

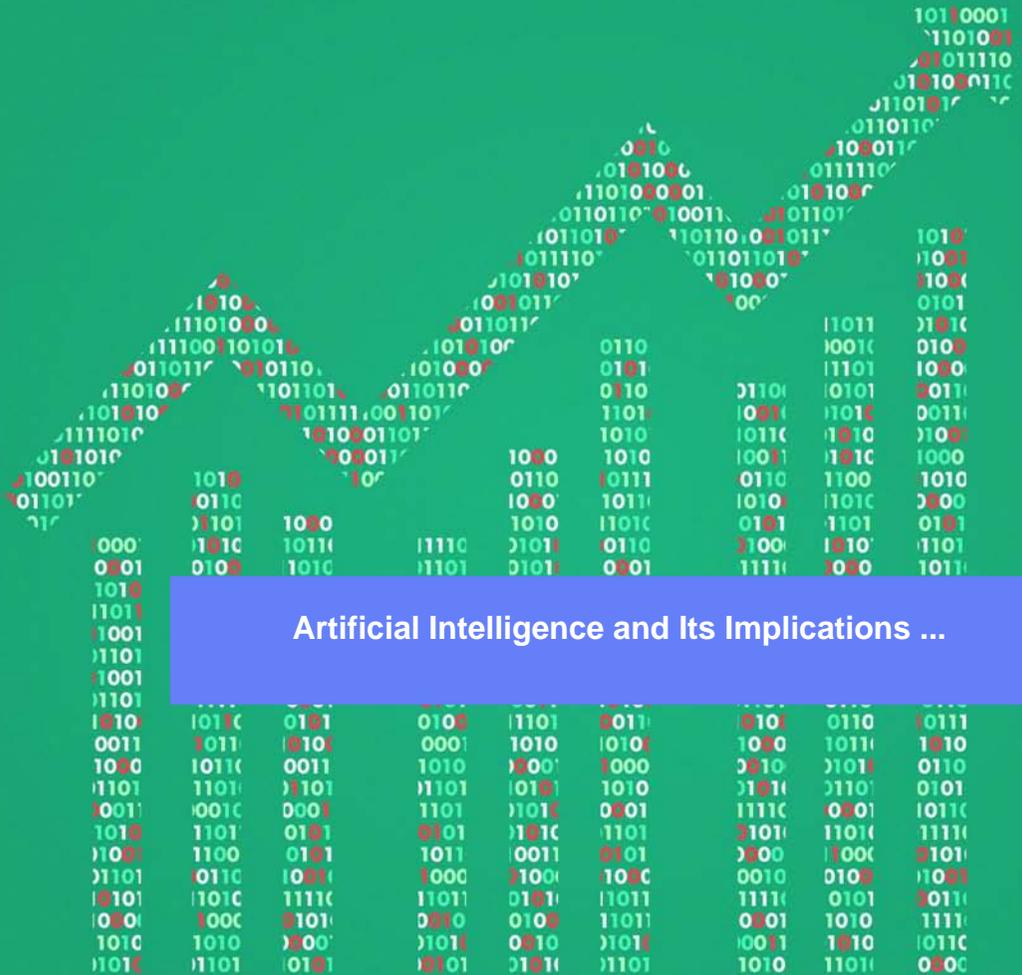


National
Bureau of
Economic
Research

THE ECONOMICS OF ARTIFICIAL INTELLIGENCE

An Agenda

Edited by Ajay Agrawal,
Joshua Gans, and Avi Goldfarb



Artificial Intelligence and Its Implications ...

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
Comment: 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
Comment: 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
Comment: 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 Its Implications for Income Distribution and Unemployment

Anton Korinek and Joseph E. Stiglitz

14.1 Introduction

The introduction of artificial intelligence (AI) is the continuation of a long process of automation. Advances in mechanization in the late nineteenth and early twentieth centuries automated much of the physical labor performed by humans. Advances in information technology in the mid- to late twentieth century automated much of the standardized data processing that used to be performed by humans. However, each of these past episodes of automation left large areas of work that could only be performed by humans.

Some propose that advances in AI are merely the latest wave in this long process of automation, and may in fact generate less economic growth than past technological advances (see, e.g., Gordon 2016). Others, by contrast, emphasize that AI critically differs from past inventions: as artificial intelligence draws closer and closer to human general intelligence, much of human labor runs the risk of becoming obsolete and being replaced by AI in all domains. In this view, progress in artificial intelligence is not only a continua-

Anton Korinek is associate professor of economics and business administration at the University of Virginia and Darden GSB and a research associate of the National Bureau of Economic Research. Joseph E. Stiglitz is University Professor at Columbia University and a research associate of the National Bureau of Economic Research.

This chapter was prepared as a background paper for the NBER conference The Economics of Artificial Intelligence. We would like to thank our discussant Tyler Cowan as well as Jayant Ray and participants at the NBER conference for helpful comments. We also acknowledge research assistance from Haaris Mateen as well as financial support from the Institute for New Economic Thinking (INET) and the Rewriting the Rules project at the Roosevelt Institute, supported by the Ford, Open Society, and the Bernard and Irene Schwartz Foundations. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <http://www.nber.org/chapters/c14018.ack>.

tion, but the culmination of technological progress; it could lead to a course of history that is markedly different from the implications of previous waves of innovation, and may even represent what James Barrat (2013) has termed “Our Final Invention.”

No matter what the long-run implications of AI are, it is clear that it has the potential to disrupt labor markets in a major way, even in the short and medium run, affecting workers across many professions and skill levels.¹ The magnitude of these disruptions will depend on two important factors: the speed and the factor bias of progress in AI.

On the first factor, measured productivity has increased rather slowly in recent years, even as the world seems to be captured by AI fever.² If AI-related innovations enter the economy at the same slow pace as suggested by recent productivity statistics, then the transition will be slower than, for example, the wave of mechanization in the 1950–1970s, and the resulting disruptions may not be very significant. However, there are three possible alternatives: First, some suggest that productivity is significantly undermeasured, for example, because quality improvements are not accurately captured. The best available estimates suggest that this problem is limited to a few tenths of a percentage point (see, e.g., the discussion in Groshen et al. [2017]). Furthermore, there are also unmeasured deteriorations in productivity, for example, declines in service quality as customer service is increasingly automated. Second, the aggregate implications of progress in AI may follow a delayed pattern, similar to what happened after the introduction of computers in the 1980s. Robert Solow (1987) famously quipped that “you can see the computer age everywhere but in the productivity statistics.” It was not until the 1990s that a significant rise in aggregate productivity could be detected, after sustained investment in computers and a reorganization of business practices had taken place. Third, it is of course possible that a significant discontinuity in productivity growth occurs, as suggested, for example, by proponents of a technological singularity (see, e.g., Kurzweil 2005).

On the second factor, the disruptions generated by AI-related innovations depend on whether they are labor-saving, using the terminology of Hicks (1932), that is, whether at a given wage the innovations lead to less demand for labor. Some suggest that artificial intelligence will mainly *assist* humans in being more productive, and refer to such new technologies as *intelligence-assisting innovation* (IA), rather than AI. Although we agree that most AI-related innovations are likely to be complementary to at least some jobs, we believe that in taking a broader perspective, progress in AI

1. For example, Frey and Osborne (2017) warn that 47 percent of jobs in the US economy are at risk of being automated by advances in AI-related fields. Areas in which human intelligence has recently become inferior to artificial intelligence include many applications of radiology, trading in financial markets, paralegal work, underwriting, driving, and so forth.

2. For example, Google Trends reveals that search interest in the topic “artificial intelligence” has quadrupled over the past four years.

is more likely to substitute for human labor, or even to replace workers outright, as we will assume in some of our formal models below.

We believe that the primary economic challenge posed by the proliferation of AI will be one of income distribution. We economists set ourselves too easy a goal if we just say that technological progress *can make everybody better off*—we also have to say *how we can make this happen*. This chapter is an attempt to do so by discussing some of the key economic research issues that this raises.³

In section 14.2 of this chapter, we provide a general taxonomy of the relationship between technological progress and welfare. We first observe that in a truly first-best economy—in which complete risk markets are available before a veil of ignorance about innovations is lifted—all individuals will share in the benefits of technological progress. However, since the real world does not correspond to this ideal, redistribution is generally needed to ensure that technological progress generates Pareto improvements. If markets are perfect and redistribution is costless, it can always be ensured that technological progress makes everybody better off. The same result holds if the costs of redistribution are sufficiently low. In all these cases, there can be political unanimity about the desirability of technological progress. However, if redistribution is too costly, it may be impossible to compensate the losers of technological progress, and they will rationally oppose progress. Even worse, if the economy suffers from market imperfections, technological progress may actually move the Pareto frontier inwards, that is, some individuals may necessarily be worse off. Finally, we observe that the first welfare theorem does not apply to the process of innovation, and as a result, privately optimal innovation choices may move the Pareto frontier inward.

In section 14.3, we decompose the mechanisms through which innovation leads to inequality into two channels. First, inequality rises because innovators earn a surplus. Unless markets for innovation are fully contestable, the surplus earned by innovators is generally in excess of the costs of innovation and includes what we call innovator rents. We discuss policies that affect the sharing of such rents, such as antitrust policies and changes in intellectual property rights. The second channel is that innovations affect market prices; they change the demand for factors such as different types of labor and capital, which affects their prices and generates redistributions. For example, AI may reduce a wide range of human wages and generate a redistribution to entrepreneurs. From the perspective of our first-best benchmark with complete insurance markets, these factor price changes represent pecuniary externalities. We discuss policies to counter the effects of the resulting factor price changes.

3. An important, and maybe even more difficult, complementary question, which is beyond the scope of this chapter, is to analyze the political issues involved.

In section 14.4, we develop a simple formal model of worker-replacing technological change, that is, we introduce a machine technology that acts as a perfect substitute for human labor. We study the implications for wages and discuss policy remedies. In the short run, an additional unit of machine labor that is added to the economy earns its marginal product, but also generates a zero-sum redistribution from labor to traditional capital because it changes the relative supply of the two. In the long run, the machine technology turns labor into a reproducible factor. Thus, in the long run, growth will likely be limited by some other irreproducible factor, and all the benefits of technological progress will accrue to that factor. However, since it is in fixed supply, it can be taxed and the proceeds can be redistributed without creating distortions. Hence a Pareto improvement is easily achieved.

In a second model, we demonstrate how changes in patent length and capital taxation can act as a second-best device to redistribute if lump sum transfers between workers and innovators are not available. A longer patent life both delays how quickly innovations enter the public domain, lowering consumer prices, and increases the incentives of innovators to produce worker-replacing machines. However, the resulting losses for workers can be made up for by imposing a distortionary tax on capital and providing transfers, so long as the supply elasticity of capital is sufficiently low.

We also discuss the implications of endogenous factor bias in technological change. Worker-replacing technological progress should make capital-saving innovations more desirable, providing some relief to workers. We also note that our economy is developing more and more into a service economy, and that the large role of government in many service sectors (e.g., education, healthcare, etc.) creates ample scope for interventions to support workers.

In section 14.5, we observe two categories of reasons for why innovation may lead to technological unemployment. The first category of reasons arises because wages cannot adjust, even in the long run: efficiency wage theory implies that employers may find it efficient to pay wages above the market-clearing level so that workers have incentives to exert proper effort. If technological progress lowers the marginal product of workers, and hence their real wage declines below their cost of living, then classic nutritional efficiency wage theories apply: unemployment would result because (in the absence of government support) workers could not survive working for the market-clearing wage and it would pay employers to raise real wages above the market-clearing level because of the resulting increase in worker productivity. The second category of technological unemployment arises as a transition phenomenon, when jobs are replaced at a faster rate than workers can find new ones. We discuss a variety of factors that may slow down the adjustment process. Efficiency wage arguments may also play an important role as a transitional phenomenon, in particular if workers' notion of fair wages is sticky. Finally, we discuss that jobs may not only provide wages but

also meaning and note that, unless societal attitudes change with the proliferation of AI, it may be welfare enhancing to subsidize jobs rather than simply redistributing resources.

In section 14.6, we take a longer-term perspective that is somewhat more speculative and discuss the potential implications of superhuman artificial intelligence. We consider two scenarios: one in which some humans use technology to enhance themselves and attain superhuman intelligence, and one in which autonomous machines that are completely separate from humans reach superhuman intelligence. In both cases, the superior productivity of superior intelligence will likely lead to vast increases in income inequality. From a Malthusian perspective, the superintelligent entities are likely to command a growing share of the scarce resources in the economy, creating the risk of pushing regular humans below their subsistence level. We discuss corrective actions that could be taken.

