference Report**



The Economics of 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://lccn.loc.gov/2018037552>

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

National Bureau of Economic Research

Officers

Karen N. Horn, <i>chair</i>	Kelly Horak, <i>controller and assistant corporate secretary</i>
John Lipsky, <i>vice chair</i>	Alterra Milone, <i>corporate secretary</i>
James M. Poterba, <i>president and chief executive officer</i>	Denis Healy, <i>assistant corporate secretary</i>
Robert Mednick, <i>treasurer</i>	

Directors at Large

Peter C. Aldrich	Diana Farrell	Michael H. Moskow
Elizabeth E. Bailey	Jacob A. Frenkel	Alicia H. Munnell
John H. Biggs	Robert S. Hamada	Robert T. Parry
John S. Clarkeson	Peter Blair Henry	James M. Poterba
Kathleen B. Cooper	Karen N. Horn	John S. Reed
Charles H. Dallara	Lisa Jordan	Marina v. N. Whitman
George C. Eads	John Lipsky	Martin B. Zimmerman
Jessica P. Einhorn	Laurence H. Meyer	
Mohamed El-Erian	Karen Mills	

Directors by University Appointment

Timothy Bresnahan, <i>Stanford</i>	George Mailath, <i>Pennsylvania</i>
Pierre-André Chiappori, <i>Columbia</i>	Marjorie B. McElroy, <i>Duke</i>
Alan V. Deardorff, <i>Michigan</i>	Joel Mokyr, <i>Northwestern</i>
Edward Foster, <i>Minnesota</i>	Cecilia Rouse, <i>Princeton</i>
John P. Gould, <i>Chicago</i>	Richard L. Schmalensee, <i>Massachusetts Institute of Technology</i>
Mark Grinblatt, <i>California, Los Angeles</i>	Ingo Walter, <i>New York</i>
Bruce Hansen, <i>Wisconsin-Madison</i>	David B. Yoffie, <i>Harvard</i>
Benjamin Hermalin, <i>California, Berkeley</i>	
Samuel Kortum, <i>Yale</i>	

Directors by Appointment of Other Organizations

Jean-Paul Chavas, <i>Agricultural and Applied Economics Association</i>	Robert Mednick, <i>American Institute of Certified Public Accountants</i>
Martin J. Gruber, <i>American Finance Association</i>	Peter L. Rousseau, <i>American Economic Association</i>
Philip Hoffman, <i>Economic History Association</i>	Gregor W. Smith, <i>Canadian Economics Association</i>
Arthur Kennickell, <i>American Statistical Association</i>	William Spriggs, <i>American Federation of Labor and Congress of Industrial Organizations</i>
Jack Kleinhenz, <i>National Association for Business Economics</i>	Bart van Ark, <i>The Conference Board</i>

Directors Emeriti

George Akerlof	Franklin Fisher	John J. Siegfried
Jagdish Bhagwati	Saul H. Hymans	Craig Swan
Don R. Conlan	Rudolph A. Oswald	
Ray C. Fair	Andrew Postlewaite	

Relation of the Directors to the Work and Publications of the National Bureau of Economic Research

1. The object of the NBER is to ascertain and present to the economics profession, and to the public more generally, important economic facts and their interpretation in a scientific manner without policy recommendations. The Board of Directors is charged with the responsibility of ensuring that the work of the NBER is carried on in strict conformity with this object.

2. The President shall establish an internal review process to ensure that book manuscripts proposed for publication DO NOT contain policy recommendations. This shall apply both to the proceedings of conferences and to manuscripts by a single author or by one or more co-authors but shall not apply to authors of comments at NBER conferences who are not NBER affiliates.

3. No book manuscript reporting research shall be published by the NBER until the President has sent to each member of the Board a notice that a manuscript is recommended for publication and that in the President's opinion it is suitable for publication in accordance with the above principles of the NBER. Such notification will include a table of contents and an abstract or summary of the manuscript's content, a list of contributors if applicable, and a response form for use by Directors who desire a copy of the manuscript for review. Each manuscript shall contain a summary drawing attention to the nature and treatment of the problem studied and the main conclusions reached.

4. No volume shall be published until forty-five days have elapsed from the above notification of intention to publish it. During this period a copy shall be sent to any Director requesting it, and if any Director objects to publication on the grounds that the manuscript contains policy recommendations, the objection will be presented to the author(s) or editor(s). In case of dispute, all members of the Board shall be notified, and the President shall appoint an ad hoc committee of the Board to decide the matter; thirty days additional shall be granted for this purpose.

5. The President shall present annually to the Board a report describing the internal manuscript review process, any objections made by Directors before publication or by anyone after publication, any disputes about such matters, and how they were handled.

6. Publications of the NBER issued for informational purposes concerning the work of the Bureau, or issued to inform the public of the activities at the Bureau, including but not limited to the NBER Digest and Reporter, shall be consistent with the object stated in paragraph 1. They shall contain a specific disclaimer noting that they have not passed through the review procedures required in this resolution. The Executive Committee of the Board is charged with the review of all such publications from time to time.

7. NBER working papers and manuscripts distributed on the Bureau's web site are not deemed to be publications for the purpose of this resolution, but they shall be consistent with the object stated in paragraph 1. Working papers shall contain a specific disclaimer noting that they have not passed through the review procedures required in this resolution. The NBER's web site shall contain a similar disclaimer. The President shall establish an internal review process to ensure that the working papers and the web site do not contain policy recommendations, and shall report annually to the Board on this process and any concerns raised in connection with it.

8. Unless otherwise determined by the Board or exempted by the terms of paragraphs 6 and 7, a copy of this resolution shall be printed in each NBER publication as described in paragraph 2 above.

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.

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 Tratjenberg, 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. Cetty, 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/benedict-evans/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.

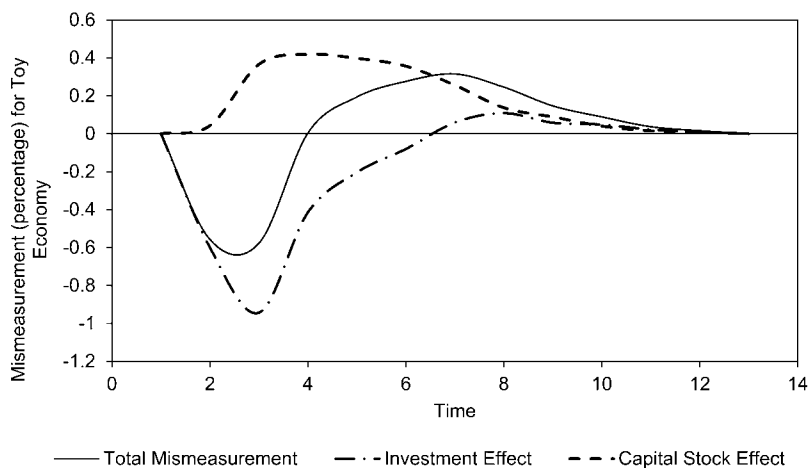


Fig. 1.9 The mismeasurement J-curve for an economy accumulating a new kind of capital

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 TFP _{Pua} 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/52480>.
- 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.
- Syerson, 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.

