

TOWARDS AN AI ECONOMY THAT WORKS FOR ALL

**A Report of the
Keystone Research Center Future of Work Project
sponsored by the Heinz Endowments**



**Harrisburg, PA
February 2019**

About the Heinz Endowments

The Heinz Endowments seeks to help the Pittsburgh region thrive as a whole and just community, and through that work to model solutions to major national and global challenges. It is based in Pittsburgh, Pennsylvania, where it concentrates on advancing a sustainable future for its community and our planet, successful learning outcomes for young people and their families, and a culture of engaged creativity for all citizens.

About the Keystone Research Center

The Keystone Research Center (KRC) was founded in 1996 to broaden public discussion on strategies to achieve a more prosperous and equitable Pennsylvania economy. Since its creation, KRC has become a leading source of independent analysis of Pennsylvania's economy and public policy. KRC is located at 412 North Third Street, Harrisburg, Pennsylvania 17101-1346. Most of KRC's original research is available on the KRC website at www.keystoneresearch.org. KRC welcomes questions or other inquiries about its work at 717-255-7181.

Authors

Stephen Herzenberg

Executive Director, Keystone Research Center

John Alic

Consultant to the Keystone Research Center

Advisory Panel

Ruth Milkman, *Chair*

Professor of Sociology, The City University of New York

Diane Bailey

Associate Professor
University of Texas at Austin

Grant Oliphant

President
Heinz Endowment

Peter Cappelli

Professor of Management
University of Pennsylvania

Adam Paulisick

Chief Product Officer
MAYA a BCG Company

Mary Gray

Senior Researcher
Microsoft Research

Michael Piore

Professor Emeritus
Massachusetts Institute of Technology

Gordon Hanson

Professor of International Economic Relations
UC San Diego

Ashley Putnam

Director
Economic Growth & Mobility Project
Federal Reserve Bank of Philadelphia

Thea Lee

President
Economic Policy Institute

Elisabeth Reynolds

Executive Director
MIT Industrial Performance Center

Victor McCrary

Vice President for Research & Graduate Programs
The University of the District of Columbia
Member, National Science Board

Leslie Richards

Secretary of Transportation
PENNDOT

Derecka Mehrens

Executive Director
Working Partnerships USA

Vien Truong

Chief Executive Officer
Green For All

Cassandra Ogren

Deputy Director of Research
International Brotherhood of Teamsters

About the Authors

Stephen Herzenberg holds a PhD in economics from MIT and since 1995 has been Executive Director of the Keystone Research Center (KRC), an economic research and policy think tank. Before KRC, Steve worked at the US Congress's Office of Technology Assessment (OTA) and the US Department of Labor (USDOL). At USDOL, Steve served as assistant to the chief negotiator of the labor side agreement to the North American Free Trade Agreement. Steve has written extensively on the challenges associated with economic restructuring, including the transition from a manufacturing-dominated U.S. economy with limited imports to a global, postindustrial, service-dominated economy; and, more recently, on the potential impact of robots and artificial intelligence. Steve's KRC writings are on line at www.keystoneresearch.org. His coauthored writings with John Alic for national audiences include *New Rules for a New Economy: Employment and Opportunity in Postindustrial America* (Cornell/ILR Press, 1998); and *US-Mexico Trade: Pulling Together or Pulling Apart?* (OTA 1992).

John Alic is a long-time student of technological change and its ramifications. He worked at the Office of Technology Assessment over most of its existence, directing studies beginning with *US Industrial Competitiveness: A Comparison of Steel, Electronics, and Automobiles* (1981). Since OTA's closure, he has consulted for government and nonprofits and written on topics including military affairs ("Technology," *SAGE Encyclopedia of War: Social Science Perspectives*, 2017), energy technology and climate change ("Energy Innovation and Policy," *Oxford Research Encyclopedia of Climate Science*, 2017), and the labor market ("What We Mean When We Talk About Workforce Skills," *Issues in Science and Technology*, Spring 2018). Alic is also author or coauthor of three books, the most recent of these *Trillions for Military Technology: How the Pentagon Innovates and Why It Costs So Much* (Palgrave Macmillan, 2007). An engineering graduate of Cornell, Stanford, and the University of Maryland, he has also taught and conducted research over the years at several universities.

Acknowledgements

The authors thank The Heinz Endowments for financial support and for their deep interest and engagement with the project. We thank the members of the Advisory Panel for generously contributing their time and insights, including at an August 9, 2018 meeting in Pittsburgh. Steve Viscelli contributed to this report, as did Keystone Research Center staff: Stephanie Frank composed the document and created the figures; Anne DeCecco and Mark Price helped with logistics and data; Erica Freeman took on some of the copy-editing; and John Neurohr took charge of arrangements for the report's release, David Cooper and John Schmitt of the Economic Policy Institute provided data for several of the figures. Two members of the Advisory Panel, Diane Bailey and Ruth Milkman, in addition to their participation in the Pittsburgh meeting, provided especially helpful comments on subsequent drafts. We are also benefited from reviews and comments by Matt Barron, Fred Block, and John Schmitt. The final content of this report is of course the sole responsibility of the two authors.

Table of Contents

Executive Summary	1
I. Introduction	6
II. Technological Change and Employment: Are the Robots Coming?	13
III. Artificial Intelligence	28
IV. Prediction: Are the Robots Coming, Again?	42
The Perils of Prediction	43
AI in the Workplace	45
Prediction: AI as Technological Innovation	50
Implications	54
V. The Plight of the American Worker	57
What Explains the Plight of the American Worker?	61
VI. Policy	70
Promote Task-Relevant Skill and Worker Resilience	80
Rebalance the Scales Between Workers and Employers	85
Economics That Works for All	92

Figures

1. Wage Growth No Longer Tracks Productivity Growth	10
2. Those at the Very Top Have More than Doubled Their Share of Income Since the 1970s	11
3. Workers With More Years of Schooling Earn More	23
4. Wage Inequality: Hourly Pay Has Stagnated for Most US Workers	58
5. The Top 0.1% Have Gotten Wealthier at the Expense of Other Americans	58
6. A Smaller Share of American Men are Working	59
7. Private and Public Sector Union Density	64