14.2 Technological Progress and Welfare: A Taxonomy

In 1930 Keynes wrote an essay on the “Economic Possibilities of our Grandchildren,” in which he described how technological possibilities may translate into utility possibilities. He worried about the quality of life that would emerge in a world with excess leisure. And he thought all individuals might face that quandary. But what has happened in recent years has raised another possibility: innovation could lead to a few very rich individuals—who may face this challenge—whereas the vast majority of ordinary workers may be left behind, with wages far below what they were at the peak of the industrial age.

So let us start by considering the arrival of a new technology that partially (or fully) replaces workers and let us ask the question: *would their standard of living necessarily decline?* We will consider this question in a number of different settings, providing a taxonomy for how technological progress might affect the welfare of different groups in society depending on the environment.

14.2.1 First Best

We start with a first-best scenario in which we assume that all markets are perfect: this includes risk markets that are free of adverse incentive effects and that allow individuals to insure against the advent of innovations “behind the veil of ignorance,” that is, before they know whether they will be workers or innovators. The main purpose for considering this idealized setting is to demonstrate that from an *ex ante* perspective, compensating workers for the losses imposed by technological progress is a question of economic efficiency not redistribution.

If risk markets were perfect and accessible to all agents before they knew their place in the economy, then all agents would be insured against any risk

that might affect their well-being, including the risk of innovation reducing the value of their factor endowment. For example, workers would be insured against the risk of declining wages.⁴ This leads us to the following observation:

OBSERVATION 1) *Consider a first-best world in which all individuals have access to a perfect insurance market “behind the veil of ignorance,” that is, before they know whether they will be innovators or workers. If an innovation occurs in such a world, the winners would compensate the losers as a matter of optimal risk sharing. As a result, technological progress always makes everybody better off, and there is political unanimity in supporting it.*

This is a powerful observation because it reminds us that if we had an ideal market, something that very much looks like redistribution would naturally emerge. In our first-best economy, there are no losers from technological progress. Losers only exist if risk markets are imperfect compared to this benchmark. In more technical language, worker-replacing technological progress imposes pecuniary externalities on workers, which lead to inefficiency when risk markets are imperfect (see, e.g., Stiglitz and Weiss 1981; Greenwald and Stiglitz 1986; Geanakoplos and Polemarchakis 1986; or more recently Dávila and Korinek 2018).

This implies that policy measures to mitigate or undo the pecuniary externalities arising from technological progress—for example, redistribution programs—make the economy’s allocation more efficient from an ex ante perspective, rather than “interfering” with economic efficiency. They bring us closer to the allocation that a well-functioning risk market would achieve. Policymakers who oppose redistribution to compensate the losers of innovation because it interferes with the free market seem to—inappropriately, in our view—take an ex post perspective, after an innovation has taken place and after individuals know their place in the economy. Even though they may pretend to preach about idealized free markets, they clearly have not understood the full implications of how an idealized free market would work, that is, that such a market would provide precisely the type of insurance that they are opposing.

In practice, workers who might be replaced by technological progress cannot purchase insurance contracts against being replaced, so in the absence of adequate government assistance, they are in fact hurt by the innovation. Of course there are good reasons for why such idealized risk markets are not present in the real world.

First, the limited lifespan of humans makes it difficult to write insurance contracts that stretch over multiple generations. Workers would have had to obtain the described insurance a long time ago, before AI was well

4. We will discuss the reasons why this is typically not the case in practice below.

conceived and its implications were clear, when the associated insurance premium would have been commensurately low. Perhaps their farsighted ancestors could have written state-contingent contracts on their behalf. Today, obtaining insurance against AI-reducing wages would require workers to pay large amounts since the possibility is very real. In short, effective insurance would have had to take place behind a “veil of ignorance” about the likely advent of AI.

To put it another way, in this perspective the first “insurable damage” to the individual occurs at the time that the probability of an innovation becomes nonnegligible, for at that time the insurance premium required for income smoothing becomes significant, and her welfare is lowered. The individual would have wanted to buy insurance against the risk that her insurance premium would go up. Thus, in a perfect market, insurance markets would have to go back at least to a date at which there was a negligible probability that the innovation occurs. This presents a problem: it may be that at the moment that the concept of AI is formulated precisely enough to be an insurable event (and therefore becomes an insurable event), it has a non-zero probability.

Second, even for more limited time periods, risk markets with respect to technological change are clearly not perfect. Among the main reasons are information problems.

Describing the State Space. This starts with the basic problem alluded to earlier of how difficult it is to describe the future state space.⁵ We cannot easily write a contract on something before it has been invented. Addressing this problem would require that an individual has to be insured against any technological event that leads to lower wages. But any such insurance contract would necessarily have adverse incentive effects.

More broadly, the curse of asymmetric information that inhibits insurance markets is as prevalent here as it is elsewhere.

Adverse Selection. Innovation leads to important adverse-selection problems. Some people in the market are more informed than others. In an ideal market, the winners of innovation would provide insurance to the losers, and the winners (e.g., entrepreneurs) would almost certainly be better informed than the losers (e.g., workers).

Moral Hazard. Innovation may also be subject to moral hazard problems, that is, the presence of insurance may affect the likelihood that the insured event occurs. Although workers are unlikely to affect the pace of innovation in AI, the actions of innovators may be, to some extent, affected. If they were to completely insure away all their returns from innovation, there

5. Interestingly, this type of information problem is easy to deal with after innovation has occurred, because then we know what has been invented and in which state we are, but very difficult to capture in ex ante contracts.

would be scant incentive to exert effort.⁶ Since, in a perfect insurance world, the winners would insure the losers, full insurance would lead to stagnation.

Insurance and Redistribution

A natural counterpart to observation 1 is that in the absence of perfect insurance markets “behind the veil of ignorance,” ensuring that progress leads to a Pareto improvement generally requires redistribution. If workers have access to some insurance against the risk of AI but not perfect insurance, this does not remove the need for redistributions.

For example, obtaining AI insurance today would require workers to pay a large premium. Of course, conceptually, if one went back in time, before AI was well conceived and its implications clear, one might argue that the premium would be low. But even that might not be so, since premiums for large events, even with small probability, can be high. In any case, at the very moment of conception of AI—the first possible moment that one could conceivably have written a policy—AI would still have distributional consequences; workers would have to pay a premium to divest themselves of this risk, and thus they would be worse off compared to the innovators, the winners.

14.2.2 Perfect Markets Ex Post and No Costs of Redistribution

Our next case pertains to a world that may be described as a second-best world without the perfect insurance markets referred to earlier, but in which, ex post, all markets are functioning well *and* there can be costless redistributions. This case covers several critical results that, although obvious at some level, often get lost in the debate about AI and technological progress more generally.

OBSERVATION 2) *If redistribution is costless and appropriate redistributions are made, then technological progress is always desirable for all agents. In that case, there is political unanimity in supporting technological progress.*

For convenience, and in conformity to conventional usage, we will refer the world with costless redistribution but otherwise perfectly functioning markets, as first best ex post; though we remind the reader that the previous analysis suggested that in a true first-best, workers would have insurance against the risks from AI, such that they would commensurately share any gains from innovation. If the world is first-best ex post in the sense thus

6. Some might argue that this problem is equally hard to deal with before or after innovation has occurred. If we tax innovators ex post, it destroys incentives just as much as if we fully insure away all returns from innovation. However, in both cases, the significance of any adverse effects is not clear. Innovators are at least partially driven by non-pecuniary motives. And partial insurance or partial redistribution are always an option. If Bill Gates had been told, ex ante, that government would take away 50 percent of his returns over \$10 billion, there is little reason to believe that it would have had any significant effect on innovation and investment. Ex post, taxing the winners in “winner takes all” games may have only small incentive effects.

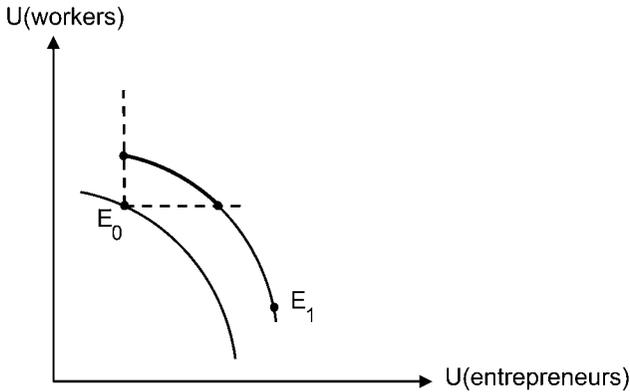


Fig. 14.1 Pareto frontier before and after innovation with costless redistribution

defined, then the utility possibilities curve (or Pareto frontier) moves out. We provide an example in figure 14.1, which depicts a utility possibilities frontier for two types of agents, workers, and entrepreneurs. In the example, technological progress increases the maximum utility level of entrepreneurs for any given level of utility of workers.⁷ Innovation has increased production possibilities, and with lump sum redistributions, an expansion in production possibilities automatically implies an expansion in utility possibilities, that is, that everybody could be better off.

The fact that they *could* be better off does not mean that they *will* be better off. That depends on institutional arrangements. In figure 14.1, we denote the initial equilibrium by E_0 and the after-innovation equilibrium by E_1 . We have deliberately not called it a competitive market equilibrium: markets do not exist in a vacuum (see e.g. Stiglitz et al. 2015). They are structured by rules and regulations, for example, concerning intellectual property rights and antitrust policies, and there may be tax and other policies in place. We thus simply refer to E_0 and E_1 as the before and after innovation (institution-given) equilibrium given the existing set of institutions. Note that in the example drawn in the figure, workers are worse off. That would normally be the case with what Hicks referred to as labor-saving innovations, that is, innovations that at a given wage lead to less demand for labor. Artificial intelligence appears to be a labor-saving innovation. In the simple formal models of worker-replacing innovations that we work out below, that is clearly the case.

This in turn has two important implications.

First, without adequate redistribution it makes sense for workers to resist the innovation. Luddism—the movement named after the possibly fictional

7. More generally, we could define a multidimensional utility possibilities frontier by adding any number of categories of individuals, or even naming the individuals.

character Ned Ludd that opposed automation in the textile sector in the late eighteenth and early nineteenth centuries in England—is a rational response for workers who are worse off from automation and who are not sufficiently compensated.

Second, in a democracy in which workers are in a majority it would make sense for enlightened innovators to support redistribution, to make sure that workers are at least not worse off. With redistribution, both innovators and workers *can* be better off. If appropriate redistribution is made so that everybody shares in the fruits of technological progress, there will again be political unanimity in supporting technological progress—progress will not be politically contentious.

There might be significant debate about how much compensation workers should receive, that is where in the “northeast corner of E_0 ” society should be. On the one hand, this debate concerns the distribution of the surplus generated by innovation. On the other hand, labor-saving innovation reduces wages, which generates a redistribution from workers to other factor owners like rentiers and capitalists, for which workers may seek compensation. This redistribution represents a pecuniary externality from the innovation, as we will discuss in further detail in section 14.3.2.

In figure 14.1, we have marked in bold that part of the postinnovation Pareto frontier that represents a Pareto improvement and lies to the northeast of E_0 . A range of philosophical principles can be adduced for determining what is a “just” division of the fruits of innovation. Behavioral economics may provide insights into what kinds of divisions might be acceptable.⁸

Of course, the innovation may not be labor saving, and the equilibrium E_1 itself could be to the northeast of E_0 . Although this case is easier, the distribution of the gains from innovation and any associated pecuniary externalities and rents may still be contentious, especially if they lead to large disparities in income. Distributive issues can also interact with production, as emphasized for example, by the efficiency wage theories that we consider in greater depth in section 14.5.

14.2.3 Perfect Markets but Costly Redistribution

There is another possibility—that as we try to redistribute, the new utility possibility curve may lie inside of the old utility possibilities frontier near the original equilibrium. This may be the case even in a world that is first-best

8. Consider a model in which workers and innovators have to agree on whether the innovation is acceptable. The innovator has the power to set the division of the gains (i.e., where along the curve Northeast of E_0 the new equilibrium lies), but the workers have the power to accept or reject. This is the standard ultimatum game, for which there is a large body of literature suggesting that at least some of the fruits of innovation have to be shared with workers. If they perceive the allocation of benefits to be unfair, they would rather be worse off (e.g., at the original point without the innovation) than at the point that just makes them indifferent to where they were before. See Fehr and Schmidt (2003).

ex post, that is, in which all the conditions for Pareto efficiency would be satisfied ex post after the innovation has taken place.