Contributors

Daron Acemoglu
Department of Economics
Massachusetts Institute of
Technology
50 Memorial Drive
Cambridge, MA 02142-1347

Philippe Aghion
Collège de France
3 Rue D'Ulm
75005 Paris, France

Ajay Agrawal
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Susan Athey
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305

James Bessen
Technology & Policy Research
Initiative
Boston University School of Law
765 Commonwealth Avenue
Boston, MA 02215

Erik Brynjolfsson
MIT Sloan School of Management
100 Main Street, E62-414
Cambridge, MA 02142

Colin F. Camerer
Department of Economics
California Institute of Technology
1200 East California Boulevard
Pasadena, CA 91125

Judith Chevalier
Yale School of Management
135 Prospect Street
New Haven, CT 06520

Iain M. Cockburn
School of Management
Boston University
595 Commonwealth Avenue
Boston, MA 02215

Tyler Cowen
Department of Economics
George Mason University
4400 University Drive
Fairfax, VA 22030

Patrick Francois
Vancouver School of Economics
University of British Columbia
IONA Building, 6000 Iona Drive
Vancouver, BC V6T 2E8, Canada

Jason Furman
Harvard Kennedy School
79 John F. Kennedy Street
Cambridge, MA 02138

Alberto Galasso
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Joshua Gans
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Avi Goldfarb
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Austan Goolsbee
University of Chicago Booth School
of Business
5807 S. Woodlawn Avenue
Chicago, IL 60637

Rebecca Henderson
Harvard Business School
Morgan Hall 445
Soldiers Field Road
Boston, MA 02163

Ginger Zhe Jin
Department of Economics
University of Maryland
3115F Tydings Hall
College Park, MD 20742-7211

Benjamin F. Jones
Department of Management and
Strategy
Kellogg School of Management
Northwestern University
2211 Campus Drive
Evanston, IL 60208

Charles I. Jones
Graduate School of Business
Stanford University
655 Knight Way
Stanford, CA 94305-4800

Daniel Kahneman
Woodrow Wilson School
Princeton University
Princeton, NJ 08544-1013

Anton Korinek
University of Virginia
Monroe Hall 246
248 McCormick Road
Charlottesville, VA 22904

Mara Lederman
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, Ontario M5S 3E6, Canada

Hong Luo
Harvard Business School
Morgan Hall 241
Soldiers Field Road
Boston, MA 02163

John McHale
108 Cairnes Building
School of Business and Economics
National University of Ireland
Galway H91 TK33, Ireland

Paul R. Milgrom
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305-6072

Matthew Mitchell
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Alexander Oettl
Scheller College of Business
Georgia Institute of Technology
800 West Peachtree Street NW
Atlanta, GA 30308

Andrea Prat
Columbia Business School
Uris Hall 624
3022 Broadway
New York, NY 10027-6902

Manav Raj
Stern School of Business
New York University
44 West Fourth Street
New York, NY 10012

Pascual Restrepo
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215

Daniel Rock
MIT Sloan School of Management
100 Main Street, E62-365
Cambridge, MA 02142

Jeffrey D. Sachs
Center for Sustainable Development,
Earth Institute
Columbia University
535 West 116th Street, MC 4327
New York, NY 10027

Robert Seamans
Stern School of Business
New York University
44 West 4th Street, KMC 7-58
New York, NY 10012

Scott Stern
MIT Sloan School of Management
100 Main Street, E62-476
Cambridge, MA 02142

Betsy Stevenson
Gerald R. Ford School of Public Policy
University of Michigan
5224 Weill Hall
735 South State Street
Ann Arbor, MI 48109-3091

Joseph E. Stiglitz
Columbia University
Uris Hall 212
3022 Broadway
New York, NY 10027

Chad Syverson
University of Chicago Booth School of
Business
5807 S. Woodlawn Avenue
Chicago, IL 60637

Matt Taddy
University of Chicago Booth School of
Business
5807 S. Woodlawn Avenue
Chicago, IL 60637

Steven Tadelis
Haas School of Business
University of California, Berkeley
545 Student Services Building
Berkeley, CA 94720

Manuel Trajtenberg
Eitan Berglas School of Economics
Tel Aviv University
Tel Aviv 69978, Israel

Daniel Trefler
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6, Canada

Catherine Tucker
MIT Sloan School of Management
100 Main Street, E62-536
Cambridge, MA 02142

Hal Varian
School of Information
University of California, Berkeley
102 South Hall
Berkeley, CA 94720-4600

Author Index

- Abadie, A., 531
Ablon, L., 448
Abrahamson, Z., 420
Abramovitz, M., 32
Acemoglu, D., 23n1, 43, 89, 90, 105, 127,
141, 152, 197, 198, 201, 202, 203,
203n4, 204, 204n5, 205, 206, 208, 210,
211, 212n7, 219, 220, 220n11, 223, 224,
225n14, 238, 240, 243, 271, 283, 293n3,
376n27, 376n28, 554
Acquisti, A., 410, 416, 424, 440n2, 444, 448,
451, 457, 457n42, 459, 483n19
Adee, S., 426
Agarwal, A., 82n12, 83
Aghion, P., 122, 172, 262, 262n19, 263, 265,
267, 268, 373n23, 465, 477, 479, 495
Agrawal, A., 5, 39, 90, 97n8, 150, 161, 167,
241, 425, 464, 501
Aguilar, M., 386n36
Airoldi, E., 80
Akerlof, G., 378
Akerman, A., 7, 303
Alexopoulos, M., 555
Allen, R. C., 204, 209
Alloway, T., 29
Alon, T., 43
Alpaydin, E., 92
Altman, S., 325
Alvarez-Cuadrado, F., 241n4
Anderson, R., 406, 416
Andrade, E. B., 592n3
Andrews, D., 30
Angermueller, C., 168
Aral, S., 43
Arntz, M., 321
Aronoff, M., 497
Arrow, K., 9, 43, 110, 118, 148, 364n11, 366,
412
Arthur, B. W., 150
Asher, S., 525
Athey, S., 68, 425, 448, 449, 510, 514, 515,
516, 517, 519, 523, 524, 525, 527, 528,
529, 530, 531, 532, 533, 534, 536, 538
Atkeson, A., 42n19
Autor, D. H., 7, 23n1, 30, 89, 198, 202n3,
203, 208, 220n11, 238, 239n3, 240, 271,
309, 322, 475, 555
Axelrod, R., 413
Ayres, R., 201
Azoulay, P., 475, 477

Babcock, L., 592n5
Bai, J., 532
Baker, D., 366n14
Baker, G., 107
Banbura, M., 537
Barkai, S., 271
Barrat, J., 350
Baslandze, S., 263, 264
Bastani, H., 529
Baumol, W., 38n14, 238
Bayati, M., 529, 531, 532, 536
Belloni, A., 93, 522
Bengio, S., 401