Boxes

II-A	Skill-Biased Technological Change: Computers in the Workplace	24
II-B	Education: No Panacea	26
III-A	Hardware and Software for Machine Learning	30
III-B	Decision-Making and Intelligence	32
III-C	Expert Systems: Buoyant Hopes Deflated	34
III-D	Autonomous Vehicles: Sensors, Software, and Safety	38
IV-A	Moore's Law: Predictable Innovation	44
IV-B	The Internet: An Unpredicted Socio-Technical-Economic Explosion	51
IV-C	Technological Innovation and Social Change	52
IV-D	Autonomous Vehicles: Social and Economic Considerations	55
V-A	Institutions and Inequality	62
V-B	Legal Institutions and the American Worker	66
V-C	The Gig Economy: 21st Century Casual Work	69
VI-A	Collective Innovation: Fountainhead of Prosperity	76

Tables

1. Major Policy Categories for Achieving an AI Economy That Works for All	72
2. Policies to Promote Work-Relevant Skills and Worker Resilience	81
3. Policies to Rebalance the Scales Between Workers and Employers	86

Executive Summary

This is the first report of a Keystone Research Center project on the “Future of Work.” The aim is to identify public policies that could help ensure that the application and diffusion of artificial intelligence (AI) over the next several decades fosters an economy in which Americans generally thrive. The project is motivated, in part, by concern that the opposite could occur: that AI will exacerbate the already high levels of income and wealth inequality in the United States. Our most important conclusion is that AI need not make our inequalities more severe. Creative public policies could lead to an AI economy “that works for the many, not just the few.”

The study design has been informed by the two authors’ experience at the one-time Office of Technology Assessment (OTA) of the US Congress. To guide the undertaking and provide feedback on its products, we recruited an advisory panel of nationally recognized academics and representatives of think tanks and the corporate, labor, and non-profit sectors. The project methodology combines interviews with technology experts, policy analysis, synthesis of research literature and, still to come, sectoral studies.

This first report contains three main parts. (1) Following an introduction, Sections II-IV consider AI’s likely impacts through the lens of technology. Section II reviews past impacts of innovations including robotics and information technology on the economy and jobs. Section III looks at AI itself, how it does and does not go beyond earlier IT in substituting for human capacities and intelligence. Section IV explores the difficulties of predicting AI’s job impacts. (2) Section V, “The Plight of the American Worker,” examines the labor market context in which AI systems will spread and the roots of the economic inequality from which the nation suffers. (3) Section VI surveys policies to influence inequality and the distribution of the benefits of productivity growth as AI spreads.

(1) The technology sections of this report conclude the following.

- Fears of massive job displacement as a result of AI and continuing advances in robotics cannot be justified in the short to medium term. System-level innovations, those drawing on multiple technologies—which for AI means not just software but a great many hardware technologies (e.g., sensors for autonomous vehicles)—almost invariably perform poorly at first and evolve more slowly than anticipated.

- Digital disruption of the US labor market, on the other hand, has been underway since the 1960s, in factories with increasing levels of automation making products designed and indeed reconceived for lower labor content, and in offices everywhere because of computerization and IT. In this light, the maturation of AI represents less a disjuncture than the most recent episode in a long-running sequence that has already disrupted millions of lives and altered in fundamental ways many tens of millions of jobs.
- AI will penetrate more broadly than previous IT systems, in part because of machine learning, processes in which computers incrementally improve their performance in specific, narrow domains by absorbing massive amounts of data. (Computers that “learn” from viewing over 100,000 X-rays, for example, already outperform most radiologists.) Even so, AI-enabled automation will remain behind human capabilities for the foreseeable future in work requiring judgment, tacit skills, and common sense—tasks that infuse many parts of most jobs.
- While AI will ultimately have widespread impacts, predictive methods are inadequate for detailed forecasts. We can only be sure that AI will cause significant disruptions—and as a society we should prepare for that.

(2) These disruptions will occur in an economy with levels of inequality as great as at any time in our nation’s history. Some economists attribute today’s inequality to previous generations of technology that favored more highly educated over less educated workers—to “skill-biased technological change” (with education levels assumed, crudely, to measure skill). We conclude that a far stronger case can be made that public policy (deregulation, including of labor markets) and institutions (e.g., the decline of collective bargaining) explain more of the increase in inequality. There is good news in that finding: AI offers an opportunity to reimagine the nation’s approach and to tip the scales back toward the interests of the 99% rather than the 1%. This conclusion also means that, as desirable as more and better education may be, additional years of schooling will not, alone, suffice to narrow income inequality. Education should be viewed as good in itself, not as a cure for labor market ills.

(3) The last section of the report aims to spur richer debate about policies to ensure that AI generates broader benefits than previous episodes of major technological change. Such debates have so far been bracketed by the belief that we shouldn’t or couldn’t do anything (a view favored by proponents of deregulation and “free” markets); and by recent enthusiasm for

a guaranteed or universal basic income (UBI) premised on ideas that machines will take over so many tasks that Americans will need means to support themselves divorced from work. The table on the next page (a longer version can be found in Section VI) demonstrates that policymakers have a great many tools between doing nothing and implementing a guaranteed income.

The starting point (first row of the table) is recognition that the impacts of AI on our labor market and society are well within the reach of public policy. Government policies shape technological innovations and applications of all sorts; the policy questions concern in whose interests government intervenes, through action or inaction. AI offers a new opportunity to reclaim public policy for the many rather than the few, so that our nation's economy evolves along a path that benefits the common good.

Many measures in the table depart from post-World War II US policies—and even more from post-Reagan-administration policies concerning regulation, trade, technology, and the labor market. With respect to category 4, for example, the United States spends remarkably little on retraining dislocated workers and helping them find new jobs (less than a quarter of the average for 29 OECD countries). The message to workers: “You’re on your own.” The United States is just as much an outlier among wealthy nations when it comes to worker rights and enforcement of these rights. Except for the brief New Deal-era interregnum from the 1930s through World War II, US laws and their interpretations by the courts have favored employers, and since the 1970s more strongly almost year by year.

The first and overriding task in addressing the implications of AI is to find responses that go beyond the marginal remedies debated over recent decades. The underlying dynamics of inequality, the imbalance in power between employers and employees, must be reversed. Rather than default deference to business, deregulation, and “the market,” policy must begin from a new starting point: that the goal is to achieve an AI economy that works for all. Markets, after all, are social constructs, tailored in large part by interested parties the ability to get their way: economics has nothing to say about distribution of the benefits of technological innovation and productivity growth. These are questions of political choice, and policy choice.