OBSERVATION 3) *If the world is first-best ex-post, but redistribution is limited or costly, then a Pareto improvement may not be possible, and some groups in society may oppose technological progress. With a sufficiently inequality-averse social welfare function, societal welfare may be reduced.*

This case is illustrated in figure 14.2. The utility possibilities frontier is constrained by the costs imposed by redistribution. Even though it might appear that innovation could make everyone better off technologically, given the existing set of institutions of that economy, it actually can't—there may not be scope for avoiding utility losses for workers.

Some economists argue that the world looks like figure 14.2, and that if we try to transfer from innovators to workers, so much output is lost that workers are still worse off. If that is the case, then one cannot say that the innovation is a Pareto improvement. One hesitates to use the word “innovation.” It is a change, perhaps a technological change, which has had the effect of making some people better off and others worse off. It is a distribution-inducing change and will be contentious.

A social welfare function that places no weight on inequality—which treats a dollar to rich innovators the same as a dollar to a poor worker—would, of course, conclude that the innovation is desirable. But with a more natural, inequality-averse social welfare function, the so-called innovation is welfare decreasing.

The workers who lose out would rationally oppose the innovation. If workers are in a majority and innovators wish to maintain their position, it would behoove the innovators to think harder about how to engage in redistribution. This is, of course, a collective action problem for innovators—for individual innovators, the contribution to economy-wide inequality is typi-

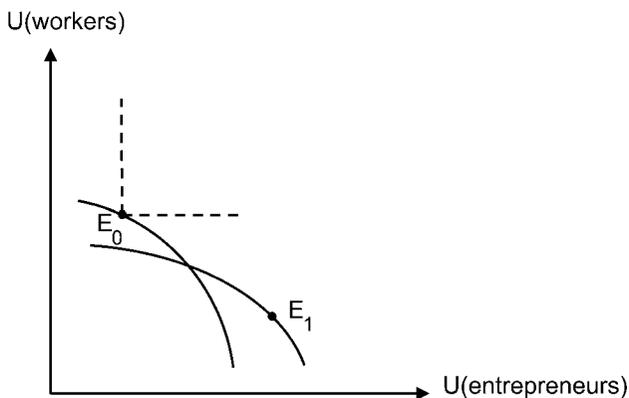


Fig. 14.2 Potential Pareto frontier with costly redistribution

cally limited, even if their collective behavior makes workers worse off. As a result, innovators often devote effort to actions that enhance their market power—lowering *real* incomes of workers still further—and to not paying taxes (both via clever tax avoidance using the existing legal framework, and via political lobbying to provide special exemptions from taxation for their industries). Disregarding, in our view unwisely, that their actions may stir up political opposition to innovation, some innovators go further and argue for weakening the progressivity of the tax system and a smaller state, so there are less public resources to provide for the well-being of the workers who are hurt by innovation.

According to a long-run version of “trickle-down” economics, repeated innovations will eventually increase the wealth of innovators so much that the benefits will trickle down to regular workers. In this view, a Pareto improvement is always possible in the long run, as in figure 14.1, even if an entire generation of workers is hurt in the short-to-medium run. This is a possibility and, in fact, the first industrial revolution may be an example. During the Industrial Revolution, workers eventually obtained enough human capital—which was publicly provided, as is in the interests of the innovators—so that the wages of almost all increased. In the current context, however, once machines are smart enough, innovators may no longer have incentives to support the public financing of human capital accumulation, and it may well be that workers’ standards of living decrease. In particular, in a political system dominated by money, the innovators, increasingly rich, may use their economic and political influence to resist redistribution. Furthermore, even if long-run trickle-down economics was correct, it may lead to tremendous suffering and social upheaval in the short run. It may also—understandably—not be very credible if innovators promise that once they are rich enough, they will support workers, but that they are not quite rich enough yet.

This leads to the important question: *How costly is redistribution in practice?* As we noted earlier, markets do not exist in a vacuum. They are structured by laws and regulations and by how those laws and regulations are enforced. The outcome is the so-called “market” distribution of income, which is then subject to taxes and transfer, leading to an after-tax distribution of income. But this conventional distinction may not be quite accurate: the rules of the game concerning redistribution affect the market income distribution, and are themselves endogenous, affected by the rules of the political game, which in turn are affected by the distribution of wealth. (See Piketty, Saez, and Stantcheva (2014) and Stiglitz 2017.) The points that we have denoted E_0 and E_1 describe the initial outcome and the outcome after technological change, *assuming that laws, regulations, institutions, and so forth remain unchanged*. But, of course, it is not reasonable to expect that they would remain unchanged with the advent of a change as significant as AI.

Setting aside the endogeneity of the rules themselves, each set of (feasible)

laws, regulations, institutions, and so on defines a feasible utility possibilities frontier. We can think of the second-best utility possibilities frontier as the outer envelope of all these frontiers. As the outer envelope, the second-best utility possibilities frontier provides more flexibility for redistribution than does that associated with one particular set of rules, regulations, and institutions. This reflects that any changes in laws, regulations, or institutions, and so forth will also have redistributive effects. Given this additional flexibility, the likelihood that a Pareto improvement as in figure 14.1 can be achieved is greater. We provide further arguments for why this is likely to be the case in section 14.4 (although we cannot entirely rule out a situation like that depicted in figure 14.2 in which a Pareto improvement is not feasible).

14.2.4 Imperfect Markets

Let us also consider a fourth case, which does not necessarily reflect the specific situation with advances in artificial intelligence, but which is important to understand and keep in mind when we evaluate technological innovations.

OBSERVATION 4) *If the economy is not first-best ex post, then the utility possibilities frontier may move inward in response to an expansion of production possibilities. Furthermore, this may even be true with costless redistribution.*

When we speak about an economy that is not first-best, we mean an economy that deviates from the Arrow-Debreu benchmark, that is, that exhibits market imperfections such as information problems, missing markets, and price and wage rigidities, which can result in aggregate demand problems, monopolies and monopsonies, and so forth. Typically, these mean that the market equilibrium is not Pareto efficient. The utility possibilities frontier represents the maximum utility of workers, given that of entrepreneurs, taking the market failures as given.

This case is illustrated in figure 14.3. The initial equilibrium is E_0 , but the innovation, which would have led to greater efficiency in the absence of these market imperfections, makes workers worse off—and even with costless redistributions, there is no way that both workers and entrepreneurs can be better off.

An example, elaborated on by Delli Gatti et al. (2012a, 2012b), were the agricultural improvements at the end of the nineteenth century and beginning of the twentieth. The result was that agricultural prices plummeted, and so too did incomes on farms and in the rural sector. But mobility is costly—moving to the urban sector required capital, and many farmers saw their capital disappear as the value of their farms decreased. Those with loans often went bankrupt. Capital market imperfections (based on information asymmetries) meant that farmers could not borrow to move to the city to where the new jobs would be created. But as incomes in the rural sector plummeted, they could not buy the goods made by the manufacturing

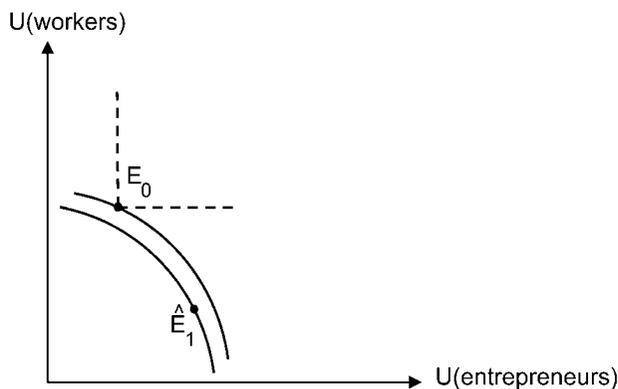


Fig. 14.3 Potential Pareto frontier with market imperfections

sector. Workers in both the rural and urban sector were worse off.⁹ This provides one interpretation of the Great Depression—in the short run, the innovations at the time proved Pareto inferior.

Another example is given by the now-standard result that free trade may lead everyone to be worse off in the absence of good risk markets (Newbery and Stiglitz 1984). That result can be interpreted as one involving technological progress. Assume that there was no way of transporting goods between two countries. A technological advance allows goods to be transported freely. Then, under the quite plausible conditions postulated by Newbery-Stiglitz, welfare (of everyone!) in both countries could decrease.

The theory of the second-best (Meade 1955; Lipsey and Lancaster 1956) reminds us that in the presence of market imperfections, improving the functioning of one market may deteriorate overall welfare. There are reasons to believe that certain innovations in financial markets, for example, structured financial products and certain derivatives like credit default swaps, especially in the absence of appropriate regulations, contributed greatly to the Great Recession (Financial Crisis Inquiry Commission 2011).¹⁰

It is important to appreciate the result described in Observation 4 to understand how crucial our institutions and our market imperfections are in determining whether and how large a benefit society will derive from innovation.

9. In the central Delli Gatti et al. model, the agricultural sector has constant returns to scale and wages in the urban sector are rigid (e.g. because of efficiency wage considerations), so that the agricultural innovations are unambiguously welfare decreasing. In one variant of the model, where urban wages are flexible, wage decreases lead to still higher unemployment. Though it is possible that entrepreneurs gain more from the wage reductions than they lose from the loss of sales, social welfare is decreased with sufficiently inequality averse social welfare functions.

10. At a theoretical level, Simsek (2013) and Guzman and Stiglitz (2016a, 2016b) have shown that opening up new markets—through financial innovation—can lead to greater volatility in consumption.

14.2.5 Ascertaining Whether the Economy Is Best Described by Observation 1, 2, 3, or 4

It is not always easy to ascertain which of the four observations in the four subsections above best describes the economy. Typically, the only thing that we can observe is that an innovation has made some individuals better off and some worse off. (In our analysis of AI, we assume that it has made workers worse off and entrepreneurs better off.) The presumption is that risk markets for innovation are highly imperfect (so Observation 1 does not apply), redistributions are costly (so Observation 2 does not strictly apply), and markets are imperfect (so Observation 3 does not strictly apply.) But the costs of redistributions may be sufficiently low and the market imperfections sufficiently small that figure 14.1 still applies: everyone could be made better off. Alternatively, redistributions may be so costly that figure 14.2 applies. Or market failures may be sufficiently large and redistributions sufficiently costly that figure 14.3 applies.

We emphasize that which situation we are in depends not just on the possibilities of ex post redistribution, but on the institutional flexibility, which determines the ex ante distribution.

As we noted, the second-best utility possibilities frontier is the outer envelope of all conceivable constrained utility possibilities frontiers, which reflect all the conceivable institutional regimes in an economy and all the market imperfections that the economy may suffer from. By *institutional regimes* we mean all explicit tax and redistribution systems (from negative income tax systems to universal basic income to the regressive tax system currently in place), intellectual property regimes, job programs, education programs, but even social norms such as those related to charitable contributions. *Market imperfections* include all the market arrangements that differ from the Arrow-Debreu “optimal” benchmark, the conditions that ensure the Pareto efficiency of the market. As we noted earlier, the term embraces imperfections in information, competition, and risk and capital markets (including “missing” markets), but also rigidities in factor reallocation or in prices that determine how easily factors and products reallocate and which may be particularly important in the context of technological progress.

Changing any of these institutions or market imperfections has an effect on workers’ welfare. In general, it may be desirable to use a package of changes to all these institutions to ensure Pareto improvements after technological change has occurred. For instance, in section 14.4.3, we show that a combination of a change in the intellectual property regime *and* a change in capital taxation can ensure that an innovation is a Pareto improvement.

Finally, we also note that the possibility of achieving a Pareto improvement depends on how broadly we define the classes of individuals that are affected by an innovation. Our earlier example differentiated society, for simplicity, into two categories, workers and entrepreneurs. More generally,

different categories of workers, for example, skilled and unskilled workers, or workers in different sectors or tasks, are differentially affected by innovation. By the same token, different categories of entrepreneurs or innovators are differentially affected by innovation—for example, a given entrepreneur will generally be worse off if she is out-competed by another’s innovation. In the limit, if we consider the welfare of every single agent in the economy, clear Pareto improvements in a strict sense will be very difficult to find. As a result, our scope of analysis has to be targeted at the level that is relevant for the question at hand.

From both a political and a macroeconomic perspective, it is desirable that our welfare analysis focuses on groups that are sufficiently broad so that they matter for the political or economic equilibrium. It may also be useful to focus on groups that can be targeted with specific policy measures. Having said that, there is also a useful role for social safety nets that insure single individuals that lose out—for example, an innovator who goes broke because he was outpaced by a competitor.