- Bengio, Y., 71, 74, 75, 75n6, 149, 168, 168n6, 172
Benzell, S. G., 90, 336
Berg, A., 373n24
Beron, K. J., 483
Bertocchi, G., 430
Bessen, J. E., 23n1, 203, 300, 413, 555
Bhatt, M. A., 592n4
Bickel, P. J., 522
Binmore, K., 589, 589n1, 593
Bjorkegren, D., 516, 519
Blake, T., 582
Blattberg, R. C., 598
Blei, D. M., 507, 510, 515, 532, 533, 534
Bloom, N., 27, 150, 223, 259n15, 268, 560
Bolton, P., 92, 95, 95n5, 96n7, 100, 112
Boppart, T., 241n4, 295
Borenstein, S., 413
Bornmann, L., 153
Bostrom, N., 3, 286, 381, 382, 382n31
Bottou, L., 79
Bousquet, O., 79
Bowen, W., 38n14
Brandeis, L. D., 431, 459
Brander, J. A., 473
Brandimarte, L., 448
Brandt, L., 471
Bresnahan, T. F., 4, 39, 42, 116, 119, 120, 169, 176n2, 310
Bridgman, B., 38
Brooks, R., 124
Broseta, B., 593n9, 594n9
Brown, J., 314
Brunskill, E., 528
Brynjolfsson, E., 23, 23n1, 28n7, 30, 39, 40, 42, 43, 47, 50, 89, 119, 120, 150, 197, 201, 204, 309, 555, 557, 560, 563
Brzeski, C., 556
Buera, F. J., 293, 293n2
Buffie, E. F., 373n24
Bughin, J., 408
Burk, I., 556
Busch, M. L., 474, 475n13
Byers, J. W., 577n5
Byrne, D. M., 29, 319

Camerer, C. F., 589, 589n1, 590, 593, 594n10, 595, 598
Campbell, K., 450
Cardarelli, R., 29
Carley, M., 130, 132
Case, A., 26
Catalini, C., 150, 425, 448, 449, 454
Cavusoglu, H., 450
Cette, G., 27
Chandler, A. D., 206
Chapelle, O., 529
Chapman, J. P., 599
Chapman, L. J., 599
Chernozhukov, V., 93, 522, 523, 527
Chetty, R., 270
Chevalier, J., 577n5
Chiou, L., 428, 448
Chong, J.-K., 593, 595
Christie, W. G., 414
Chui, M., 331
Clark, C., 295
Cockburn, I., 150
Coey, D., 514
Cohen, J., 555
Cohen, L., 501
Comin, D., 241n4, 295
Corrigan, B., 597
Costa-Gomes, M. A., 593, 593n9, 594n9
Courville, A., 71, 75n6
Cowen, T., 27, 150
Cowgill, B., 562
Cranor, L. F., 424, 458
Crawford, V. P., 593, 593n9, 594n9
Crisuolo, C., 30

Dana, J., 599
Danaylov, N., 253
Dasgupta, P., 366n15
Datta, A., 433
Dauth, W., 556
David, P. A., 4, 41, 42n19, 119
Davies, R. B., 483
Dávila, E., 354
Dawes, R. M., 597, 598, 599, 599n15
Dawsey, K., 29
Deaton, A., 26, 85
Della Vigna, S., 411
Delli Gatti, D., 361, 362n9, 380
De Loecker, J., 30
Dennis, B. N., 293n2
Dewatripoint, M., 112
Diamond, A., 531
Dietvorst, B. J., 426
Dimakopoulou, M., 529
Dimico, A., 430
Dobson, W., 485n23
Dogan, M., 107n9
Doi, E., 599
Domingos, P., 93
Doms, M., 562

- Dorantes, C., 450
 Dorn, D., 208
 Dosi, G., 366n14
 Doudchenko, N., 531
 Dover, Y., 577n5
 Drandakis, E., 376n27
 Dubé, J.-P., 410
 Dudik, M., 528
 Duffo, E., 539
 Dunne, T., 562
 Dupuit, J., 296, 298
 Duranton, G., 480, 567
 Dwork, C., 453
 Dzamba, M., 116

 Eckles, D., 538
 Edwards, D. D., 598
 Edwards, J. S., 598
 Edwards, L., 294
 Edwards, W., 595
 Eeckhout, J., 30
 Egami, N., 510
 Einhorn, H. J., 598, 599
 Elkin-Koren, N., 583
 Elsby, M. W. L., 270, 329
 Engel, E., 295
 Engerman, S. L., 430
 Engstrom, R., 537
 Erlingsson, U., 453
 Ethier, W. J., 476
 Etzioni, O., 410
 Evenson, R. E., 140
 Ezrachi, A., 414

 Fajgelbaum, P., 480
 Farrell, J., 424
 Faure-Grimaud, A., 92, 95, 95n5, 96n7,
 100, 112
 Faust, D., 597
 Fehr, E., 358n8
 Feldman, M. P., 560
 Feldstein, M., 29
 Feraud, R., 529
 Fernald, J. G., 27, 29, 32, 319
 Feurer, M., 70
 Filippas, A., 577n5
 Finklestein, A., 503
 Fleming, L., 161, 479
 Florêncio, D., 452
 Foellmi, R., 295
 Forbes, S., 108
 Ford, M., 201, 204, 301
 Forsythe, R., 590n2

 Fortunato, M., 25
 Fredriksson, P. G., 483
 Frey, C. B., 291, 321, 331, 350n1, 553, 555,
 556
 Friedman, J., 93
 Frosst, N., 76n7
 Fudenberg, D., 413, 594
 Fujii, H., 465
 Fung, A., 459
 Furman, J. L., 30, 122, 140, 150, 553

 Gaarder, I., 7, 303
 Gabaix, X., 604
 Gaggi, P., 303
 Gal, M. S., 583
 Gal, P., 30
 Galasso, A., 495, 496, 497, 498, 501
 Gans, J., 5, 39, 90, 97n8, 171, 425, 454, 464,
 550
 Garicano, L., 267
 Geanakoplos, J., 354, 368n19
 Geank, 369
 Gehring, J., 25
 Gentszkow, M., 411
 Gibbons, R., 107
 Glaeser, E. L., 516, 537, 548
 Goel, S., 516
 Goeree, J., 593
 Goh, G., 151n1, 168
 Goldberg, L. R., 597, 601n19
 Golden, J. M., 577n5
 Goldenshluger, A., 529
 Goldfarb, A., 5, 39, 90, 97n8, 148, 150,
 161, 167, 425, 427, 448, 464, 482, 483,
 483n19, 550, 593n8
 Goldin, C., 209, 322
 Goldman, M., 536
 Gomez-Uribe, C., 603
 Good, I. J., 238, 253
 Goodfellow, I., 66, 71, 75n6, 401
 Goolsbee, A. D., 310, 312, 314
 Gopalan, P., 510
 Gordon, R. D., 305, 310, 349
 Gordon, R. J., 27, 150, 175, 210, 223,
 259n15, 264
 Graepel, T., 429
 Graetz, G., 201, 274, 319, 554
 Graff Zivin, J. S., 475, 477
 Graham, M., 459
 Green, J., 502
 Greenstein, S., 119
 Greenwald, B., 354, 364n11, 366n15,
 368n19, 369, 370, 412