Major Policy Categories for Achieving an AI Economy That Works for All	
<i>Category</i>	<i>Goals</i>
1. Advance the narrative that we can have an economy that works for all.	Create space for policy forceful enough to achieve equitable growth and elevate an economic policy framework that supports policymakers and society in the pursuit of equitable growth.
2. Promote trust-busting 2.0.	Reassert, beginning with more aggressive enforcement of existing antitrust laws, government's obligation to regulate business enterprises for the public good and to benefit workers and consumers.
3. Promote stakeholder over shareholder capitalism.	Give communities and workers a stronger voice in corporate decision-making as AI technologies reshape the economy and society.
4. Promote work-relevant skills and worker resilience for dislocated, young, and employed Americans.	Institutionalize with adequate funding systems that enable workers to learn, and relearn, work-relevant skills that lead to family sustaining jobs. Reduce fear of displacement and its economic and social costs; boost job satisfaction, business performance, and productivity growth.
5. Rebalance the scales between employees and employers.	Reverse the deeply rooted US elevation of employer interests over employees; make real the nearly extinguished rights of workers to act collectively to raise wages, reduce inequality, and restore mobility and indeed democracy itself.
6. Create enough paid, family-sustaining, jobs to go around.	Reduce work time and establish a right to work (in the original meaning of the phrase) backed by public job creation; avoid mass unemployment, maintain the dignity of work and the satisfaction that comes from meaningful work for a living wage.
7. Share AI productivity gains through tax policy and social policy (e.g., income guarantees).	Recognize that the collective sources of AI innovations—which don't just spring from the heads of a few entrepreneurs—justify wide distribution of the benefits; reduce after-tax inequality; avoid the societal stagnation that accompanies intergenerational social immobility.
8. Promote responsive democracy.	Reverse slide toward oligarchy; build grassroots political power to implement policies in other categories.

The second part of this project will examine AI applications and impacts in specific sectors—possibly including transportation, retail and distribution, health care, manufacturing, and construction, among others. The analysis will explore alternative scenarios, focusing on policies and institutions that will benefit the great mass of ordinary Americans who work for a living, and their families. By being concrete about pathways to shared prosperity, in its second stage our project will seek to increase the chances that the United States will find and follow those pathways.

I. Introduction

That an enormous Proportion of Property vested in a few Individuals is dangerous to the Rights, and destructive of the Common Happiness, of Mankind

~ Proposed amendment to Pennsylvania State Constitution, attributed to Benjamin Franklin, 1776 ¹

This is a first report with preliminary analysis and findings from a project that will offer a menu of policy responses to artificial intelligence (AI). In the words of Grant Oliphant, president of The Heinz Endowments, the project’s sponsor, the goal should be to ensure that AI will “work for the many, not just the few.”² As the Executive Summary indicates, our primary concern has been the prospective impacts of AI—more broadly, information technology (IT)—on labor demand and inequality.

AI is not one thing, but many. Combined with innovations in hardware including robotic manipulators and sensors, and new forms of work organization, IT has contributed to job displacement, wage stagnation, and inequality for decades without generating anything beyond piecemeal policy responses. Absent more effective policies, AI and AI systems—AIs for short—seem poised to deepen, broaden, and further accelerate these impacts, trapping, in combination with economic deregulation (e.g., of transportation and telecommunications), and globalization (especially in manufacturing) even more of the US workforce in low-quality jobs, poorly paid and lacking in opportunities for advancement.³ Downward mobility for many has been coupled with rapid increases in income and wealth for those at the top. These trends are not a

¹ Franklin’s text, intended as Article 16 of the Pennsylvania Declaration of Rights for the constitution then being drafted, did not survive into the final version. See “Revisions of the Pennsylvania Declaration of Rights, [between 29 July 1776 and 15 August 1776],” Founders Online, National Archives, <http://founders.archives.gov/documents/Franklin/01-22-02-0314>. Original source: *The Papers of Benjamin Franklin*, Vol. 22, March 23, 1775, through October 27, 1776, William B. Willcox, ed. (New Haven: Yale University Press, 1982), pp. 529-533.

² The quotation is from p. 4 of “Remarks by Grant Oliphant,” President, The Heinz Endowments, at Metro 21: Smart Cities Institute Launch, Pittsburgh, PA, March 2, 2018; www.heinz.org/UserFiles/File/GO%20speech%20Metro21%20Smart%20Cities%203_2_18.pdf.

³ For a brief introduction to the sorts of concerns widely expressed over the last few years, see Erik Brynjolfsson and Tom Mitchell, “What Can Machine Learning Do? Workforce Implications,” *Science*, Vol. 358, No. 6370 (2017), pp. 1530-1534.

consequence of inexorable economic or technological forces. They reflect political dynamics and policy choices. Some of these go back many decades and others have emerged more recently, during and after the presidency of Ronald Reagan. Most of all, perhaps, they reflect the hostility of American employers to labor, omnipresent from the beginnings of big business well over a century ago.

In theory, a new wave of labor-displacing technological change could be cause for celebration. “Productivity growth,” after all, is another name that economists use for software and/or hardware that enable fewer work hours to produce goods or services of the same or superior quality—or indeed products with new and unprecedented attributes such as smartphones. And rapid productivity growth can undergird overall improvements in living standards. The catch—the reason having machines do more of our work might not be reason for celebration—is distribution. If income continues to be distributed in highly unequal fashion, and we put in place no policies to create paid work for all those who want it, or an adequate living for those without enough paid work, a boost in productivity growth could simply make our current inequalities more savage. Behind hand-wringing that “the robots are actually coming this time,” and the surprising (in the US context anyway) attention paid to the idea of a “universal basic income” (UBI)—government-provided income that would not be conditioned on whether people work—lie an implicit expectation that America will not get right the distribution of the benefits from AI. The fear, to illustrate the point with a concrete example, is that American long-haul truck drivers, whose hourly compensation (counting the hours they spend waiting for the next load) has declined by more than half since the industry deregulation in the late 1970s, may now have their jobs eliminated or, as our colleague Steve Viscelli anticipates, find themselves driving for even lower pay as part of last mile delivery.⁴ AIs, many observers feel, will be used to kick the American worker when she’s down. The only way to head off such outcomes is through policies aimed at more equitable outcomes—starting with rebalancing the scales between

⁴ Viscelli concludes that in road transport, AI-enabled semiautonomous and fully autonomous trucks will cut into relatively good driving jobs more than bad jobs, resulting in lower average wages for truck drivers and more onerous working conditions. See Steve Viscelli, *Driverless? Autonomous Trucks and the Future of the American Trucker* (Berkeley and San Jose: University of California, Center for Labor Research and Education and Working Partnerships USA, September 2018); online. On the compensation fall of long-haul truck drivers since deregulation, see Michael H. Belzer and Stanley A. Sedo, “Why Do Long Distance Truck Drivers Work Extremely Long Hours?” *Economic and Labour Relations Review*, Vol. 29, No. 1, 2018, pp. 59-79.

employers and employees. We outline such policies in the last section of this report and will expand upon them in the second stage of the project.