14.2.6 Endogenous Technological Progress

A fifth and last point to emphasize is that there is no first welfare theorem for endogenous innovation. Generally speaking, the private returns to innovation in an economy differ from the social returns.¹¹

OBSERVATION 5) *The privately optimal choice of innovation may move the utility possibilities frontier inward, even if redistribution is costless.*

This implies that there may be benefits from intervening in the innovation process to generate Pareto improvements, for example, by making it less labor saving (see e.g. Stiglitz 2014b). Again, this does not specifically refer to advances in artificial intelligence—it will probably not apply to most examples of innovation in AI—but it is easy to think of examples where privately optimal innovation may generate Pareto deteriorations, for example, in the context of high-frequency trading in financial markets (see, e.g., Stiglitz 2014c).

14.2.7 Relationship between Technological Progress and Globalization

Many of the effects of technological change in general, and AI in particular, are similar to those of globalization. Indeed, globalization can be viewed as a change in technology, that of trading with the rest of the world.

11. It is hard to know who first had this insight. Certainly, Thomas Jefferson, America’s third president, recognized it when he said that knowledge is like a candle: when it lights another, the light of the first candle is not diminished. In the economics literature, it was clearly articulated by Arrow (1962) and Stiglitz (1987a). For a more recent statement of why social and private returns to innovation differ, see Stiglitz and Greenwald (2015). These results hold regardless of the intellectual property regime. Poorly designed intellectual property regimes can (and do) impair innovation. For a simple theoretical model, see Stiglitz (2014a); for empirical evidence, see Williams (2010).

In particular, trade of advanced countries with developing countries is “labor saving” (in the sense of Hicks): the demand for unskilled workers, or workers in general, decreases, at any given wage, implying that while the production possibilities curve moves out, and the utility possibilities curve may move out, the new equilibrium entails workers being worse off, as in figures 14.1 and 14.2. (In the absence of good risk markets, as we noted, everyone can be worse off, as in figure 14.3). Thus, the issue of whether globalization is welfare enhancing comes back to the question addressed in this chapter: is it possible to ensure, either through redistributive taxes or changes in institutions/rules, that workers are not made worse off. Again, there is a presumption that the gains to capital (or enterprises) could be taxed to provide the requisite redistributions.¹²

As we discuss in greater detail below, one of the side effects of innovation and intellectual property rights (IPR) is the creation of market power, resulting in ex post inefficient outcomes. Similarly, one of the consequences of globalization is to weaken the market power of workers. This is important because there is ample evidence that labor markets are far from perfectly competitive. The requisite compensation and/or offsetting changes in institutional rules to ensure that globalization represents a Pareto improvement may thus have to be all the greater.

14.3 Technological Progress and Channels of Inequality

There are two main channels through which technological progress may affect the distribution of resources and thus inequality: first, through the surplus earned by innovators and second, through effects on other agents in the economy.

14.3.1 Surplus Earned by Innovators

Technology is an information good, which implies that it is *nonrival*, but it may be *excludable*. Nonrival means that information can be used without being used up—in principle, many economic actors could use the same technology at the same time. If information about an innovation is widely shared, it can be used by all of society and provide welfare benefits to anybody who uses it. The excludable nature of information means, however, that others can be prevented from either obtaining or using a technology, for example, by withholding it from the public (e.g., as a business secret) or

12. Although a given country that opens up to trade is always made better off in a first-best world, ensuring a global Pareto improvement after a country reduces its trade barriers may be even more difficult than after technological progress has occurred, since changes in trade barriers affect international terms of trade and lead to redistributions across all other countries that can only be undone via cross-country transfers. (See, e.g., Korinek 2016). Furthermore, within each country, gains from trade inherently require changes in relative prices, which means that large redistributions are even more likely than in the case of technological progress.

by using social institutions such as intellectual property rights (e.g., copyrights or patents). This excludability may provide innovators with market power that enables them to charge a positive price for the innovation and earn a surplus.

Society faces a difficult trade-off in determining how to engineer the optimal level of innovation. In a first-best world, there are no agency problems in the process of innovation, and an optimal solution would be for the public to fund innovations and make them freely available to all (see, e.g., Arrow 1962). In fact, this model of financing innovation is common for basic research and has given rise to some very significant innovations in history, including the invention of the internet. A closely related solution is the production of innovations for nonpecuniary rewards, such as, for example, the prevalence of open source technology, which is widespread in the context of software and even artificial intelligence.¹³

However, in many circumstances, private agents are superior in producing innovation, and when they fund innovation, they expect to earn a return. The surplus earned by innovators then plays an important economic role because it rewards innovators for what they accomplish—it represents the economic return to innovation activity. However, as a result there is generally some market power associated with innovations, especially when there is a system of IPR in place, and this generally leads to inefficiencies compared to the first-best allocation in which innovations are distributed as public goods.¹⁴

We distinguish the following two cases, which determine whether innovators earn rents, that is, payoffs in excess of the cost of their innovative activity:

First, if entry into innovation activity is restricted, then the surplus or net income earned by innovators is generally greater than the costs of innovation activity. A natural example of such restrictions is when only a small number of people are endowed with special skills that enable them to innovate. These innovators then earn rents based on their exclusive abilities.

Restrictions to innovative activity may also arise from market structure: in markets with Bertrand competition, the first entrant who develops a costly innovation may enjoy a monopoly position because any potential competitor knows that if she enters, the incumbent will cut prices to marginal cost so that she cannot recoup the investment into an innovation.¹⁵

13. This approach relies on individuals or companies that are willing to innovate in exchange for nonpecuniary rewards such as prestige or, alternatively, on a calculated decision that providing free technology will steer potential customers or employees toward an innovator's platform, as seems to be the case in the field of AI.

14. For discussions of the merits of alternative ways of funding and incentivizing innovation, see Dosi and Stiglitz (2014), Baker, Jayadev, and Stiglitz (2017), Stiglitz (2008), and Korinek and Ng (2017).

15. See Stiglitz (1987b) and Dasgupta and Stiglitz (1988). When the number of firms is limited and there is Cournot competition, there will also be rents associated with innovation. For more general theoretical discussions of industrial structure and innovation, see Dasgupta and Stiglitz (1980a 1980b) and Stiglitz and Greenwald (2015, ch. 5).

Second, if innovative activity is contestable, that is, if there is a sufficiently large set of potential innovators with equal skills, then the expected rents to innovative activity are competed down to zero, that is, the marginal entrant into innovative activity is indifferent between innovating or not.¹⁶ However, given that the payoffs to innovation are highly stochastic, there will be winners and losers *ex post*. In the context of new technologies, the distribution of payoffs seems to be increasingly skewed, with a small number of entrepreneurs earning gigantic payoffs and the vast majority earning little in return for their efforts. This gives rise to significant inequality even among innovators.¹⁷

In either case, the returns earned by an innovator may not correspond closely to the social returns to the innovation; in particular, some of the returns may reflect the capture of profits that would otherwise have gone to other entrepreneurs.

Policies to Share the Surplus of Innovators

There is a growing consensus that one of the sources of the growth of inequality is the growth of rents, including the rents that innovators earn in excess of the cost of innovation (see e.g. Korinek and Ng 2018). Taxing and redistributing such rents has an important role in ensuring that AI and other advances in technology are Pareto improving. Also, anti-trust policies may lower such rents, ensuring that the benefits of innovations are more widely shared, as more competition lowers consumer prices from which all benefit. From the perspective of low wage workers who lose from innovation, targeted expenditure programs financed by high rent taxes may be of greater benefit than the lowering of prices, the benefit of which will go disproportionately to those who have high spending power.

Moreover, changes in intellectual property rights (IPRs) affect who receives the benefits of innovation—and thus the “incidence” of innovation, since IPRs are instrumental in providing extended market power to innovators.¹⁸

Additionally, public research—with government or the public at large appropriating the returns, rather than allowing private firms to do so—together with stronger competition policies, might reduce the scope for

16. Given the difficulty of predicting the success of innovative activity or of even assigning success probabilities, it is questionable how efficiently this mechanism works in practice. For example, there may be excessive entry because of overoptimism by some potential entrepreneurs, or there may be insufficient entry because of imperfect insurance markets for risk averse entrepreneurs. If some are better at innovating than others (and know it), then these individuals will enjoy inframarginal innovation rents (on average.)

17. If there are diminishing returns to the allocation of resources to innovation, the efficient resource allocation will entail there being rents associated with innovation. Who captures these rents will be affected by the institutional (including tax) structure. In the absence of adequate mechanisms for the public capturing those rents, there may be excessive investments in innovation, a standard result in the economics of the commons.

18. Especially when there is Bertrand competition, the benefits of innovation may be quickly shared with consumers upon the termination of patents.

monopolies capturing large fractions of the returns to innovation, and thus enhance the likelihood that AI will be Pareto improving.

Workers may also note that the innovations that ultimately led to AI—including those created by private entrepreneurs—build on significant public support. Society as a whole, but not necessarily this generation of innovators, paid for this knowledge, and should therefore share in the surplus generated by the innovation. One proposal to ensure that workers share in the benefits of innovation—and are less likely to lose from it—is to give workers shares in enterprises to ensure that their welfare goes up in tandem with that of shareholders/innovators as a whole.

14.3.2 Effects on Others

Innovation also leads to large redistributions among others in the economy who are not directly involved in the process of innovation, for example, workers who experience a sudden increase or decline in the demand for their labor. These redistributions can thus be viewed as externalities from innovation, and they are one of the main reasons why innovation raises concerns about inequality. We distinguish two categories of such externalities, pecuniary and nonpecuniary externalities. We discuss both in detail in the following.

Pecuniary Externalities: Price and Wage Changes

Among the most prominent implications of technological change is that it affects the prices of factors of production (including wages) and of produced goods. Hicks (1932) already observed that innovations generally change the demand for factors and will, in equilibrium, lead to factor price changes, especially changes in wages. The price and wage changes that result from innovations represent pecuniary externalities. Traditional general equilibrium theory, following Arrow and Debreu, emphasized that pecuniary externalities are fully consistent with Pareto efficiency. However, the benchmark of Pareto efficiency is blind to the distribution of income. Even if the equilibrium reached after an innovation is Pareto efficient, the pecuniary externalities lead to redistributions and imply that there are winners and losers.¹⁹

If—as many technologists predict—artificial intelligence directly replaces human labor, the demand for human labor will go down, and so will wages. More generally, innovations typically reduce demand for specific types of labor with specific human capital. For example, self-driving cars will likely depress the wages of drivers, or radiology-reading AI may lower the wages of traditional radiologists. Conversely, AI has certainly led to an increase in demand for computer scientists and has greatly increased their wages, in par-

19. Greenwald and Stiglitz (1986) and Geanakoplos and Polemarchakis (1986) demonstrated that pecuniary externalities also matter for efficiency when there are market imperfections such as imperfect information and incomplete markets; market equilibrium will as a result not even be Pareto efficient.

ticular in subfields that are directly related to AI. Since AI is a general purpose technology, there are reasons to believe that advances in AI will reverberate throughout many different sectors and lead to significant changes in wages throughout the economy in coming decades. Similar arguments can be made about the demand for and the value of different types of specific capital, as well as the demand for and prices of particular products.

Even though there are frequently losers, technological progress by definition shifts out the production possibilities frontier. This implies that the total dollar gain of the winners of progress exceeds the dollar loss of the losers.²⁰ In section 14.4 below, we will use this property of technological progress to argue that under relatively broad conditions this should enable the redistribution that is necessary to ensure that innovation leads to a Pareto improvement: the gains that arise to some factor owners as a result of technological progress are excess returns that are like unearned rents and could be taxed away without introducing distortions into the economy.

Although we noted that pecuniary externalities are generally viewed as Pareto efficient, there are two reasons for why they are likely to be associated with inefficiency in practice. First, Greenwald and Stiglitz (1986) and Geanakoplos and Polemarchakis (1986) demonstrated that pecuniary externalities matter for efficiency when there are market imperfections such as imperfect information and incomplete markets; market equilibrium will as a result not be Pareto efficient. Compared to the benchmark of idealized insurance markets “behind the veil of ignorance” that we discussed in section 14.2, the pecuniary externalities from innovation are clearly inefficient. Additional market imperfections are likely to lead to additional inefficiencies. Second, if the pecuniary externalities from innovation give rise to the need for redistributive policies that are costly to perform, the policy response will generate additional inefficiencies.

Policies to Counter Wage Declines

Aside from lump sum transfers, there are a range of further policies to counter the wage declines that are experienced by workers who are displaced by machines, even for low-skill jobs. These include wage subsidies and earned income tax credits. If bargaining power in labor markets is biased toward employers, an increased minimum wage can also help ensure that no one who works full time is in poverty. Furthermore, ensuring high aggregate demand—and thus a low unemployment rate—also increases the bargaining power of workers and leads to higher wages.

Other policies aimed at increasing the demand for especially low-skill labor include any measures that raise the wages of workers that are substitutes, for example, higher wages in the public sector as well as an increase in

20. If lump sum transfers were feasible, the winners could compensate the losers. However, in the absence of such compensation, social welfare may be lower.

public investments and other public expenditures; all of these policies help to drive up wages in the economy more generally.