- Gregory, T., 321
Griliches, Z., 116, 120, 170
Grissen, D., 516, 519
Groover, M., 201, 205
Groshen, E. L., 350
Gross, R., 451
Grossman, G. M., 464, 476, 477, 478, 479, 480, 482
Grove, W. M., 597
Guerrieri, V., 293n3
Gurun, U., 501
Gutiérrez, G., 30
Guvenen, F., 29
- Ha, Y., 599
Hahn, J., 522
Haile, P. A., 515
Hainmueller, J., 531
Hall, B. H., 132
Hall, R. E., 47
Hansen, C., 93, 522
Hanson, G. H., 208
Hanson, R., 386n36
Harari, Y., 383
Hardt, M., 80
Harrell, E., 443, 444
Hart, O., 269
Hartford, J., 68, 75, 79n9, 526, 594
Haskel, J., 18
Hastie, T., 93, 526
Hatzius, J., 29
Haugeland, J., 69
Hawkins, J., 93
Hay, B., 494n2, 499
Hazan, E., 408
He, K., 75
Heckman, J. J., 85, 182
Heifels, A., 116
Helpman, E., 169, 464, 476, 477, 478, 479, 480
Hemous, D., 241
Hendel, I., 412
Henderson, R., 42, 43, 150
Herley, C., 452
Herrendorf, B., 204, 241n4
Hersh, J., 537
Hicks, J., 350, 368
Hillis, A., 516
Hinton, G. E., 68, 74, 76n7, 124, 125, 149, 168, 168n6, 172
Hirschman, A., 178
Hitt, L., 42, 43, 120
Ho, T.-H., 592n3, 593, 595
Hobijn, B., 270, 329
Hoch, S. J., 598
Hochreiter, S., 73, 76
Hodas, N., 151n1, 168
Hoffman, D., 444, 457
Hoffman, M., 447n17
Hofman, J. M., 510
Hogarth, R. M., 598
Holmes, T. J., 43
Holmstrom, B., 112, 269
Holt, C., 593
Hoos, H., 575
Hornik, K., 74, 525
Hortaçsu, A., 43, 593n8
Horton, J. J., 577n5
Hotborn, T., 525
Howitt, P., 122, 262n19, 263, 477
Hubbard, F., 494n2, 500
Huber, P., 494
Hunt, N., 603
Hutter, F., 575
Hutter, M., 237n1
- Imbens, G. W., 68, 517, 522, 523, 524, 525, 529, 530, 536, 538
Iriberry, N., 593, 594n10
Irwin, D. A., 477
İşçan, T. B., 293n2
- Jaffe, A. B., 132, 150
Jaravel, X., 312
Jarrell, G., 502
Jayadev, A., 366n14
Jean, N., 537
Jewitt, I., 112
Jha, S., 94, 95
Jiang, N., 528
Jin, G. Z., 442n4
Joachims, T., 528
Johnson, E. J., 589, 593, 603
Jones, B. F., 127, 150, 155, 161, 259, 259n15, 373n23
Jones, C. I., 46, 150, 152, 159, 160, 170n9, 171n9, 240, 252, 259n15, 373n23
Jones, R. W., 472
Jordan, M. I., 509, 532
Jovanovic, B., 41
- Kaboski, J. P., 293, 293n2
Kahn, L. B., 447n17
Kahneman, D., 596, 600, 604
Kaldor, N., 240
Kallus, N., 528

- Kaplan, J., 3
 Kaplow, L., 500
 Kapur, D., 150
 Karabarounis, L., 270, 329
 Karagözoğlu, E., 589
 Katz, L. F., 7, 209, 322
 Kehoe, P. J., 42
 Kehrig, M., 270
 Kendrick, J. W., 43
 Kennan, J., 590n2
 Kennedy, C., 376
 Keynes, J. M., 379
 Kislev, Y., 140
 Kitagawa, T., 527
 Klayman, J., 599
 Kleinberg, J., 15, 507, 516, 518, 548, 588, 597
 Klenow, P. J., 310, 312, 477
 Klette, T. J., 479
 Kleven, H. J., 270
 Knetsch, L., 604
 Ko, M., 450
 Kogler, D. F., 560
 Koh, D., 329
 Kohno, T., 446
 Kolb, D. A., 207
 Kollmeyer, C., 293n2
 Komarova, T., 538
 Kominers, D., 501, 516
 Kongsamut, P., 241n4, 295
 Korinek, A., 354, 365, 366n14, 371, 373, 383, 384
 Korolova, A., 453
 Kortum, S. S., 259, 259n15, 479
 Kosinski, M., 429
 Kotlikoff, L. J., 336, 344
 Krajbich, I., 592n6
 Kravitz, L., 325
 Krisiloff, M., 325
 Krizhevsky, A., 68, 74, 125
 Krueger, A. B., 7, 482
 Krugman, P. R., 464, 472, 479, 480
 Krussell, P., 265
 Künzel, S., 527
 Kurakin, A., 401
 Kurakin, P. R., 464
 Kurzweil, R., 238, 253, 253n12, 350, 373n23, 381, 382
 Kuznets, S., 205, 206

 Lada, A., 540
 Laffont, J.-J., 514
 LaGarda, G., 336
 Laibson, D., 606
 Lambrecht, A., 425, 426
 Lancaster, K., 362
 Landes, D., 206
 Lane, J., 554, 562
 Langford, J., 528
 Lanier, J., 84
 Lashkari, D., 241n4, 295
 Lawrence, R. Z., 294
 LeCun, Y., 75, 149, 168, 168n6, 172
 Lederman, M., 108, 550
 Legg, S., 237n1
 Leung, M. K. K., 121
 Levin, J., 514, 515
 Levine, D. K., 43
 Levine, S., 38
 Levinthal, D., 161
 Levitt, S. D., 43
 Levy, F., 37, 202n3, 238, 239n3
 Lewis, E., 562
 Lewis, G., 526, 527
 Lewis, M., 596
 Lewis, R. A., 541n2
 Lewis, W. A., 294
 Leyton-Brown, K., 569, 575, 594
 Li, D., 447n17, 529
 Li, L., 528, 529, 580, 581
 Liang, A., 588, 594
 Lim, K., 465, 479
 Lindor, R., 494
 Lipsey, R., 362
 List, J. A., 43
 Litan, R. E., 441n2, 494
 Liu, Y., 38
 Long, N., 241n4
 Lovallo, D., 600
 Lowenstein, G., 448, 592n5
 Lu, Y., 529
 Luo, H., 495, 496, 497, 498, 501
 Lusinyan, L., 29
 Lusted, L. B., 94

 Malthus, T. R., 384
 Mamoshina, P., 168n7
 Managi, S., 465
 Mandel, M., 555
 Manning, A., 379
 Mantoux, P., 200, 203
 Manuelli, R. E., 205, 243n6
 Manyika, J., 331
 Marchant, G., 494n2
 Marco, A., 130, 132
 Markov, J., 93