Fear of technological unemployment is not new. It has been a persistent concern since the Great Depression in the 1930s (Section II). Since the 1950s, information technology—beginning with hulking mainframe computers to automate paperwork such as financial record-keeping and preparation of payrolls, reducing clerical employment, and more recently digital ubiquity in forms ranging from internet commerce to smartphone apps—have upset industries and transformed jobs and work throughout the economy. In this light, AI is simply the next stage in a long-running process. Already, AI systems handle some tasks and jobs quite well. Others, especially those that call for the sort of common sense that humans bring to everyday life as well as their workplaces, remain at least for now beyond even the most advanced AI (Section III).

For such reasons, and because of commonalities between AIs and other complex and systemic technological innovations (Section IV), we believe the short to medium term implications of AI have sometimes been exaggerated. Still, in the longer term AI seems bound to cause major disruption in the labor market. After all, this is what happened with past generations of IT, office automation and computerized factory production, which have displaced millions of workers since the 1950s and 1960s—and which did so without anything so powerful as emerging approaches like machine learning promise.

While we can expect major impacts, projections concerning new and emerging technologies must be taken as no more than suggestive—little more than gazing into foggy crystal balls. Predictions raise awareness and may help focus policy attention but they invariably go astray. In any event, whatever AI brings next will not arrive overnight. The United States can and should use that time to prepare. As yet, there are few signs that it will. While governments everywhere struggle to grapple with the implications of technical change—with the positives and negatives of Schumpeterian creative destruction—the United States has rarely addressed labor market disruptions with truly serious intent.

Compared with other high-income nations the United States stands out as exceptional with respect to what can only be termed a gross imbalance in policies affecting the labor market. Federal government policies cater unabashedly to business interests and the well-off

while offering little beyond modest levels of unemployment benefits and mostly unrealizable promises of upward mobility to those who experience the destructive aspects of innovation and productivity growth. The chief exceptions, long past, came as part of the response to the Great Depression. The measures devised at that time put many Americans back to work, and World War II completed the recovery. The policy thrust behind those measures, increasing consumer demand to drive down unemployment and ensuring fair wages for workers, whether earning the minimum wage or in a unionized factory, delivered broadly shared prosperity until the 1970s (Section V). The basic New Deal policy approach then withered away as a result of renewed employer assaults on workers, supported by political conservatives in Washington and elsewhere—a counterattack that began even before World War II ended.

The plight of America’s working families since the 1970s stems in considerable part from free market ideology that licenses unequal sharing of the rewards of productivity growth and innovation. Although government supports much of the technology base from which innovation springs (including that for AI), private firms conceive and develop the new products and processes that result. Advocates of “free” enterprise therefore have a credible argument when they seek minimal restrictions by government. Less credible is their claim that a free hand for technological advance and business innovation will yield prosperity that redounds to the benefit of all. Also not credible is the corollary argument that innovation can be attributed to a visionary few—entrepreneurs, financiers, bold captains of business and industry, the occasional Edisonian inventor (Section VI).

Technological advance and wealth creation through productivity growth are in fact collective activities. They emerge from deep within society and the economy, with contributions, if unrecognized, from all levels of the workforce. Government, moreover, has always given plenty of help to commerce and industry, with new boosts after World War II through generous funding for technology and science.⁵ The processes are fundamentally incremental and largely although not entirely anonymous. All this was more or less obvious at a time when visitors from abroad admired Yankee ingenuity as they now remark on Silicon Valley entrepreneurship, the era


⁵ Carter Goodrich, *Government Promotion of American Canals and Railroads, 1800-1890* (New York: Columbia University Press, 1960); Vernon W. Ruttan, *Is War Necessary for Economic Growth? Military Procurement and Technology Development* (New York: Oxford University Press, 2006).

of the machine shops rather than microchips. Now they are far from self-evident: no one, after all, can “see” what goes on within a transistor gate much less a machine learning program.

Our argument is that the United States can have both rapid, wealth-creating productivity-enhancing innovation and a more equal society—these are not in tension, notwithstanding claims to the contrary by free-enterprise myth-makers.

Figure 1 shows that compensation for typical workers tracked productivity growth from the late 1940s to 1973—then quite suddenly stopped doing so. Over the years 1947-1973, real incomes roughly doubled for families at all income levels, low, middle, and high.⁶ Wages for most workers flat-lined in the 1970s. At about the same time, pay packages for top executives, which had largely tracked those for the rest of the workforce since before World War II, began

Figure 1
Wage Growth No Longer Tracks Productivity Growth




Note: Compensation (wages and benefits) for private-sector nonsupervisory workers. Productivity growth calculated as increase in output of goods and services minus depreciation.

Source: Economic Policy Institute, Washington, DC, based on data from the Department of Labor, Bureau of Labor Statistics and Department of Commerce, Bureau of Economic Analysis. <https://www.epi.org/productivity-pay-gap/>

⁶ See Current Population Survey, Bureau of the Census, “Table F-1. Income Limits for Each Fifth and Top 5 Percent of Families (All Races): 1947 to 2017,” online, which shows remarkably similar increases for all groups until 1973.

Figure 2

Those at the Very Top Have More than Doubled Their Share of Income Since the 1970s



Note: Income received by individuals or married couples filing jointly, before taxes and transfers.

Source: Based on data from Estelle Sommeiller and Mark Price. *The New Gilded Age: Income Inequality in the U.S. by State, Metropolitan Area, and County* (Washington, DC: Economic Policy Institute, July 19, 2018).

to swell.⁷ Income inequality rose, with large gains for the top 0.1% (Figure 2). Given these dynamics, if AI leads to large-scale disruptions of the US job market, we will be unable to manage them without confronting the fundamental imbalances in our economy, public policy, and politics. Put differently, policymakers cannot significantly influence the outcomes by playing only at the margins.