Policies that could be used to finance such measures include carbon taxes, which would encourage resource-saving innovation at the expense of labor-saving innovation. It would thus simultaneously address two of most serious global problems, global climate change and inequality.²¹

Furthermore, the elimination of tax deduction for interest and the imposition of a tax on capital would increase the cost of capital and induce more capital augmenting innovation rather than labor saving innovation.²²

Non-Pecuniary Externalities

Innovation may also generate nonpecuniary externalities on agents other than the innovator. Classic examples for this are technological externalities—for example, if an innovation produces public goods or generates or alleviates pollution. In markets that deviate from the Arrow-Debreu benchmark, a variety of nonpecuniary effects may arise: for example, innovation may affect quantities demanded, or the probability of buying or selling a good or factor, including the probability of being unemployed.

Some effects are such that they can be interpreted either as pecuniary or nonpecuniary externalities. For example, product innovations can be interpreted as a price changes—the price of the newly invented good changes from infinity to some positive value—or as a change in the price of the consumption services provided by the good. Alternatively, they can also be interpreted in a nonpecuniary manner by viewing a product (such as a smartphone) as providing a bundle of services to consumers that can only be bought in fixed proportion (e.g., since we cannot separately purchase different functions of the smartphone). In that view, an innovation represents a change in the structure of incomplete markets because it changes the bundle of consumption service available from a product. Similarly, changes in job quality can be interpreted by viewing each job as a vector of transactions that are only available in predetermined bundles, and the innovation changes the elements in the bundle that are available. It is well known that changes in the degree of market incompleteness for such bundles give rise to externalities (a specific application of Greenwald and Stiglitz 1986).

14.4 Worker-Replacing Progress and Redistribution

This section considers a stark form of technological progress that we term worker-replacing technological progress. We develop two simple models

21. As we noted above, there is no first fundamental welfare theorem for innovation, and indeed, there is a presumption that the market is biased toward labor-saving innovation relative to innovations directed toward “saving the planet.” (See Stiglitz 2014b.)

22. The allocation of resources to capital augmenting technological change depends on the after-tax share of capital. An increase in the relative cost of capital will increase the capital share if the elasticity of substitution is less than unity. Most of the empirical evidence suggests that this is the case.

to analyze the two channels generating inequality that we discussed in the previous section. In sections 14.4.1 and 14.4.2, we consider the pecuniary externalities (redistributions) generated by worker-replacing progress, both from a static and a dynamic perspective. In section 14.4.3., we focus on the distribution of the surplus accruing to innovators in a model in which the surplus is determined by the level of patent protection. Furthermore, in section 14.4.4 we discuss the implications of endogenous factor bias in technological progress.

14.4.1 Static Pecuniary Externalities of Worker-Replacing Progress

For sections 14.4.1 and 14.4.2, we consider the simple model of worker-replacing technological change of Korinek and Stiglitz (2017). We assume a production technology that combines capital and labor in a constant-returns-to-scale (CRS) function, but where labor consists of the sum of human and machine labor. Assuming that human and machine labor enter the production function additively means that they are perfect substitutes for each other. The details of the baseline model are presented in box 14.1.

We analyze three questions: What does worker-replacing technological change do to wages in the short run and in the long run? And what can policy do about it?

First, we look exclusively at the short run before any of the other factors have adjusted:

OBSERVATION 6) *Machine Labor and Factor Earnings (in the short run): adding a marginal unit of machine labor reduces human wages, but increases returns of complementary factors in a zero-sum manner.*

Intuitively, what happens if we add one unit of machine labor is that first, that unit will earn its marginal return, but second, there is also a redistribution from labor to capital, which now becomes relatively scarcer. The gains of capital are exactly the losses of the existing stock of labor.

The redistribution generated by technological progress can be thought of as a pecuniary externality, as we emphasized earlier. The income losses of wage earners and the income gains of other factors owners are inefficient compared to the first-best benchmark considered in section 14.2.1. In the given example, the owners of capital have obtained windfall gains but have not done anything to earn these higher return. A compensatory transfer from capital owners to workers simply undoes these windfall gains and leaves them equally well off as they were before.

More generally, adding machine labor creates a redistribution away from human labor toward complementary factors. This result holds for any CRS production function no matter what the complementary factor, for instance, whether it is capital or land or unskilled versus skilled labor or entrepreneurial rents. Policy can undo these redistributions by taxing windfall gains while leaving the price system to work at the margin. The result also holds for decreasing-returns-to-scale production functions if we interpret the profits

Box 14.1**Machine Labor and Factor Earnings**

Assume a constant-returns-to-scale production function that produces output Y by combining capital K with labor, consisting of the sum of human labor H and machine labor M :

$$Y = F(K, H + M).$$

In this formulation human labor and machine labor are perfect substitutes, so machine technology is what we call worker-replacing.

In the competitive equilibrium, the wage is determined by the marginal product of labor,

$$w = F_L.$$

Proposition 1: Machine Labor and Factor Earnings: adding a marginal unit of machine labor reduces human wages but increases the returns to capital in a zero-sum manner, in addition to increasing output by the marginal product of labor, which is equal to the wage.

Proof: Using Euler's Theorem, we rewrite the production function:

$$(H + M)F_L(\cdot) + KF_K(\cdot) = F(K, H + M).$$

We can now ascertain the effect of an additional unit of M :

$$F_L + (H + M)F_{LL} + KF_{KL} = F_L,$$

$$\text{or, simplified, } \underbrace{(H + M)F_{LL}}_{\text{decline in wage bill}} + \underbrace{KF_{KL}}_{\text{increase in return to K}} = 0.$$

Source: Korinek and Stiglitz (2017).

earned by the owner of the technology as compensation for the implicit factor “entrepreneurship,” which takes part in the zero-sum redistribution.

Let us also emphasize that taxes on previously accumulated factors that suddenly earn an unexpected excess return are nondistortionary. This means that at least in principle, there is a role for implementing costless redistribution and generating a Pareto improvement. (In practice, there are some natural caveats to this result. For example, it relies on the assumption that we can distinguish between previously installed capital that earns windfall gains and new capital that would be distorted if it were taxed.)

14.4.2 Dynamic Implications of Worker-Replacing Progress

In the longer run, worker-replacing technological change will lead to significant economic change. It implies that the biggest constraint on output—the scarcity of labor—is suddenly relaxed. As a result, greater amounts of complementary factors, here capital, are accumulated.

OBSERVATION 7) Machine Labor and Abundance of Labor: If not only capital, but also labor, is reproducible at sufficiently low cost, then the economy will grow exponentially in AK fashion, driven purely by factor accumulation, even in the absence of further technological change.

In Korinek and Stiglitz (2017), we describe the dynamics of this transition as machines made by machines get increasingly efficient or, equivalently, as the cost of producing machines decreases. We identify a singularity point at which it becomes cost effective for machines to start to fully replace human labor.²³ In the simplest case, when complementary factors such as capital adjust without friction, the human wage may actually be unchanged because capital K grows in proportion to effective labor ($H + M$) so that the marginal productivity of labor and the wage remain unchanged. In other words, investment is allocated between conventional machines and human-replacing robots in such a way that the return is equal to the intertemporal marginal rate of substitution. Under the assumption that workers only care about their absolute income, not their relative income, this outcome would not be too bad: in absolute terms, even though the human labor share would go to zero as an increasing fraction of the labor in the economy is performed by machines, workers are no worse off as a result of AI.

When factors are slow to adjust, the pattern of transition can be complex, with demand for human labor typically going down temporarily.²⁴ In general, the pattern of adjustment depends on how fast the capital stock versus the stock of labor adjust. (For example, if the capital stock rises in anticipation of an increased supply of machine labor in the future that has not yet materialized, then human wages may even go up at intermediate stages.)²⁵

However, the following observation describes that in the long run, workers are actually worse off as a result of machine labor if there are nonreproducible complementary factors that are in scarce supply, such as land or other natural resources.

23. This singularity captures the important economic aspects of what technologists such as Vernor Vinge (1993) or Ray Kurzweil (2005) call the technological singularity. A similar point is also made in Aghion, Jones, and Jones (2017).

24. Berg, Buffie, and Zanna (2018) shows that it may actually take decades for the economy's complementary capital stock to adjust after major revolutions in labor-saving technology.

25. This assumes that capital is "putty-putty," that is, that capital investments made before AI arrives are equally productive after AI, as would be the case if humans and robots were in fact identical.

OBSERVATION 8) *Machine Labor and Return of Scarcity: if there are non-reproducible complementary factors, they eventually limit growth; human real wages fall, and the owners of nonreproducible factors absorb all the rents.*

Intuitively, as the supply of effective labor proliferates due to the introduction of machine labor, agents in the economy will compete for scarce nonreproducible resources like land, driving up their price.

A similar argument holds for nonreproducible consumption goods: even if all factors in the production process are reproducible so that productive output in the economy exhibits AK-style growth and workers' product wages remain unchanged, competition for fixed resources that are part of their consumption basket, such as land used for housing, may lead workers to eventually be worse off. This may be particularly important in urban settings where, say, economic activity occurs at the center. Rich rentiers may occupy the more desirable locations near the center, with workers having to obtain less expensive housing at the periphery, spending more time commuting. The advent of AI will thus lower their utility.

However, just as in the earlier case, at the margin, the redistribution from workers to nonreproducible factor owners is zero sum. Since taxes on nonreproducible factors are by definition nondistortionary, there is scope for nondistortionary redistribution.

OBSERVATION 9) *Nonreproducible Factors and Pareto Improvements: so long as nondistortionary taxes on factor rents are feasible, labor-replacing innovation can be a Pareto improvement.*

14.4.3 Redistributing the Innovators' Surplus via Changes in Institutions

If outright redistribution is not feasible or limited, there may be other institutional changes that result in market distributions that are more favorable to workers. For example, intervention to steer technological progress may act as a second-best device.

In this section, we provide an example in which a change in intellectual property rights—a shortening of the term of patent protection—effectively redistributes some of the innovators' surplus to workers (consumers) to mitigate the pecuniary externalities on wages that they experience, with the ultimate goal that the benefits of the innovation are more widely shared. If an innovation results in a lower cost of production, then the innovator enjoys the benefits of the innovation in the form of higher profits during the life of the patent; but after the expiration of the patent, society enjoys the benefits in terms of lower prices. The trade-off is that shortening the life of the patent may reduce the pace of innovation. But in the spirit of the theory of the second-best, there is generally an "optimal" patent life, in which there is still some innovation, but in which the well-being of workers is protected.

Box 14.2

Intellectual Property Regime and Redistribution

Consider an economy with a unit mass of workers $H = 1$, in which the capital stock supplied each period $K(\tau)$ is a function solely of a distortionary capital tax τ , the proceeds of which are distributed to workers, and the effective stock of machine labor $M(z)$ is an increasing function of patent life z .

A worker's total income I consists of her wage plus the revenue of the capital tax,

$$I = w + \tau K(\tau).$$

For any level of $M(z)$, we define $\tau(M)$ as the value of the capital tax that keeps workers just as well off as they were before the introduction of machine labor.

Proposition 1. As long as elasticity of capital supply is not too large, we can always increase z from $z = 0$ and compensate workers by raising the capital tax τ .

Steady-State Dynamics

Consider an intertemporal setting in which the growth rate $g = g(z, \tau)$ is a function of the length of the patent z and the tax rate τ , by which we now denote the tax rate on innovators. Assume that the share of output that is invested is a function of the growth rate ($i(g)$) and that the fraction of output not spent on investment that is appropriated by the innovator is $b(z, \tau)$. In steady state, the present discounted value of the income of workers can be approximated as

$$\text{PDV} = (1 - i(g)) [1 - (1 - \tau)b(z, \tau)] / (r - g),$$

where r is the discount rate. If we choose $\{z, \tau\}$ to maximize the PDV, in general, the optimum will not be a corner solution in which any innovation hurts workers.

Proposition 2. In general, the optimal $\{z^*, \tau^*\}$ entails $g > 0$.

It is easy to write down sufficient conditions under which Proposition 2 holds: setting τ^* equal to zero, all that we require is that $|g_z|$ is not too large relative to $|b_z|$.

With network externalities the innovator may be able to maintain a dominant position even after the end of the patent, and may continue to earn the surplus from her innovation. With taxes on monopoly profits, it should be possible to ensure that the innovations are Pareto improving and that even human worker-replacing technological change can improve the well-being of workers.