- Markusen, J. R., 476
Marthews, A., 425
Marx, M., 479
Massey, C., 426
Masterov, D. V., 578, 579
Matsuyama, K., 294, 295
Mayer, U. F., 578, 579
Mayer, W., 471
Mayzlin, D., 577n5
McAfee, A., 23, 23n1, 30, 40, 50, 89, 150, 201, 204, 309, 555
McCall, J. J., 582
McClelland, J. L., 602
McDonald, A. M., 424, 458
McElheran, K., 560
McFadden, D., 85, 514
Mcguckin, R. H., 562
McHale, J., 150, 241
McLaren, J., 469
McSweeney, T., 583
Meade, J. E., 362
Meadows, M., 420
Meehl, P. E., 596, 597
Meltz, M. J., 488
Mestieri, M., 241n4, 295
Mian, A., 208
Michaels, G., 201, 274, 319, 554
Mikolov, T., 76
Milgrom, P. R., 43, 111, 567, 569, 574, 577
Miller, A., 428, 429, 448, 483
Miller, S. M., 201
Milliment, D. L., 483
Minsky, M., 28, 208
Miremadi, M., 331
Mishel, L., 322
Misra, S., 410
Mitchell, T., 39, 557, 563
Mnih, V., 84
Mobius, M. M., 510
Mogstad, M., 7, 303
Mojon, B., 27
Mokyr, J., 29, 121, 150, 175, 201, 206, 209
Monro, S., 78
Moore, G. E., 384n35, 498
Moore, M., 496
Morris, D. Z., 37
Mortensen, D. T., 379, 582
Muellbauer, J., 85
Mullainathan, S., 169n8, 511, 512, 518, 588
Murdoch, J. C., 483
Murnane, R. J., 202n3, 238, 239n3
Murphy, K., 107
Murray, C., 325
Murray, F., 141
Mussa, M. L., 471
Mutz, R., 153
Myers, A., 130
Myerson, R. B., 568
Naecker, J., 592
Naik, N., 516, 537
Nakamura, L., 29
Nave, G., 590
Neal, R. M., 74
Neelin, J., 589n1
Neiman, B., 270, 329
Nekipelov, D., 538
Nelson, P., 581
Nelson, R., 118, 161
Newbery, D., 362
Newell, A., 123
Newhouse, D., 537
Ng, A., 93, 509, 532
Ng, S., 532
Ngai, L. R., 241n4, 294
Nickell, S., 293n2
Nielsen, M., 152, 161
Nilsson, N., 122, 207, 208
Nordhaus, W. D., 28, 152, 172, 238
North, D. C., 577
Nosko, C., 577n5, 582
O'Dea, B., 583
Odlyzko, A., 441n2
Oettle, A., 241
Olmstead, A. L., 204, 205, 206
Olsen, M., 241
O'Mahony, S., 141
O'Neil, C., 433
O'Reilly, B., 499
Orlikowski, W. J., 43
Orszag, P., 30
Osborne, M. A., 291, 321, 331, 350n1, 553, 555, 556
Osindero, S., 74
Oskamp, S., 601
Ossard, H., 514
Östling, R., 593n8
Ostrovsky, M., 568
Pajarinen, M., 556
Pál, J., 510
Palfrey, T., 593
Parchomovsky, G., 494

- Pate, R. H., 420
 Pelzman, S., 502
 Peretto, P. F., 241
 Peterson, N., 599
 Peysakhovich, A., 540, 592
 Phelps, E., 376n27
 Philippon, T., 30
 Pierce, D. G., 413
 Pihur, V., 453
 Piketty, T., 360
 Pissarides, C. A., 241n4, 294, 379
 Polemarchakis, H., 354, 368n19, 369
 Polinsky, M., 494n2
 Porter, M. E., 481, 494
 Poschke, M., 241n4
 Posner, R. A., 439
 Prantl, S., 263
 Pratt, G. A., 40
 Proserpio, D., 577n5
 Puga, D., 480

 Raghaven, M., 518
 Ramaswamy, R., 291, 293n2
 Ramlogan, R., 480
 Rao, J. M., 516, 536, 541n2
 Rasmussen, W. D., 200, 203, 205n6, 206
 Rauch, J. E., 469
 Rebelo, S., 241n4, 295
 Recht, B., 80
 Redding, S., 293n2
 Reinsdorf, M. B., 29
 Rennie, J., 80
 Restrepo, P., 23n1, 43, 90, 105, 127, 152, 197,
 201, 202, 203, 203n4, 204, 204n5, 206,
 208, 210, 211, 212n7, 219, 220, 220n11,
 223, 224, 225n14, 238, 240, 241, 243,
 271, 283, 376n27, 379, 554
 Rhode, P. W., 204, 205, 206
 Rhodes, E., 325
 Rivera-Batiz, L. A., 477, 478
 Rivkin, J., 161
 Robbins, H., 78
 Roberts, J., 43, 111
 Robinson, P. M., 524, 527
 Rock, D., 557
 Rodrik, D., 293
 Roesner, F., 446
 Rogerson, R., 204, 241n4
 Romanosky, S., 444, 450, 457, 457n42
 Romer, P. M., 46, 122, 149, 153, 155n2, 159,
 169, 171, 171n9, 172, 255, 477, 478
 Rosenberg, N., 150, 169

 Rosenblatt, F., 73, 124
 Rosenfeld, J., 322
 Rossi-Hansberg, E., 476
 Roth, A. E., 539, 567, 584, 589
 Rousseau, P. L., 41
 Rouvinen, P., 556
 Rowthorn, R., 291, 293n2
 Rubin, D. B., 517, 522, 527, 529
 Rubinstein, A., 427
 Ruffin, R. J., 472
 Ruiz, F. J., 515, 533, 534
 Rumelhart, D. E., 73, 77, 124, 602

 Sabour, S., 76n7
 Sachs, J. D., 336, 344
 Saez, E., 270, 360
 Şahin, E., 270, 329
 Salakhutdinov, R. R., 125
 Salomons, A., 23n1, 309, 555
 Samuelson, P., 376
 Santaaulalia-Llopolis, R., 329
 Saon, G., 25
 Sawyer, J., 597
 Saxenian, A.L., 4, 478
 Schierholz, H., 322
 Schmidhuber, J., 73, 76
 Schmidt, K. M., 358n8
 Schmitt, J., 322, 379
 Schmitz, J. A., 43
 Schmucki, R., 497
 Schultz, P. H., 414
 Schumpeter, J., 148
 Schwartz, M., 568
 Scotchmer, S., 121, 496, 502
 Scott, S. L., 528
 Seater, J. J., 241
 Segal, I., 414, 569, 574
 Seira, E., 514
 Seshadri, A., 205, 243n6
 Shah, A. K., 517
 Shaked, A., 589n1
 Shavell, S., 494n2
 Shaw, J. C., 123
 Shiller, B. R., 410
 Shroff, R., 516
 Silver, D., 63, 66, 453
 Simcoe, T., 487
 Simmons, J. P., 426
 Simon, H. A., 107, 123, 207–8
 Sims, C., 113
 Simsek, A., 362n10
 Singer, Y., 80