To encourage policymakers to be bold, Section VI outlines an ambitious array of policy options—including rebalancing the employer-employee relationship; creating enough jobs to go around through shorter work time (i.e., spreading around the tasks that AI and robots don’t do) and public employment creation; funding a universal system of income and reemployment supports for dislocated workers; and investing in learning models (such as apprenticeship) to develop and enhance occupational skills and competencies, fostering worker resilience. Our policy recommendations also highlight the need to evaluate policy not on the basis of the “free-

⁷ Carola Frydman and Raven E. Saks, “Executive Compensation: A New View from a Long-Term Perspective, 1936–2005,” *Review of Financial Studies*, Vol. 23, No. 5 (2010), pp. 2099–2138.

market” orientation of the past four decades but with the goal of a better life for all Americans. “Markets” and “deregulation” have their place but are not ends in themselves. We need a policy perspective focused on the right objective—“economics that works for all.”

The national debate about the future of work provides an opportunity to reexamine where our country is headed. Silicon Valley’s UBI talk has legitimized the search for ways to create a decent society in a potentially brave new world of AI. This project aims to take up that invitation, and to begin the grounded analysis necessary to imagine, envision, and then build a future with more, not less, human dignity.

II. Technological Change and Employment: Are the Robots Coming?

[T]hinking ‘robots,’ capable not only of taking over the traditional manual efforts on the part of human beings, but capable of injecting some substantial measures of memory and of even ‘thinking’ processes ... have naturally captured the public imagination.

One of these machines, called Erma—which might well sound like a hurricane to those affected—has recently been getting a lot of publicity because it can reputedly do the chores of 50 bank clerks.

*~ Wright Patman, Chairman, Subcommittee on Economic Stabilization,
US Congress, 1955¹*

Patman delivered the above remarks in 1955, a half-dozen years before General Motors installed its first robot. At the end of that decade, in 1969, GM opened a heavily automated assembly plant in Lordstown, Ohio, in the news again recently because of its announced closure. The Lordstown assembly line, equipped with spot-welding robots to build the new Chevrolet Vega, moved at unprecedented speed. Walkouts by members of the United Automobile Workers (UAW) that followed are probably more widely recalled than the Vega itself. Robots, even “thinking” robots—provided we use the term as loosely as Congressman Patman did, simply to refer to digital computers—have been around for a long time. Over the decades, they have altered jobs and tasks throughout the economy. Concerned statements by public officials such as Patman notwithstanding, the US government has never confronted the ongoing dynamics with much seriousness.

Well before the severity of the 1930s depression became fully apparent, British economist J.M. Keynes wrote, “We are being afflicted with a new disease ... technological unemployment. This means unemployment due to our discovery of means of economising the use of labour outrunning the pace at which we can find new uses for labour.”² As Keynes well knew, technological unemployment was being not so much discovered as rediscovered, having

¹ *Automation and Technological Change*, Hearings before the Subcommittee on Economic Stabilization of the Joint Committee on the Economic Report, Congress of the United States, Eighty-Fourth Congress, October 14, 15, 17, 18, 24, 25, 26, 27, and 28, 1955 (Washington, DC: Government Printing Office, 1955), p. 6.

² John Maynard Keynes, “Economic Possibilities for Our Grandchildren (1930),” in *Essays in Persuasion* (London: Macmillan, 1931), ebook pp. 290-304; quotation from p. 296. For a broad review from the perspective of economics highlighting indicators and analytical questions related to robotics and AI, see Jason Furman and Robert Seamans, “AI and the Economy,” Working Paper 24689, National Bureau of Economic Research, Cambridge, MA, June 2018.

been around at least since the beginnings of the factory system in England at the time of the first Industrial Revolution. What was new was that the effects were no longer concentrated in a handful of industries but seemed universal. In the United States, fears of technological unemployment had been widespread in the 1930s but receded as World War II put millions to work producing matériel for the Allied military and millions of others into uniform.³ Even so, by the time of the Lordstown strike, Americans everywhere could see that new technologies were both altering and destroying jobs.

Less visibly, technological change was also spawning new jobs, new occupations, and new industries too, with reductions in costs for many products old and new consequent benefits for consumers regardless of income level. Prices came down with costs as productivity increased. Wages rose and living standards too. When the auto industry was young, Henry Ford slashed prices through mechanization and moving assembly lines. Sales boomed, Ford paid his workers more in part so more of them would buy the cars they themselves were making. Over much the same period productivity soared in agriculture. Mechanization, hybrid seeds, fast-growing livestock breeds, and many other innovations—widely diffused through government extension programs, state and federal, that helped farmers learn what worked in their local conditions—led to greatly reduced labor intensity, wider selections for consumers in the stores, and lower prices for many food products.⁴

Over recent decades, most of the expansion in industry and employment has come in the service-producing sector of the economy. Here too, costs and prices have come down, as they have for air travel and telecommunications. At the same time, computerized decision-making and decision support systems encroached on low-paying jobs, such as taking hotel and airline reservations, and began to threaten better paid occupations such as law, in which “serious consideration to the practical application of the computer [began] in the mid 1950s.”⁵ Well before the Lordstown strike, UAW president Walter Reuther worried about the prospect of mass layoffs—not of factory workers, but of white-collar clerical employees displaced by document

³ Amy Sue Bix, *Inventing Ourselves Out of Jobs: America's Debate Over Technological Unemployment, 1929-1981* (Baltimore, MD: Johns Hopkins University Press, 2000).

⁴ Wallace E. Huffman and Robert E. Evenson, *Science for Agriculture: A Long-Term Perspective* (Ames: Iowa State University Press, 1993).

⁵ Robert P. Bigelow, “The Use of Computers in the Law,” *Hastings Law Journal*, Vol. 24 (1973), pp. 707-731; quotation from p. 707.

scanners, typing machines, and equipment like ERMA.⁶ The acronym stood for Electronic Recording Machine, Accounting. ERMA automatically read bank checks imprinted with magnetic characters and recorded them on magnetic tape for further processing.⁷

Integrated circuits, the building blocks of today's digital systems, were introduced around 1960, and another decade passed before prices dropped to levels affordable for anything but military and space systems. Microprocessors did not arrive until the 1970s. Sensors and electromechanical actuators were expensive too, with limited capabilities. For such reasons, early industrial robots were simple, even crude by today's standards, dumb and slow. They were fixed in place, their motions few, cycle times short, and assembly lines moved at constant speed. Even so, robots could take over repetitive tasks. Many of these, such as welding, were dirty and dangerous. Just as steam shovels and bulldozers had earlier displaced arduous manual labor in digging ditches, sewers, and canals, jobs were lost, but to no great regret on the part of many except those directly affected, as at Lordstown.