14.4.4 Factor-Biased Technological Change

So far, we have simply assumed that technological change—the introduction of AI—is worker replacing. But advances in technology also make some machines more productive and others obsolete, affecting the (marginal) return to traditional capital.²⁶ It is thus useful to think of the world as having three groups: capitalists, workers, and innovators. Intellectual property rights (and antitrust laws) determine the returns to innovators, but the nature of technological change in a competitive market determines the division between workers and capitalists.

A long-standing literature, going back to Kennedy (1964), von Weizacker (1966), and Samuelson (1965), describes the endogenous determination of the factor bias of technological progress.²⁷ The central result is that as the share of labor becomes smaller, the bias shifts toward capital-augmenting technological progress. If the world works as these models suggest, this should limit the decline in the share of labor (at least in a stable equilibrium) and in inequality.²⁸ As the share of labor decreases, the incentive to produce worker-replacing innovation such as AI decreases. But the relevant discounted future wage share near the point of singularity—the point where it is cost effective for machines to fully replace human labor and produce more machines all by themselves—may be sufficiently great that there is nonetheless an incentive to pass the point of singularity.

Let us assume that land becomes the binding constraint once human labor is fully replaceable by machine labor. In that case, provided the elasticity of substitution between land and the other production factors—capital cum labor—is less than unity, the share of land increases over time, generating the result (analogous to that where labor is the binding constraint in the standard literature) that in the long run, all technological progress is land augmenting. If the production function is constant returns to scale in land,

26. As we noted earlier, IA (intelligence assisting) innovation may increase the productivity of humans, and thus increase the demand for humans if the elasticity of substitution is less than unity.

27. Important contributions were also made by Drandakis and Phelps (1965). More recently, there has been some revival of the literature, with work of Acemoglu (2002), Stiglitz (2006, 2014b) and Acemoglu and Restrepo (2018), among others.

28. One can describe dynamics with standard wage-setting mechanisms. The system is stable so long as the elasticity of substitution between factors is less than unity (Acemoglu 1998; Stiglitz 2006, 2014b).

labor (including machine labor) and traditional capital, then the long-run rate of growth is determined by the pace of land-augmenting technological change.

Role of the Service Sector

Currently, progress in AI focuses on certain sectors of the economy, like manufacturing. Partly because of the resulting lower cost of manufacturing, and partly because of the shape of preferences, the economy is evolving toward a service-sector economy. (If there is differential productivity across sectors, and the elasticity of demand for the innovation sector is not too high, then production factors will move out of that sector into other sectors. This is even more so if preferences are nonhomothetic, for example, demand for food and many manufactured goods having an income elasticity less than unity.) Among the key service sectors are education, health, the military, and other public services. The value of those services is in large part socially determined, that is, by public policies not just a market process. If we value those services highly—pay good wages, provide good working conditions, and create a sufficient number of jobs—this will limit increases in the inequality of market income. Governments typically play an important role in these sectors, and their employment policies will thus play an important role in the AI transition. Many of these service-sector jobs have limited skill requirements. However, higher public-sector wages will—through standard equilibrium effects—also raise wages in the private sector, will improve the bargaining position of workers, and will result in such jobs having higher “respect.” All of this will, of course, require tax revenues. If the elasticity of entrepreneurial services is low, for example, if entrepreneurs are driven partly by nonpecuniary motives, we can impose high taxes to finance these jobs.

14.5 Technological Unemployment

Unemployment is one of the most problematic societal implications of technological progress—new technology often implies that old jobs are destroyed and workers need to find new jobs. Economists, of course, understand the “lump-of-labor fallacy”—the false notion that there is a fixed number of jobs, and that automating a given job means that there will forever be fewer jobs left in the economy. In a well-functioning economy, we generally expect that technological progress creates additional income, which in turn can support more jobs.

However, there are two sound economic reasons for why technological unemployment may arise: first, because wages do not adjust for some structural reason, as described, for example, by efficiency wage theories, and second as a transition phenomenon. The two phenomena may also interact

in important ways, for example, when efficiency wage considerations slow down the transition to a new equilibrium. We discuss the two categories in turn in the following subsections.

The unemployment implications are especially problematic when technological progress is labor saving, which—by definition—requires that either wages have to fall or that other complementary factors like capital have to adjust enough for labor market equilibrium to be restored at or above the historic wage.

14.5.1 Efficiency Wage Theory and Nonadjustment of Wages

The first category of technological unemployment arises when wages do not adjust for structural reasons. Efficiency wage theory emphasizes that productivity depends on wages and so employers may have reasons to pay wages above the market-clearing level. The original efficiency wage paper (Stiglitz 1969) noted one of the reasons for this: that income disparities can weaken worker morale. Akerlof and Yellen (1990) have formalized this into the “fair wage hypothesis.”

If fairness considerations are significant enough, and workers think that a decrease in their wages is “unfair” (e.g., because the income of entrepreneurs increases so entrepreneurs could easily “afford” pay increases), it means that the scope of labor-saving progress that shifts the utility possibilities curve out *without redistributions* is very limited. Similar results hold if workers’ well-being and efforts are related to relative incomes. The new utility possibilities curve may lie outside the old one to the “north” of E_0 , that is, there is scope for a Pareto improvement in principle; but it may lie inside of the old utility possibilities curve near E_1 , that is, the utility possibilities of workers decrease for a given level of utility of entrepreneurs because workers reduce their effort so much that the effective labor supply declines—any gains from technology are more than offset by increased shirking. Shapiro and Stiglitz (1984) emphasize that paying a wage above the market-clearing level reduces shirking, leading to unemployment.

An even more daunting example of efficiency wages may arise if automation continues and the marginal product of labor for low-skill workers falls below their cost of living at what they view as their basic subsistence (even if they exert their best effort). Unless basic social services are provided to such workers, a nutritional efficiency wage model applies in that case, similar to what Stiglitz (1976) described for developing countries: employers could not pay a market-clearing wage because they know that this would be insufficient for their employees to provide for themselves and remain productive.²⁹ We will follow up on this theme in the final section of our chapter.

In traditional efficiency wage models, the unemployment effects of effi-

29. Even worse outcomes could emerge in the presence of imperfect capital markets, if expenditures on health and nutrition at one date affect productivity at later dates.

ciency wages are permanent, part of the long-run equilibrium. For example, if technological change leads to greater inequality (or better information about the existing level of inequality), morale effects and the resulting efficiency wage responses imply that the equilibrium level of unemployment rises.

However, efficiency wage arguments may also contribute to slowing down the transition to a new equilibrium after an innovation, as we will explore subsection 14.5.2.

Minimum Wages and Nonadjustment of Wages

An alternative reason why wages may not adjust to the market-clearing level are minimum wage laws. Basic economics implies that there will be unemployment if wages are set to an excessive level. Although this is a theoretical possibility, recent experience in the United States has repeatedly shown that modest increases in minimum wages from current levels have hardly any employment effects but raise the income of minimum wage workers, which may have positive aggregate demand effects since low-income workers have a high marginal propensity to consume (see, e.g., Schmitt 2013). From an economic theory perspective, these observations are possible because wages are not determined in a purely Walrasian manner—there is a significant amount of bargaining involved when prospective employers and employees match—and increases in minimum wages substitute for the lacking bargaining power of workers (see, e.g., Manning 2011).

14.5.2 Technological Unemployment as a Transition Phenomenon

The second category of technological unemployment is as a transition phenomenon, that is, when technological change makes workers redundant at a faster pace than they can find new jobs or that new jobs are created. This phenomenon was already observed by Keynes (1932). It is well understood that there is always a certain “natural” or “equilibrium” level of unemployment as a result of churning in the labor market. In benchmark models of search and matching to characterize this equilibrium level of unemployment (see Mortensen and Pissarides 1994, 1998), employment relationships are separated at random, and workers and employers need to search for new matches to replace them. The random shocks in this framework can be viewed as capturing, in reduced form, phenomena such as life cycle transitions but also technological progress in individual firms. In this view, an increase in the pace of technological progress corresponds to a higher job separation rate and results in a higher equilibrium level of unemployment.

The transition may be especially prolonged if technology implies that the old skills of workers become obsolete and they need to acquire new skills and/or find out what new jobs match their skills (see, e.g., Restrepo 2015).

Even if in the long run workers adjusted to AI, the transition may be difficult. Artificial intelligence will impact some sectors more than others,

and there will be significant job dislocation. As a general lesson, markets on their own are not good at structural transformation. Often, the pace of job destruction is greater than the pace of job creation, especially as a result of imperfections in capital markets, inhibiting the ability of entrepreneurs to exploit quickly new opportunities as they are opened up.

The Great Depression as an Example of Transitional Unemployment

The Great Depression can be viewed as being caused by rapid pace of innovation in agriculture (see Delli Gatti et al. 2012a). Fewer workers were needed to produce the food that individuals demanded, resulting in marked decline in agriculture prices and income, leading to a decline in demand for urban products. In the late 1920s, these effects became so large that long-standing migration patterns were reversed.

What *might* have been a Pareto improvement turned out to be an immiserizing technological change, as both those in the urban and rural sector suffered.

The general result is that noted earlier: with mobility frictions and rigidities (themselves partly caused by capital market imperfections, as workers in the rural sector couldn't obtain funds to obtain the human capital required in the urban sector and to relocate) technological change can be welfare decreasing. The economy can be caught, for an extensive period of time, in a low-level equilibrium trap, with high unemployment and low output.

In the case of the Great Depression, government intervention (as a by-product of World War II) eventually enabled a successful structural transformation: the intervention was not only a Keynesian stimulus, but facilitated the move from rural farming areas to the cities where manufacturing was occurring at the time and facilitated the retraining of the labor force, helping workers acquire the skills necessary for success in an urban manufacturing environment, which were quite different from those that ensured success in a rural, farming environment. It was, in this sense, an example of a successful industrial policy.

There are clear parallels to the situation today in that a significant fraction of the workforce may not have the skills required to succeed in the age of AI.

Transitional Efficiency Wage Theory

Efficiency wage arguments may also slow down the transition to a new equilibrium after technological progress. For example, if worker morale depends on last period's wages, it may be difficult to reduce wages to the market-clearing level after a labor-saving innovation, and unemployment may persist for a long time.³⁰

30. In the limiting case, employers may simply keep wages fixed to avoid negative morale effects, and unemployment would persist forever—or until some offsetting shock occurs.

14.5.3 Jobs and Meaning

The potentially widespread destruction of jobs can have large human consequences that go beyond just economics because jobs provide not only income but also other mental services such as meaning, dignity, and fulfillment to humans. Whether this is a legacy of our past, and whether individuals could find meaning in other forms of activities, mental or physical, is a matter of philosophical debate.

If workers derive a separate benefit from work in the form of meaning, then job subsidies are a better way of ensuring that technological advances are welfare enhancing than simply providing lump sum grants (e.g., through the provision of a universal basic income), as some are suggesting in response to the inequalities created by AI.

This discussion is, of course, a departure from the usual neoclassical formulation, where work only enters negatively into individual's well-being. There are some that claim that individuals' deriving dignity and meaning from work is an artifact of a world with labor scarcity. In a workerless AI world, individuals will have to get their identity and dignity elsewhere, for example, through spiritual or cultural values. The fact that most humans can find a meaningful life after retirement perhaps suggests that there are good substitutes for jobs in providing meaning.

14.6 Longer-Term Perspectives: AI and the Return of Malthus?

There is a final point that is worth discussing in a chapter on the implications of artificial intelligence for inequality. This point relates to a somewhat longer-term perspective. Currently, artificial intelligence is at the stage where it strictly dominates human intelligence in a number of specific areas, for instance playing chess or Go, identifying patterns in x-rays, driving, and so forth. This is commonly termed *narrow* artificial intelligence. By contrast, humans are able to apply their intelligence across a wide range of domains. This capacity is termed *general* intelligence.

If AI reaches and surpasses human levels of general intelligence, a set of radically different considerations apply. Some techno-optimists predict the advent of general artificial intelligence for as early as 2029 (see Kurzweil 2005), although the median estimate in the AI expert community is around 2040 to 2050, with most AI experts assigning a 90 percent probability to human-level general artificial intelligence arising within the current century (see Bostrom 2014). A minority believes that general artificial intelligence will never arrive. However, if human-level artificial general intelligence is reached, there is broad agreement that AI would soon after become super-intelligent, that is, more intelligent than humans, since technological progress would likely accelerate, aided by the intelligent machines. Given these

predictions, we have to think seriously about the implications of artificial general intelligence for humanity and, in the context of this chapter, for what it implies for our economy as well as for inequality.