- Slovic, P., 596
Smith, A., 590
Smith, M. D., 43
Smith, N., 29
Sokoloff, K. L., 430
Solomonoff, R. J., 253n12
Solove, D., 447, 455
Soloveichik, R., 29
Solow, R. M., 24, 46, 350
Somanchi, S., 444
Song, J., 30
Sonnenschein, H., 589n1
Sopher, B., 590n2
Sorensen, O., 161
Spencer, B. J., 473
Spezio, M., 594n10
Spiegel, M., 589n1
Spiegel, Y., 412
Spier, K., 494, 499, 502
Spiess, J., 169n8, 511, 512
Spindler, M., 522
Srivastava, N., 80
Stahl, D. O., 593
Stantcheva, S., 360
Stein, A., 494
Stern, A., 325
Stern, S., 122, 141
Stevenson, B., 194n5
Stigler, G. J., 439, 582
Stiglitz, J. E., 30n9, 354, 360, 362, 364,
364n11, 366n14, 366n15, 368n19, 369,
370, 370n21, 371, 373, 376n27, 376n28,
378, 412
Stillwell, D., 429
Stinchcombe, M., 74
Stivers, A., 442n4
Stole, L., 268
Strehl, A., 527
Streitwieser, M. L., 562
Strotz, R. H., 427
Stucke, M. E., 414
Stutzman, F., 451
Sufi, A., 208
Summers, L. H., 175
Sutskever, I., 68, 74, 125
Sutton, J., 467, 589n1
Swaffield, J., 293n2
Swaminathan, A., 528
Sweeney, L., 433
Swire, P. P., 441n2
Syverson, C., 23, 29, 43, 210, 319, 320
Taddy, M., 83, 526, 527
Tadelis, S., 107, 577n5, 578, 579, 580, 581,
582
Tang, J., 453
Taylor, C. R., 416, 424, 441n2, 448, 459,
483n19
Tegmark, M., 382, 384n34
Teh, Y.-W., 74
Telang, R., 444, 450, 457, 457n42
Teodoridis, F., 150, 161
Tetenov, A., 527
Thaler, R., 604
Thomas, K., 443
Thomas, P., 528
Thompson, W. R., 82
Thrnton, B., 598
Tibshirani, R., 93, 525, 527
Tirole, J., 106, 112, 264, 268, 269
Topol, E. J., 94, 95
Tory, J., 485n23
Toulis, P., 80
Trajtenberg, M., 4, 39, 116, 119, 132, 150,
169, 176n2
Trefler, D., 465, 467, 479, 485n23
Troske, K. R., 562
Tschantz, M. C., 433
Tucker, C. E., 148, 425, 426, 427, 428, 429,
448, 449, 482, 483, 483n19
Turing, A. M., 123, 385
Tuzel, S., 210
Tversky, A., 596
Ugander, J., 538
Uyarra, E., 480
Vadlamannati, K. C., 483
Valentinyi, Á., 204, 241n4
Van der Laan, M. J., 522, 527
Van Seijen, H., 62, 84
Varian, H. R., 93, 310, 410, 413, 424, 425,
440, 511
Venables, A. J., 480
Vesteger, M., 420
Vickrey, W., 567
Vijverberg, W. P. M., 483
Vincent, N., 270
Vines, P., 446
Vinge, V., 238, 253, 253n12, 373n23
Viscusi, K., 496, 498
Vishnu, A., 151n1, 168
Von Hippel, E., 499

- Von Weizacker, C. C., 376
Von Winterfeldt, D., 595
Vuong, Q., 514
- Wager, S., 523, 525, 527, 528
Wagman, L., 416, 424, 441n2, 448, 459, 483n19
Wallach, I., 116
Wan, M., 534
Wang, J., 475, 477, 594n10
Warren, S. D., 431, 459
Waseem, M., 270
Wattal, S., 450
Weil, D., 459
Weingast, B. R., 577
Weiss, A., 354
Weitzman, M., 149, 151, 157, 171, 171n9, 172, 241
Western, B., 322
Westlake, S., 18
Whinston, M. D., 148
White, H., 74, 511
Williams, H., 122, 364n11
Williams, R., 124
Wilson, P. W., 593
Winter, S., 161
Wolfers, J., 195n5
Wooldridge, J. M., 522
Wright, G., 42n19
Wright, G. C., 303
- Wright, J. R., 594
Wu, L., 43
- Xiao, M., 593n8
Xie, D., 241n4, 295
Xu, H., 577
- Yakovlev, E., 538
Yang, S., 42, 47
Yellen, J., 378
Yeomans, M., 517
Yildirim, P., 107n9
Yu, M., 465, 479
Yudkowsky, E., 253n12
- Zanna, L.-F., 373n24
Zeevi, A., 529
Zeileis, A., 525
Zeira, J., 152, 198, 212n7, 238, 239, 241
Zervas, G., 577n5
Zhang, M. B., 210
Zheng, Y., 329
Zhou, D., 529, 581
Zhou, X., 580
Zhu, X., 471
Zierahn, U., 321
Zubizarreta, J. R., 523
Zweimüller, J., 295
Zwiebel, J., 268

Subject Index

Note: Page numbers followed by “f” or “t” refer to figures or tables, respectively.

- adoption, technological: implications of speed of, for job market and inequality, 310–12
- adversarial artificial intelligence, 401
- aggregate productivity statistics, technologies and, 26–28
- AI. *See* artificial intelligence (AI)
- AlphaGo (Google), 63
- Amazon Go concept, 67
- applied artificial intelligence, 208
- artificial intelligence (AI), 1; and automation of production, 239–50; as basis for learning, 120–21; benefit of more, 318–20; bibliometric data on evolution of, 128–32; capital shares and, 270–74, 272–73f; in context, 84–85; defined, 3–4, 62–67, 93, 122, 237, 468; economies of scale from data and, 468–69; economies of scale from overhead of developing AI capabilities, 469–70; evolution of, 122–25; expected productivity effects of, 45–46; firm organization and, 264–70; future of research on economics of, 17; as general purpose technology, 4–7, 39–41; in idea production function, 250–52; impact of, on innovation, 125–28; impact of long-term decline in labor force participation rate and, 323–25; implications of, 349–53; income distribution and, 351; inequality and, 320–23; internal agreements and, 463; international macroeconomics and, 488; knowledge externalities and, 470–71; likely productivity effects of, and acceleration of, 45–46; longer-term prospects of, 381–86; market structure and, 262–63; as “next big thing,” 175; political economy of, 11, 394–95; prediction costs and, 92–93; privacy and, 425–26; privacy concerns and, 423–24; for promoting trust in online marketplaces, 576–81; recent approach to, 93; for reducing search frictions, 581–83; return of Malthus and, 381–86; revolution, international effects of, 393–94; in Schumpeterian model with creative destruction, 276–79; sectoral reallocation and, 263–64; statistics on, 465–66, 466t; studies on economic effect of, 556–58; theory of privacy in economics and, 424–26; as tool, 16–17; world’s largest companies and exposure to, 465–67, 467t. *See also* machine learning (ML)
- artificial intelligence capital, measuring, 46–50
- artificial intelligence–general purpose technology (GPT) era: education strategies for, 179–82; human-enhancing innovations vs. human-replacing innovations