Even in the 1950s, as Congressman Patman's remarks indicate, the meanings attached to robotics and automation extended to offices and paperwork. Punched card equipment had for decades been used for keeping business records and for tabulating information such as census returns. The first known application of electronic digital computing in private industry began around 1950 when a small British firm, J. Lyons, designed and built its own hardware and software for inventory management.⁸ Lyons supplied some 200 tea shops in and around London. Each day the company took orders for next-day deliveries of baked goods, blended teas (available in hundreds of varieties), and other products. The Lyons computer kept track of orders and stock on hand, much like the IT systems employed later by Walmart and other wholesale and retail firms. All such systems, in effect, seek to automate information processing earlier the province of clerical personnel working with paper and pen, telephone and telegraph, adding machines, rotary calculators, and punched cards. Procedures then and now follow methods—rooted in accounting, industrial engineering, and what today is called management science—for

⁶ Bix, *Inventing Ourselves Out of Jobs*, p. 242.

⁷ Amy Weaver Fisher and James L. McKenney, "The Development of the ERMA Banking System: Lessons from History," *IEEE Annals of the History of Computing*, Vol. 15, No. 1 (1993), pp. 44-57. The acronym is usually capitalized.

⁸ John Aris, "Inventing Systems Engineering," *IEEE Annals of the History of Computing*, Vol. 22, No. 3 (2000), pp. 4-15.

getting the right inputs to ensure the right products are delivered to the right customers at the right time, while seeking to minimize costs and maximize revenues. A century ago, department stores such as Macy's and catalog retailers such as Sears Roebuck employed veritable armies of clerks. From the 1950s onward, more and more of this work was automated.

In the early years, this was mostly a matter of big corporations putting paperwork on big computers; smaller firms that could not afford mainframes gave their business to providers that specialized in standardized business applications such as payroll services.⁹ Banks and insurance companies bought or leased electronic machinery for bookkeeping, preparing monthly statements, soon including credit card statements, and calculating interest payments and billings. Computers did the arithmetic, crediting one account and debiting another with mistake-free accuracy, calculating totals and preparing statements. IT eliminated most sources of error other than incorrect inputs of hand-entered data. These errors remained possible, however, since as late as the 1990s optical character recognition systems could read no more than around 40% of handwritten dollar amounts on bank checks, leaving the rest to be handled by human operators (who had to check the machine-read figures too).¹⁰ Unlike the ATMs (automated teller machines) that began to appear outside banks in 1969, this first generation of automation was invisible to ordinary customers, who still received paper in the mail. Even so, bank customers benefited as bookkeeping errors that could be frustrating to find and correct diminished. Consumers also benefited from the convenience of ATMs. As with back-office check processing, the original intent of the ATM was to reduce labor costs.¹¹ That changed once banks realized that the more ATMs they scattered through local markets, the more customers they could expect to attract. Recognizing ATMs as a source of competitive advantage, they opened more branches and hired more tellers to staff them; employment in commercial banking drifted downward for a while, rose quite steeply in the late 1980s, then began to decline again as banks renewed their efforts to

⁹ James W. Cortada, "Commercial Applications of the Digital Computer in American Corporations, 1945-1995," *IEEE Annals of the History of Computing*, Vol. 18, No. 2 (1996), pp. 18-29. As late as 1965, US firms bought only about 5600 computers, nearly all of them mainframes (and many of these for technical calculations in engineering and science). Marcus E. Einstein and James C. Franklin, "Computer Manufacturing Enters a New Era of Growth," *Monthly Labor Review*, September 1986, pp. 9-16.

¹⁰ "Technology and Industrial Performance in the Service Sector: Field Research Report." Keystone Research Center, Harrisburg, PA, December 1997, p. 69.

¹¹ Ray M. Haynes, "The ATM at Age Twenty: A Productivity Paradox," *National Productivity Review*, Vol. 9 (1990), pp. 273-280.

curtail labor costs.¹²

Over the following decades, rather than robots replacing assembly line workers one-for-one and machines such as ERMA replacing clerical employees 50 at a time, whole new work systems came into being and products themselves began to be redesigned for easier automation. At first, computer-based airline reservations and ticketing systems at travel agencies made agents' jobs easier—they could spend less time on the phone with airlines and more time with customers. Later the entire industry all but vanished as web-enabled end-customers became their own travel agents. Fixed prices, meanwhile, gave way to algorithms (computational procedures) designed to continuously adjust prices as airlines sought to fill all seats on all flights to maximize revenues and boost profitability (empty seats cost as much as full ones but bring in nothing).

In manufacturing industries, such as food production and chemicals, electronic sensors and controls hooked to minicomputers—smaller and far less costly than mainframes, hence suited to dedicated applications—began to take the place of gauges, meters, switches, and valves earlier watched over by skilled workers. Still less expensive microprocessor-based control units followed, opening even more applications. Programming of numerically-controlled machine tools grew easier as did modifying control routines that told industrial robots what to do.

Until the early 1980s, makers of desktop computers tried to sell them as “personal computers” to hobbyists and for the home. They did not target or even recognize the office market until businesses began buying PCs to replace electric typewriters (some of which by then came with memories and some editing capability), dedicated word processors (single-purpose machines preloaded with proprietary text editing software), and electrical and electronic calculators. Vast new markets then opened, in large part because of complementary products from entrepreneurial software firms that developed cheap, easy-to-use packages for word processing (WordStar, WordPerfect) and spreadsheets (Visicalc, Lotus 1-2-3). A PC and the necessary software cost less than a dedicated word processor and the PC could do more. Neither PC manufacturers nor software firms saw this coming.¹³ Then in the 1990s, companies began

¹² Teresa L. Morisi, “Commercial Banking Transformed by Computer Technology,” *Monthly Labor Review*, August 1996, pp. 30-36.