Assuming that our social and economic system will be maintained upon the advent of artificial general intelligence and superintelligence,³¹ there are two main scenarios. One scenario is that man and machine will merge, that is, that humans will “enhance” themselves with ever more advanced technology so that their physical and mental capabilities are increasingly determined by the state of the art in technology and AI rather than by traditional human biology (see, e.g., Kurzweil, 2005). The second scenario is that artificially intelligent entities will develop separately from humans, with their own objectives and behavior (see, e.g., Bostrom 2014; Tegmark 2017). As we will argue below, it is plausible that the two scenarios might differ only in the short run.

First Scenario: Human Enhancement and Inequality

The scenario that humans will enhance themselves with machines may lead to massive increases in human inequality, unless policymakers recognize the threat and take steps to equalize access to human enhancement technologies.³² Human intelligence is currently distributed within a fairly narrow range compared to the distance between the intelligence of humans and that of the next-closest species. If intelligence becomes a matter of ability to pay, it is conceivable that the wealthiest (enhanced) humans will become orders of magnitude more productive—“more intelligent”—than the unenhanced, leaving the majority of the population further and further behind. In fact, if intelligence enhancement becomes possible, then—unless preemptive actions are taken—it is difficult to imagine how to avoid such a dynamic. For those who can afford it, the incentive to purchase enhancements is great, especially since they are in competition with other wealthy humans who may otherwise leapfrog them. This is even more so in an economy which is, or is perceived to be, a winner-take-all economy and/or in which well-being is based on relative income. Those who cannot afford the latest technology will have to rely on what is in the public domain, and if the pace of innovation increases, the gap between the best technology and what is publicly available will increase.

A useful analogy is to compare human enhancement technology to health care—technology to *maintain* rather than *enhance* the human body. Dif-

31. Researchers who work on the topic of AI safety point out that there is also a risk of doomsday scenarios in which a sufficiently advanced artificial intelligence eradicates humanity because humans stand in the way of its goals. See, for example, Bostrom (2014) who elaborates on this using the example of a “paperclip maximizer”—an AI that has been programmed to produce as many paperclips as possible, without regard for other human goals, and who realizes that humans contain valuable raw materials that should better be transformed into paperclips.

32. In many respects, the issues are parallel to those associated with performance-enhancing drugs. In sports, these have been strictly regulated, but in other arenas, they have not.

ferent countries have chosen significantly different models for how to provide access to health care, with some regarding it as a basic human right and others allocating it more according to ability to pay. In the United States, for example, the expected life spans of the poor and the wealthy have diverged significantly in recent decades, in part because of unequal access to health care and ever more costly new technologies that are only available to those who can pay. The differences are even starker if we look at humanity across nations, with the expected life span in the richest countries being two-thirds longer than in the least developed countries (see, e.g., UN 2015). Like with health care, it is conceivable that different societies will make significantly different choices about access to human enhancement technologies.

Once the wealthiest enhanced humans have separated sufficiently far from the unenhanced, they can effectively be considered as a separate species of artificially intelligent agents. To emphasize the difference in productivities, Yuval Harari (2017) has dubbed the two classes that may result “the gods” and “the useless.” In that case, the long-run implications of our first scenario coincide with the second scenario.

Second Scenario: Artificially Intelligent Agents and the Return of Malthus

We thus turn to the scenario that artificially intelligent entities develop separately from regular (or unenhanced) humans. One of the likely characteristics of any sufficiently intelligent entity—no matter what final objectives are programmed into it by evolution or by its creator—is that it will act by pursuing intermediate objectives or “basic drives” that are instrumental for any final objective (Omohundro 2008). These intermediate objectives include self-preservation, self-improvement, and resource accumulation, which all make it likelier and easier for the entity to achieve its final objectives.

It may be worthwhile pursuing the logic of what happens if humans do not or cannot assert ownership rights over artificially intelligent or superintelligent entities.³³ That would imply that sufficiently advanced AI is likely to operate autonomously.

To describe the resulting economic system, Korinek (2017) assumes that there are two types of entities, unenhanced humans and AI entities, which are in a Malthusian race and differ—potentially starkly—in how they are affected by technological progress. At the heart of Malthusian models is the notion that survival and reproduction requires resources, which are poten-

33. If humans and artificially intelligent entities are somewhat close in their levels of intelligence, it may still be possible for humans to assert ownership rights over the AI—in fact, throughout the history of mankind, those determining and exerting property rights have not always been the most intelligent. For example, humans could still threaten to turn off or destroy the computers on which AI entities are running. However, if the gap between humans and superintelligent AI entities grows too large, it may be impossible for humans to continue to exert control, just like a two-year-old would not be able to effectively exert property rights over adults.

tially scarce.³⁴ Formally, traditional Malthusian models capture this by describing how limited factor supplies interact with two related sets of technologies, a production and a consumption/reproduction technology: First, humans supply the factor labor, which is used in a production technology to generate consumption goods. Second, a consumption/reproduction technology converts consumption goods into the survival and reproduction of humans, determining the future supply of the factor labor.

Throughout human history Malthusian dynamics, in which scarce consumption goods limited the survival and reproduction of humans, provided a good description of the state of humanity, roughly until when Malthus (1798) published his *Essay on the Principle of Population* to describe the resulting Iron Law of Population. Over the past two centuries, humanity, at least in advanced countries, was lucky to escape its Malthusian constraints: capital accumulation and rapid labor-augmenting technological progress generated by the Industrial Revolution meant that our technology to produce consumption goods was constantly ahead of the consumption goods required to guarantee our physical survival. Moreover, human choices to limit physical reproduction meant that the gains of greater productivity were only partly dissipated in increased population. However, this state of affairs is not guaranteed to last forever.

Korinek (2017) compares the production and consumption/reproduction technologies of humans and AI entities and observes that they differ starkly: On the production side, the factor human labor is quickly losing ground to the labor provided by AI entities, captured by the notion of *worker-replacing technological progress* that we introduced earlier. In other words, AI entities are becoming more and more efficient in the production of output compared to humans. On the consumption/reproduction side, the human technology to convert consumption goods such as food and housing into future humans has experienced relatively little technological change—the basic biology of unenhanced humans is slow to change. By contrast, the reproduction technology of AI entities—to convert AI consumption goods such as energy, silicon, aluminum into future AI—is subject to exponential progress, as described, for example, by Moore’s Law and its successors, which postulate that computing power per dollar (i.e., per unit of “AI consumption good”) doubles roughly every two years.³⁵

34. If AI directs its enhanced capabilities at binding resource constraints, it is conceivable that such constraints might successively be lifted, just as we seem to have avoided the constraints that might have been imposed by the limited supply of fossil fuels. At present, humans consume only a small fraction—about 0.1 percent—of the energy that earth receives from the sun. However, astrophysicists such as Tegmark (2017) note that according to the laws of physics as currently known, there will be an ultimate resource constraint on superintelligent AI given by the availability of energy (or, equivalently, matter, since $E = mc^2$) accessible from within our event horizon.

35. The original version of Moore’s Law, articulated by the cofounder of Intel, Gordon Moore (1965), stated that the number of components that can be fit on an integrated circuit (IC) would double every year. Moore revised his estimate to every two years in 1975. In recent

Taken together, these two dynamics imply—unsurprisingly—that humans may lose the Malthusian race in the long run, unless counteracting steps are taken, to which we will turn shortly. In the following paragraphs we trace out what this might entail and how we might respond to it. (Fully following the discussion requires a certain suspension of disbelief. However, we should begin by recognizing that machines can already engage in a large variety of economic transactions—trading financial securities, placing orders, making payments, and so forth. It is not a stretch of the mind to assume that they could in fact engage in all of what we now view as economic activities. In fact, if an outside observer from a different planet were to witness the interactions among the various intelligent entities on earth, it might not be clear to her if, for example, artificially intelligent entities such as Apple or Google control what we humans do [via a plethora of control devices called smartphones that we carry with us] or whether we intelligent humans control what entities such as Apple and Google do. See also the discussion in Turing [1950].) The most interesting aspects of the economic analysis concern the transition dynamics and the economic mechanisms through which the Malthusian race plays out.

In the beginning, those lacking the skills that are useful in an AI-dominated world may find that they are increasingly at a disadvantage in competing for scarce resources, and they will see their incomes decline, as we noted earlier. The proliferation of AI entities will at first put only modest price pressure on scarce resources, and most of the scarce factors are of relatively little interest to humans (such as silicon), so humanity as a whole will benefit from the high productivity of AI entities and from large gains from trade. From a human perspective, this will look like AI leading to significant productivity gains in our world. Moreover, any scarce factors that are valuable for the reproduction and improvement of AI, such as human labor skilled in programming, or intellectual property, would experience large gains.

As time goes on, the superior production and consumption technologies of AI entities imply that they will proliferate. Their ever-increasing efficiency units will lead to fierce competition over any nonreproducible factors that are in limited supply, such as land and energy, pushing up the prices of such factors and making them increasingly unaffordable for regular humans, given their limited factor income. It is not hard to imagine an outcome where the AI entities, living for themselves, absorb (i.e., “consume”) more and more of our resources.

Eventually, this may force humans to cut back on their consumption to the point where their real income is so low that they decline in numbers.

years, companies such as Intel have predicted that the literal version of Moore’s Law may come to an end over the coming decade, as the design of traditional single-core ICs has reached its physical limits. However, the introduction of multidimensional ICs, multicore processors and other specialized chips for parallel processing implies that a broader version of Moore’s Law, expressed in terms of computing power per dollar, is likely to continue for several decades to come. Quantum computing may extend this time span even further into the future.

Technologists have described several dystopian ways in which humans could survive for some time—ranging from uploading themselves into a simulated (and more energy-efficient) world,³⁶ to taking drugs that reduce their energy intake. The decline of humanity may not play out in the traditional way described by Malthus—that humans are literally starving—since human fertility is increasingly a matter of choice rather than nutrition. It is sufficient that a growing number of unenhanced humans decide that given the prices they face, they cannot afford sufficient offspring to meet the human replacement rate while providing their offspring with the space, education, and prospects that they aspire to.

One question that these observations bring up is whether it might be desirable for humanity to slow down or halt progress in AI beyond a certain point. However, even if such a move were desirable, it may well be technologically infeasible—progress may have to be stopped well short of the point where general artificial intelligence could occur. Furthermore it cannot be ruled out that a graduate student under the radar working in a garage will create the world's first superhuman AI.

If progress in AI cannot be halted, our description above suggests mechanisms that may ensure that humans can afford a separate living space and remain viable: because humans start out owning some of the factors that are in limited supply, if they are prohibited from transferring these factors, they could continue to consume them without suffering from their price appreciation. This would create a type of human “reservation” in an AI-dominated world. Humans would likely be tempted to sell their initial factor holdings, for two reasons: First, humans may be less patient than artificially intelligent entities. Second, superintelligent AI entities may earn higher returns on factors and thus be willing to pay more for them than other humans. That is why, for the future of humanity, it may be necessary to limit the ability of humans to sell their factor allocations to AI entities. Furthermore, for factors such as energy that correspond to a flow that is used up in consumption, it would be necessary to allocate permanent usage rights to humans. Alternatively, we could provide an equivalent flow income to humans that is adjusted regularly to keep pace with factor prices.³⁷

14.7 Conclusions

The proliferation of AI and other forms of worker-replacing technological change can be unambiguously positive in a first-best economy in which individuals are fully insured against any adverse effects of innovation, or if it is coupled with the right form of redistribution. In the absence of such

36. See, for example, Hanson (2016). In fact, Aguiar et al. (2017) document that young males with low education have already shifted a considerable part of their time into the cyber world rather than supplying labor to the market economy—at wages that they deem unattractive.

37. All of this assumes that the superintelligent AI entities don't use their powers in one way or another to abrogate these property rights.

intervention, worker-replacing technological change may not only lead to workers getting a diminishing fraction of national income, but may actually make them worse off in absolute terms.

The scope for redistribution is facilitated by the fact that the changes in factor prices create windfall gains on the complementary factors, which should make it feasible to achieve Pareto improvements. If there are limits on redistribution, the calculus worsens and a Pareto improvement can no longer be ensured. This may lead to resistance from those in society who are losing. As a result, there is a case for using as broad of a set of second-best policies as possible, including changes in intellectual property rights, to maximize the likelihood that AI (or technological progress more generally) generate a Pareto improvement.

Artificial intelligence and other changes in technology necessitate large adjustments, and while individuals and the economy more broadly may be able to adjust to slow changes, this may not be so when the pace is rapid. Indeed, in such situations, outcomes can be Pareto inferior. The more willing society is to support the necessary transition and to provide support to those who are “left behind,” the faster the pace of innovation that society can accommodate, and still ensure that the outcomes are Pareto improvements. A society that is not willing to engage in such actions should expect resistance to innovation, with uncertain political and economic consequences.