- artificial intelligence—general purpose technology (GPT) era (*continued*)
 for, 184–85; professionalization of personal services strategies for, 182–84; top skills required for employment in, 180–81, 181t
- artificial intelligence revolution, international effects of, 393–94
- Atomwise, case of, 115–16, 120, 154
- automatic teller machines (ATMs), security policy and, 416
- automation, 3–4, 105–6; basic model, 336–41; Baumol's cost disease and, 241–50; to date, and capital shares, 270–74; decline in labor share and, 329–31; deepening of, 198, 204–5, 216–17; economic adjustment and, 208–9; employment and, 190–91; excessive, 224–26; model of, 211–14; of production, and artificial intelligence, 239–50; productivity and, 210–11; sector of economy affected by, 330–33; studies on employment on, 555–58; wages and, 200–211; winners, 190; work and, 200–211; Zeira model of growth and, 239–41
- average treatment effects, 522–24
- bandits (algorithms), problem of, 528–29
- Baumol's cost disease. *See* cost disease, Baumol's
- BenchSci search technology, 153
- buy/make decisions (firm boundaries), 107–8
- capital accumulation, 198, 204–5, 216
- capital shares, and automation to date, 270–74
- causal inference, new literature on, 519–34
- Children's Online Privacy Protection Act of 1998 (COPPA), 454
- cloud-computing facilities, 402–3
- cluster policies, 480–81
- clusters, regional, theory of, 479
- CNNs (convolutional neural networks), 75–76, 75n6
- collusion, strategies for facilitating, 413–14
- combinatorial-based knowledge production function, 154–61; potential uses of new, 170–71; with team production, 161–67
- Communications Act (1986), 456
- competition policy, innovations and, 141–43
- complexity, 103–4
- consumer attitude, 448–49
- consumer privacy: challenging issues in, 457–59; consumer attitude and, 448–49; consumer risk and, 443–48; data risk and, 442–43; nature of problem of, 442; policy landscape in United States, 454–57; supply side actions and, 450–54. *See also* privacy
- consumer surplus, 11; distribution of, 391–93
- contracting, 106–7
- convolutional neural networks (CNNs), 75–76, 75n6
- cooperation, evolution of, 414
- cost disease, Baumol's, 8–9, 238–39; automation and, 241–50
- creative destruction, 260–61
- data, 61; acquisition methods, 403–4; decreasing marginal returns of, 406, 407f; economics of, 14; importance of, 13–14; important characteristics of, 404–6; localization rules, trade policies and, 485–86; persistence of predictive power of, 427–28; privileged access to government, trade policies and, 486–87; types of, and generation of spillovers, 431–34
- data access, 405–6
- data generation, as pillar of artificial intelligence, 62f, 65–66
- data ownership, 405–6
- data persistence, 426–27; predictive power and, 427–28
- data pipeline, 402
- data pyramid, 404, 405f
- data repurposing, 428–31
- data security: challenging issues in, 457–59; policy landscape in United States, 454–57
- data spillovers, 431–34
- data warehouses, 402–3
- decision-making, baseline model for, 95–103; complexity and, 103–8
- deepening of automation, 198, 204–5, 216–17
- Deep Genomics, 154
- deep learning, 3, 71–77, 400; as general purpose invention in method of invention, 139–43; as general purpose technology, 133–39; as new discovery tool, 167–69; patent systems and, 142

- deep learning networks, 94
 deep learning techniques, 25
 deep neural networks (DNNs), 25–26, 61,
 63; structure in, 76–77
 demand, importance of, 301–2
 destruction, creative, 260–61
 difference-in-difference models, 530–31
 digital information, 334–35
 direct network effects, 412
 displacement effect, 8, 198, 208, 214
 DNNs. *See* deep neural networks (DNNs)
 domain structure, as pillar of artificial intel-
 ligence, 62f, 63–65
 double machine learning, 523–24
- economic growth: artificial intelligence and,
 262–70; prospects for technology-
 driven, 149–53; Zeira model of auto-
 mation and, 239–41. *See also* growth
 economics, impact of machine learning on
 practice of, 15–16
 education, factory model of, 180
 Electronic Communications Privacy Act of
 1986: (ECPA), 456
 employment: automation and, 190–91;
 levels of, and new technologies, 220–21;
 long-run vs. short run, 192–94; studies
 on automation and, 555–57; work out-
 side of, 194–95
 evolution of cooperation, 414
- factory model of education, 180
 Federal Trade Commission (FTC), 454–55
 firm boundaries (make/buy decisions),
 107–8
 firm-level data: need for, 558–59; strategies
 for collecting, 561–62
 firm-level research questions, 560–61
 firms: artificial intelligence and, 262–70;
 impact of machine learning on, 12
- general purpose machine learning (GPML),
 67–71
 general purpose technology (GPT), 2, 65,
 119–20, 169–70; artificial intelligence
 as, 4–7, 39–41; deep learning as, 133–
 39; viewing today's Solow paradox
 through previous history of, 44–45
 generative adversarial networks (GANs),
 66–67
 GPML (general purpose machine learning),
 67–71
- GPT. *See* general purpose technology (GPT)
 Gramm-Leach-Bliley Act (GLBA), 454
 growth: impact of artificial intelligence on,
 7–9. *See also* economic growth
- Health Insurance Portability and Account-
 ability Act of 1996 (HIPAA), 454
 HEI (human-enhancing innovations),
 184–85
 heterogeneous treatment effects, 524–28
 hierarchical Poisson factorization, 510
 HIPAA. *See* Health Insurance Portabil-
 ity and Accountability Act of 1996
 (HIPAA)
 HRI (human-replacing innovations), 184–85
 human-enhancing innovations (HEI), 184–85
 human-replacing innovations (HRI), 184–85
- idea production function, artificial intelli-
 gence in, 250–52
 implementation/restructuring lags, as expla-
 nation for Solow paradox, 31–36
 incentive auctions, machine learning and,
 569–76
 income, artificial intelligence and, 189–90
 income distribution: artificial intelligence
 and, 351; impact of AI on, 11
 income inequality: artificial intelligence and,
 320–23; impact of AI on, 7–8, 11–12;
 speed of technological adoption and,
 310–12
 income redistribution, political economy of,
 394–95
 indirect network effects, 412
 industrial regulation, trade policies and, 487
 inequality. *See* income inequality
 information technology (IT), 24
 innovation, 115–18; competition policy and,
 141–43; early stage, 121–22; impact
 of artificial intelligence on, 125–28;
 institutions and, 141–43; management
 and organization of, 140–41; product
 liability and, 494
 institutions, innovations and, 141–43
 intelligence-assisting innovation (IA),
 350–51
 International Federation of Robotics
 (IFR), 16
 international macroeconomics, artificial
 intelligence and, 488
 international trade, economics of data
 and, 14