¹³ Martin Campbell-Kelly, “Not Only Microsoft: The Maturing of the Personal Computer Software Industry, 1982-1995,” *Business History Review*, Vol. 75, No. 1 (2001), pp. 103-145; Thomas Haigh, “Remembering the Office of the Future: The Origins of Word Processing and Office Automation,” *IEEE Annals of the History of Computing*, Vol. 28, No. 4 (2006), pp. 6-31.

to connect what had been stand-alone PCs, workstations, and terminals through networks, many of them local and hard-wired. Once dial-up connections for wide-area networks gave way to broadband links, speed and capacity increased, costs plummeted, and decentralized computing spread rapidly.

Decentralized IT altered the task content of many jobs. Others simply melted away. These patterns, heralded as an emerging information economy, brought consequences for work and workers that from the beginning have been difficult to weigh in terms of opportunities, occupations, skills, and tasks. Over the years 2000 to 2017, for example, employment of telemarketers fell by about 270,000—jobs lost to web-based advertising and order-taking—while employers added some 860,000 new jobs for customer service representatives, many hired to handle telephone and online queries.¹⁴ New IT-based work systems evolved unpredictably, with much trial and error. As a share of total employment, clerical and administrative support occupations peaked in 1980 and began to drift downward.¹⁵

The consequences can be more easily seen in particular industries. Self-service spread through gasoline stations in the 1960s and scannable bar codes appeared in food stores over the following decade:

By “reading” the food industry’s new University Product Code (UPC) symbol [the computer] system makes checkout faster and easier for grocery checkers.... For added customer convenience, as a product is scanned the system’s point of sale terminal displays its name and price on a lighted panel, and prints a receipt. The receipt lists the name or a description of each item, its price, whether it is taxable or not, the total charge and any special price modifications such as the food stamp credits available, the time of the sale and the date.

One firm’s president went on to say

“We also expect sizable labor savings because a sale or other price change can be made by changing the price on a shelf label and in the computer rather than by remarking each item.”¹⁶

Hand counts of inventory on the shelves and in stockrooms gave way to scanning; accuracy rose, labor intensity fell, and retailers and wholesalers alike could keep supply and demand in better

¹⁴ Employment of telemarketers fell from 462,000 in 2000 to 190,000 in 2017; employment of customer service representatives rose from 1.91 million to 2.77 million. Figures from the Bureau of Labor Statistics (BLS) Occupational Employment surveys for 2000 and 2017, online.

¹⁵ Ian D. Wyatt and Daniel E. Hecker, “Occupational Changes During the 20th Century,” *Monthly Labor Review*, March 2006, pp. 35-57.

¹⁶ John M. Dunnagan, “Park ‘N Shop Going Electronic with New IBM Grocery Checkstand System,” *Computers and People*, January 1974, p. 42.

balance. Over the next decade, powerful retail chains like Walmart required suppliers to place bar codes on every product or package: no bar code, no shelf space.¹⁷ Supermarkets and big-box dry goods stores found they could keep many more items in stock without driving up their wage bills. Retailers started to push in-store costs off onto customers by installing terminals for self-scanning at checkout. Employment in retailing peaked around 2000 and began to fall, dropping sharply during the 2008 recession. Jobs then rebounded, rising to a little above the 2000 level (15.9 million compared to 15.3 million). This was an increase of only 3.8%, far less than the increase in US employment overall—a little over 13%. Labor productivity in retail trade had grown dramatically, by nearly 60%, a consequence of both structural change (big-box stores and no-box stores such as Amazon) and technological and organizational change (self-service, online sales).¹⁸ There is little reason to expect future job growth in retailing to keep pace with the rest of the economy; total employment in the industry could shrink—and not just in relative terms.

In Europe more than in the United States, jobless growth raised concern starting in the 1980s, with IT and telecommunications widely blamed when employment did not recover as expected following recessions.¹⁹ In the United States, especially, many employers bought into “business process reengineering,” aiming to shrink or replace entire organizational layers. By flattening their internal hierarchies, they expected to cut out many mid-level workers in administrative and managerial positions. As IT systems became more capable, they would take over tasks not just from front-line workers but better paid mid-level employees. Payrolls would be slashed, and many firms also sought strategic advantages as in the ATM and airline pricing examples. Reengineering failures were common, some of them abysmal and leading

¹⁷ Paul R. Messinger and Chakravarthi Narasimhan, “Has Power Shifted in the Grocery Channel?” *Marketing Science*, Vol. 14, No. 2 (1995), pp. 189-223; Emek Basker and Timothy Simcoe, “Upstream, Downstream: Diffusion and Impacts of the Universal Product Code,” Working Paper 24040, National Bureau of Economic Research, Cambridge, MA, November 2017.

¹⁸ Employment from BLS, Current Population Survey, as of July 2018. Labor productivity, 2000-2017, from BLS. Also see *Productivity and Costs by Industry: Wholesale Trade, Retail Trade, and Food Services and Drinking Places Industries – 2017*, USDL-18-1181 (Washington, DC: Bureau of Labor Statistics, July 19, 2018). For snapshots of employment change in retailing over the past two decades, see Michael H. Strople, “From Supermarkets to Supercenters: Employment Shifts to the One-Stop Shop,” *Monthly Labor Review*, February 2006, pp. 39-46; and Michael D. McCall, “Deep Drop in Retail Trade Employment During the 2007-09 Recession,” *Monthly Labor Review*, April 2011, pp. 45-48.

¹⁹ Chris Freeman and Luc Soete, *Work for All or Mass Unemployment? Computerised Technical Change into the 21st Century* (London: Pinter, 1994).

to abandonment of systems that had cost many millions of dollars to design and install.²⁰ Yet businesses learned from failures and successes as did the consultants who worked on many of these projects. Performance improved, and many mid-level workers who had once felt secure based on their tenure and experience lost their jobs. Other firms kept their old structures more or less intact, while firing employees en masse, bringing in younger people at lower wages or moving wholesale to parts of the country (or abroad) with lower prevailing wage levels and state governments unfriendly to labor. Many began to bring in temporary contract workers and part-timers without benefits, not even sick days.²¹ In one way or another, these examples illustrate how business objectives and business practices interact with technological capabilities to alter employment patterns, workplace tasks, and the skill requirements of jobs. The overall dynamics are not that different from those in the first Industrial Revolution when the driving force was steam power rather than telecommunications and digital electronics.