References

- Acemoglu, Daron. 1998. “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality.” *Quarterly Journal of Economics* 113 (4): 1055–89.
- . 2002. “Directed Technical Change.” *Review of Economic Studies* 69 (4): 781–809.
- Acemoglu, Daron, and Pascual Restrepo. 2018. “The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment.” *American Economic Review* 108 (6): 1488–542.
- Aghion, Philippe, Benjamin Jones, and Charles Jones. 2017. “Artificial Intelligence and Economic Growth.” NBER Working Paper no. 23928, Cambridge, MA.
- Aguiar, Mark, Mark Bils, Kerwin Kofi Charles, and Erik Hurst. 2017. “Leisure Luxuries and the Labor Supply of Young Men.” NBER Working Paper no. 23552, Cambridge, MA.
- Akerlof, George, and Janet Yellen. 1990. “The Fair Wage-Effort Hypothesis and Unemployment.” *Quarterly Journal of Economics* 105 (2): 255–83.
- 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.
- Baker, Dean, Arjun Jayadev, and Joseph E. Stiglitz. 2017. “Innovation, Intellectual Property, and Development: A Better Set of Approaches for the 21st Century.” AccessIBSA: Innovation & Access to Medicines in India, Brazil & South Africa.

- Barrat, James. 2013. *Our Final Invention: Artificial Intelligence and the End of the Human Era*. New York: St. Martin's Press.
- Berg, Andrew, Edward F. Buffie, and Luis-Felipe Zanna. 2018. "Should We Fear the Robot Revolution? (The Correct Answer is Yes)." *Journal of Monetary Economics* 97. www.doi.org/10.1016/j.jmoneco.2018.05.012.
- Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford: Oxford University Press.
- Dasgupta, Partha, and Joseph E. Stiglitz. 1980a. "Uncertainty, Industrial Structure and the Speed of R&D." *Bell Journal of Economics* 11 (1): 1–28.
- . 1980b. "Industrial Structure and the Nature of Innovative Activity." *Economic Journal* 90 (358): 266–93.
- . 1988. "Potential Competition, Actual Competition and Economic Welfare." *European Economic Review* 32:569–77.
- Dávila, Eduardo, and Anton Korinek. 2018. "Pecuniary Externalities in Economies with Financial Frictions." *Review of Economic Studies* 85 (1): 352–95.
- Delli Gatti, Domenico, Mauro Gallegati, Bruce C. Greenwald, Alberto Russo, and Joseph E. Stiglitz. 2012a. "Mobility Constraints, Productivity Trends, and Extended Crises." *Journal of Economic Behavior & Organization* 83 (3): 375–93.
- . 2012b. "Sectoral Imbalances and Long-run Crises." In *The Global Macro Economy and Finance*, edited by Franklin Allen, Masahiko Aoki, Jean-Paul Fitoussi, Nobuhiro Kiyotaki, Robert Gordon, and Joseph E. Stiglitz. International Economic Association Series. London: Palgrave Macmillan.
- Dosi, Giovanni, and Joseph E. Stiglitz. 2014. "The Role of Intellectual Property Rights in the Development Process, with Some Lessons from Developed Countries: An Introduction." In *Intellectual Property Rights: Legal and Economic Challenges for Development*, edited by Mario Cimoli, Giovanni Dosi, Keith E. Maskus, Ruth L. Okediji, Jerome H. Reichman, and Joseph E. Stiglitz, 1–53. Oxford: Oxford University Press.
- Drandakis, Emmanuel, and Edmund Phelps. 1965. "A Model of Induced Invention, Growth and Distribution." *Economic Journal* 76:823–40.
- Fehr, Ernst, and Klaus M. Schmidt. 2003. "Theories of Fairness and Reciprocity—Evidence and Economic Applications." In *Advances in Economics and Econometrics*, Econometric Society Monographs, Eighth World Congress, vol. 1, edited by Mathias Dewatripont, Lars Peter Hansen, and Stephen J Turnovsky, 208–57. Cambridge: Cambridge University Press.
- Financial Crisis Inquiry Commission. 2011. "The Financial Crisis Inquiry Report." Final Report of the National Commission. January. <http://www.gpoaccess.gov/fcic/fcic.pdf>.
- Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The Future of Employment: How Susceptible Are Jobs to Computerisation?" *Technological Forecasting and Social Change* 114:254–80.
- Geanakoplos, John, and Herakles Polemarchakis. 1986. "Existence, Regularity, and Constrained Suboptimality of Competitive Allocations When the Asset Market Is Incomplete." In *Uncertainty, Information and Communication: Essays in Honor of KJ Arrow*, edited by W. Heller, R. Starr, and D. Starrett, 65–96. Cambridge: Cambridge University Press.
- Gordon, Robert. 2016. *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*. Princeton, NJ: Princeton University Press.
- Greenwald, Bruce, and Joseph E. Stiglitz. 1986. "Externalities in Economics with Imperfect Information and Incomplete Markets." *Quarterly Journal of Economics* 101 (2): 229–64.
- Groshen, Erica L., Brian C. Moyer, Ana M. Aizcorbe, Ralph Bradley, and David M. Friedman. 2017. "How Government Statistics Adjust for Potential Biases from

- Quality Change and New Goods in an Age of Digital Technologies: A View from the Trenches.” *Journal of Economic Perspectives* 31 (2): 187–210.
- Guzman, Martin, and Joseph E. Stiglitz. 2016a. “Pseudo-Wealth and Consumption Fluctuations.” NBER Working Paper no. 22838, Cambridge, MA.
- Guzman, Martin, and Joseph E. Stiglitz. 2016b. “A Theory of Pseudo-Wealth.” In *Contemporary Issues in Macroeconomics: Lessons from The Crisis and Beyond*, edited by Joseph E. Stiglitz and Martin Guzman. IEA Conference Volume no. 155-II. Basingstoke, UK: Palgrave Macmillan.
- Hanson, Robin. 2016. *The Age of Em*. Oxford: Oxford University Press.
- Harari, Yuval N. 2017. *Homo Deus: A Brief History of Tomorrow*. New York: Harper.
- Hicks, John. 1932. *The Theory of Wages*. London: Macmillan.
- Kennedy, Charles. 1964. “Induced Bias in Innovation and the Theory of Distribution.” *Economic Journal* LXXIV:541–47.
- Keynes, John Maynard. 1932. “Economic Possibilities for our Grandchildren.” In *Essays in Persuasion*, 358–73. San Diego, CA: Harcourt Brace.
- Korinek, Anton. 2016. “Currency Wars or Efficient Spillovers? A General Theory of International Policy Cooperation.” NBER Working Paper no. 23004, Cambridge, MA.
- . 2017. “The Rise of Artificially Intelligent Agents.” Working paper, Johns Hopkins University.
- Korinek, Anton, and Ding Xuan Ng. 2018. “Digitization and the Macroeconomics of Superstars.” Working paper, Johns Hopkins University and University of Virginia.
- Korinek, Anton, and Joseph E. Stiglitz. 2017. “Worker-Replacing Technological Progress.” NBER Working Paper no. 24174, Cambridge, MA.
- Kurzweil, Ray. 2005. *The Singularity Is Near: When Humans Transcend Biology*. New York: Viking.
- Lipsey, Richard, and Kelvin Lancaster. 1956. “The General Theory of Second Best.” *Review of Economic Studies* 24 (1): 11–32.
- Malthus, Thomas Robert. 1798. *An Essay on the Principle of Population*. Project Gutenberg.
- Manning, Alan. 2011. “Imperfect Competition in Labour Markets.” In *Handbook of Labor Economics*, vol. 4, edited by O. Ashenfelter and D. Card. North-Holland: Amsterdam.
- Meade, James E. 1955. *Trade and Welfare*. Oxford: Oxford University Press.
- Moore, Gordon E. 1965. “Cramming More Components onto Integrated Circuits.” *Electronics* 38 (8): 114:ff.
- Mortensen, Dale T., and Christopher A. Pissarides. 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *Review of Economic Studies* 61 (3): 397–415.
- . 1998. “Technological Progress, Job Creation, and Job Destruction.” *Review of Economic Dynamics* 1 (4): 733–53.
- Newbery, David, and Joseph E. Stiglitz. 1984. “Pareto Inferior Trade.” *Review of Economic Studies* 51 (1): 1–12.
- Omohundro, Stephen M. 2008. “The Basic AI drives.” In *Artificial General Intelligence 2008: Proceedings of the First AGI Conference*, edited by Pei Wang, Ben Goertzel, and Stan Franklin, 483–92. Amsterdam: IOS.
- Piketty, Thomas, Emmanuel Saez, and Stefanie Stantcheva. 2014. “Optimal Taxation of Top Labor Incomes: A Tale of Three Elasticities.” *American Economic Journal: Economic Policy* 6 (1): 230–71.
- Restrepo, Pascual. 2015. “Skill Mismatch and Structural Unemployment.” Working paper, Massachusetts Institute of Technology.
- Samuelson, Paul. 1965. “A Theory of Induced Innovations along Kennedy-Weisacker Lines.” *Review of Economics and Statistics* XLVII:444–64.

- Schmitt, John. 2013. *Why Does the Minimum Wage Have No Discernible Effect on Employment?* Washington, DC: Center for Economic and Policy Research.
- Shapiro, Carl, and Joseph E. Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review* 74 (3): 433–44.
- Simsek, Alp. 2013. "Speculation and Risk Sharing with New Financial Assets." *Quarterly Journal of Economics* 128 (3): 1365–96.
- Solow, Robert. 1987. "We'd Better Watch Out." *New York Times Book Review*, July 12, 36.
- Stiglitz, Joseph E. 1969. "Distribution of Income and Wealth among Individuals." *Econometrica* 37 (3): 382–97.
- . 1976. "The Efficiency Wage Hypothesis, Surplus Labour and the Distribution of Income in LDCs." *Oxford Economic Papers* 28:185–207.
- . 1987a. "On the Microeconomics of Technical Progress." In *Technology Generation in Latin American Manufacturing Industries*, edited by Jorge M. Katz, 56–77. New York: St. Martin's Press.
- . 1987b. "Technological Change, Sunk Costs, and Competition." *Brookings Papers on Economic Activity* 3:883–947.
- . 2006. "Samuelson and the Factor Bias of Technological Change." *Samuelsonian Economics and the Twenty-First Century*, edited by M. Szenberg et al., 235–51. New York: Oxford University Press.
- . 2008. "The Economic Foundations of Intellectual Property." *Duke Law Journal* 57 (6): 1693–724.
- . 2014a. "Intellectual Property Rights, the Pool of Knowledge, and Innovation." NBER Working Paper no. 20014, Cambridge, MA.
- . 2014b. "Unemployment and Innovation." NBER Working Paper no. 20670, Cambridge, MA.
- . 2014c. "Tapping the Brakes: Are Less Active Markets Safer and Better for the Economy?" Presentation at the Federal Reserve Bank of Atlanta 2014 Financial Markets Conference.
- . 2017. "Pareto Efficient Taxation and Expenditures: Pre- and Redistribution." NBER Working Paper no. 23892, Cambridge, MA.
- Stiglitz, Joseph E., and Bruce Greenwald 2015. *Creating a Learning Society: A New Approach to Growth, Development, and Social Progress*. New York: Columbia University Press.
- Stiglitz, Joseph E., with Nell Abernathy, Adam Hersh, Susan Holmberg, and Mike Konczal. 2015. *Rewriting the Rules of the American Economy*, A Roosevelt Institute Book. New York: W. W. Norton.
- Stiglitz, Joseph E., and Andrew Weiss. 1981. "Credit Rationing in Markets with Imperfect Information." *American Economic Review* 71 (3): 393–410.
- Tegmark, Max. 2017. *Life 3.0: Being Human in the Age of Artificial Intelligence*. New York: Knopf.
- Turing, Alan M. 1950. "Computing Machinery and Intelligence." *Mind* 59 (236): 433–60.
- United Nations Department of Economic and Social Affairs (UN). 2015. "United Nations World Population Prospects: 2015 Revision." https://esa.un.org/unpd/wpp/publications/files/key_findings_wpp_2015.pdf.
- Vinge, Vernor. 1993. "The Coming Technological Singularity: How to Survive in the Post-Human Era." In *Proc. Vision 21: Interdisciplinary Science and Engineering in the Era of Cyberspace*, 11–22. NASA: Lewis Research Center.
- von Weizacker, C. C. 1966. "Tentative Notes on a Two-Sector Model with Induced Technical Progress." *Review of Economic Studies* 33:245–51.
- Williams, Heidi. 2010. "Intellectual Property Rights and Innovation: Evidence from the Human Genome." NBER Working Paper no. 16213, Cambridge, MA.

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
İsçan, 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