- invention of a method of inventing (IMI), 120–21, 124
- inverted-U pattern, 293; simple model of, 297–301
- JDM (judgment and decision-making) research, 596–98
- job displacement, 310–12
- job losses, 291
- job markets, speed of technological adoption and, 310–12
- jobs, impact of artificial intelligence on, 7–8, 9–11
- judgment: in absence of prediction, 96–101; as complements/substitutes to prediction, 102–3; creating role for, 91; prediction and, 91–92
- judgment and decision-making (JDM) research, 596–98
- knowledge creation, 477–79
- knowledge externalities, and artificial intelligence, 470–71
- knowledge spillovers, 479
- labor: comparative advantage of, and new tasks, 217–18; model of demand for, 211–14
- labor demand, technology and, 214–21
- labor productivity growth, technologies and, 26–28
- learning by doing, 412–13
- liability, innovation and, 494; empirical evidence on, 496–98; theoretical model of, 494–96
- liability, tort, development of artificial intelligence technologies and, 498–502
- machine learning (ML), 3, 24–25; applications of, 401–2; defined, 509–10; double, 523–24; early use cases of, 510–15; for finding behavioral variables, 587–96; general purpose, 67–71; human prediction as imperfect, 496–603; impact of, 12–13, 15–17, 507–9; incentive auctions and, 569–76; new literature on, 519–34; overview, 399–406; predictions about impact of on economics, 534–42; regulation and, 12–15; supervised, 511–12; unsupervised, 510–11; vertical integration and, 408–9. *See also* artificial intelligence (AI)
- machine learning–provision industries, 414–15
- machine learning–using industries, 408f; boundaries and, 409–10; firm size and, 409–10; price differentiation and, 410–11; pricing and, 410; returns to scale and, 411–14; structure of, 406–8
- macroeconomics, international, artificial intelligence and, 488
- make/buy decisions (firm boundaries), 107–8
- Maluuba, 62–63
- market design, introduction to, 567–69
- massive open online courses (MOOCs), 181
- matrix completion problem, 531–32
- matrix factorization, 51, 511
- meta technologies, 153
- ML. *See* machine learning (ML)
- model averaging, 511
- MOOCs (massive open online courses), 181
- neural networks, 72–74, 123, 124–25, 510, 511
- new economic geography (NEG), 479–80
- online marketplaces, using artificial intelligence to promote trust in, 576–81
- optical lenses, invention of, 121
- optimism, sources of technological, 24–26
- Pareto improvement, 363
- patent systems, deep learning and, 142
- Pen Register Act (1986), 456
- policy analysis, using methods of prediction for, 516–19
- political economy: of artificial intelligence, 11; of income redistribution, 394–95; of technological disruptions, 176–79
- prediction: in absence of judgment, 101–2; artificial intelligence as tool for, 16–17; as complements/substitutes to judgment, 102–3; costs of, and AI, 92–93; judgment and, 91–92; using methods of, in policy analysis, 516–19
- premature deindustrialization, 393–94
- price discrimination, artificial intelligence and, 604
- principal components analysis, 74, 92, 510
- privacy, 13–14; artificial intelligence and,

- 425–26; current models of economics and, 424–25; data spillovers and, 431.
See also consumer privacy
- privacy policy, 416–17
- privacy regulation, trade policies and, 482–85
- privileged access to government data, trade policies and, 486–87
- production, automation of, and artificial intelligence, 239–50
- productivity: automation and, 210–11; missing growth of, and new technologies, 223–26
- productivity effects, 198, 203–4, 214–16
- productivity growth: low current, reasons why it is consistent with future technological growth, 41–44; rates of, technologies and, 26–28; slow, and future productivity growth, 31–36
- productivity optimism, technology-driven case for, 36–39
- radiology, case of, 94–95
- random forest, 511
- R&D. *See* research and development (R&D)
- recommender systems, 603
- regional clusters, theory of, 479
- regression trees, 511
- regularization on norm of matrix, 510
- regularized regression, 511
- regulation, 12–15; machine learning and, 12–15
- reinforcement learning, 25, 66, 81–84, 400–401
- reinstatement effect, 8, 198, 206
- research and development (R&D), 336; productivity, effects of rise in, 341–43
- research tools, economics of new, 118–22
- robotics, 123–24; tort law and, 493–94.
See also robots
- robots, studies on, 554–55. *See also* robotics
- Romer/Jones knowledge production function, 151–52
- scale, economies of, and artificial intelligence, 470
- Schumpeterian model with artificial intelligence, 276–79
- scientific discovery, rate of, 6
- scientists, role of, 472–73; superstar, 474–76
- scope, economies of, and artificial intelligence, 470
- search frictions, artificial intelligence for reducing, 581–83
- security policy, 416
- singularities, 253–61; examples of technological, 254–58; objections to, 258–61
- skills: mismatch of technologies and, 221–23; technologies and, 209
- Solow paradox, 24; potential explanations for, 28–31
- source code, trade policies and, 487–88
- spectrum reallocation, 569–76, 572f
- spillovers, data, 431
- spreadsheet software, invention and impact of, 90
- stochastic gradient descent optimization, 77–81
- strategic trade policy, 473–74
- structural change, 293–96
- structural models, 532–34
- superstar scientists, role of, 474–76
- supervised machine learning, 511; methods for, 511–12
- supplementary analysis, 530
- support vector machines, 511
- symbolic processing hypothesis, 123
- symbolic systems, 123
- tasks, new, 205–7; comparative advantage of labor and, 217–18; creation of, 198; model of, 211–14
- technological changes: factor-biased, 376–77; and levels of employment, 220–21; types of, 212–13
- technological disruptions, political economy of, 176–79
- technological growth, reasons why it is consistent with low current productivity growth, 41–44
- technological optimism, sources of, 24–26
- technological progress: channels of inequality and, 365–70; determining scenarios that best describe economy, 363–64; endogenous, 364–65; first-best scenario, 353–56; imperfect markets scenario, 361–62; perfect markets but costly redistribution scenario, 358–61; perfect markets ex post and no costs of redistribution scenario, 356–58; welfare and, 353–65; worker-replacing, redistribution and, 370–77
- technological singularities, examples of, 254–58

- technological unemployment, 377–81
- technologies: future progress of, and low current productivity growth, 41–44; labor demand and, 214–21; mismatch of skills and, 221–23; skills and, 209
- technology-driven economic growth, prospects for, 149–53
- tort law, robotics and, 493
- tort liability, 14
- total factor productivity growth, 32
- trade models, basic, 476–77
- trade policies: data localization rules and, 485–86; industrial and strategic, case for, 471–81; industrial regulation and, 487; privacy regulation and, 482–85; privileged access to government data and, 486–87; role of university-related talent, 472–76; source code and, 487–88; strategic, 473–74
- UBI. *See* universal basic income (UBI)
- unemployment, technological, 377–81
- universal basic income (UBI), 312–14; cost of replacing current safety net with, 325–26
- unsupervised machine learning, 510–11
- USA Patriot Act (2001), 456
- vertical integration, machine learning and, 408–9
- vertical research spillovers, 122
- wages, automation and, 200–211
- welfare, technological progress and, 353–65
- work, automation and, 200–211
- worker-replacing technological progress: dynamic implications of, 373–74; redistributing innovators' surplus and, 374–76; redistribution and, 370–77; static pecuniary externalities of, 371–72
- Zeira model of automation and growth, 239–41
- “zero-shot” learning systems, 66