Health care, now accounting for about 18 percent of US GDP (measured by expenditures, in the absence of output indicators), has been the major exception to the trends sketched above. The industry includes many small providers—individual and group practices, specialized clinics, care centers, and hospices—along with hospitals of all sizes, from a dozen beds to a thousand and more.²² Mergers over the past two decades have left regions once served by many independently owned and operated hospitals and clinics with near monopolies. (The University of Pittsburgh Medical Center now runs some 36 hospitals and hundreds of clinics in the United States and several foreign countries.) With fewer independents and greater concentration, pressures for quality of care and control of costs—already weak since health care providers release little information on either prices or quality—have diminished still further: “research evidence shows that hospitals and doctors who face less competition charge higher prices to private payers, without accompanying gains in efficiency or quality.”²³ Meanwhile, public

²⁰ Paul A. Strassmann, *The Business Value of Computers* (New Canaan, CT: Information Economics Press, 1990).

²¹ *Contingent and Alternative Employment Arrangements—May 2017*, USDL-18-0942 (Washington, DC: Bureau of Labor Statistics, June 7, 2018).

²² Martin Gaynor, Kate Ho, and Robert J. Town, “The Industrial Organization of Health-Care Markets,” *Journal of Economic Literature*, Vol. 53, No. 2 (2015), pp. 235-284.

²³ Martin Gaynor, “Examining the Impact of Health Care Consolidation,” Statement before the Committee on Energy and Commerce Oversight and Investigations Subcommittee, US House of Representatives, February 14, 2018; quotation from p. 5. Gaynor, in his opening summary (p. 2), states “Prices are high and rising, there are incomprehensible and egregious pricing practices, quality is suboptimal, and the sector is sluggish and unresponsive,

policies have left many parts of the country with few insurance providers, and even those classed as not-for-profit entities have come to be managed much like their profit-seeking counterparts.²⁴

Health care lags well behind other industries in adoption of IT. Physicians enjoy far more power than most workers to fend off infringement on their preferred ways of practicing; although some do favor clinical protocols and AI support, others rail against “cookbook” care delivery and “medical Taylorism.”²⁵ In some contrast to advances in robot-assisted surgery, which have been remarkable, physicians’ offices, clinics, and hospitals are still full of paper files.²⁶ Decision support remains mostly a matter of checking for errors such as drug interactions. This is highly desirable, because such errors are common.²⁷ Still, much of health care informatics seems primitive compared to IT elsewhere in the economy, a continuing disappointment to enthusiasts who have been touting computer-assisted medicine for decades.²⁸ Unlike administrative records in other industries, patient files, accounts and billings, and insurance reimbursements include a great deal of specialized information, including handwritten physician notes and image scans (X-rays, MRIs). Neither clinical nor administrative information has been standardized, with exceptions as required for Medicare and Medicaid reimbursements, and legislated patient privacy protections have slowed adoption of electronic health records.²⁹ The consequences include high labor costs; gross inefficiencies, including missing or lost patient records, many of them with

in contrast to the innovation and dynamism which characterize much of the rest of our economy.” Also see Matt Schmitt, “Multimarket Contact in the Hospital Industry,” *American Economic Journal: Economic Policy*, Vol. 10, No. 3 (2018), pp. 361-387.

²⁴ “As Blue Cross plans evolved, they became more like other companies. Today, although many still are considered to be non-profit under state laws, most Blue Cross plans operate as though they were for-profit insurance companies, and they compete directly in the market for profits.” Max J. Skidmore, “Health Care in America, and Everywhere Else: A Review Essay,” *Poverty & Public Policy*, Vol. 2, No. 1 (2010), pp. 185-194; quotation from p. 188.

²⁵ Pamela Hartzband and Jerome Groopman, “Medical Taylorism,” *New England Journal of Medicine*, Vol. 374, No. 2 (2016), pp. 106-108.

²⁶ On robotic surgery, see, e.g., Nabil Simaan, Rashid M. Yasin, and Long Wang, “Medical Technologies and Challenges of Robot-Assisted Minimally Invasive Intervention and Diagnostics,” *Annual Review of Control, Robotics, and Autonomous Systems*, Vol. 11 (2018), pp. 15.1-15.26.

²⁷ David Blumenthal, Elizabeth Malphrus, and J. Michael McGinnis, eds, *Vital Signs: Core Metrics for Health and Health Care Progress* (Washington, DC: National Academies Press, 2015); Mary Dixon-Woods and Graham A. Martin, “Does Quality Improvement Improve Quality?” *Future Hospital Journal*, Vol. 3, No. 3 (2016), pp. 191-194.

²⁸ W.F. Bauer, D.L. Gerlough, and J.W. Granholm, “Advanced Computer Applications,” *Proceedings of the IRE*, Vol. 49, No. 1 (1961), pp. 296-304.

²⁹ *Realizing the Full Potential of Health Information Technology to Improve Healthcare for Americans: The Path Forward* (Washington, DC: Executive Office of the President, President’s Council of Advisers on Science and Technology, December 2010).

detailed information on previous illnesses and treatments; and stand-alone IT systems unable to share information with one another. For such reasons, clerical employment per unit of revenue in health care is reported to be “almost eight times more than other industries.”³⁰ Some 2.8 million people now work for health care providers in occupations classed as “office and administrative support.”³¹ A considerable fraction of these jobs could disappear over the next several decades as IT system integration spreads within health care, as it must: expenditures simply cannot continue to rise as they have over the past several decades. The alternative is further rationing of services. Other wealthy nations, after all, achieve comparable or superior health outcomes while spending on the order of half as much as the United States.³²

Two accompanying forces interacted with technological change to alter the US labor market: economic deregulation and globalization. Beginning in the Carter administration and accelerating under Ronald Reagan, deregulation of industries including transportation (air travel, railroads, trucking), telecommunications, and financial services, banking especially, ended the relatively relaxed competitive conditions earlier prevailing. Managements began trying to shave fractions of a cent from the costs of shipping freight, placing telephone calls, handling bank transactions. They laid off workers; clamped down on wage increases; redoubled resistance to labor unions; turned to part-time, temporary, and contract workers; and cut training expenditures.

In manufacturing, as globalization gathered speed symptoms of deindustrialization seemed everywhere by the late 1970s. Prefigured by earlier contraction in labor-intensive industries such as furniture, footwear, and apparel, other manufacturers and their suppliers shed large numbers of high-wage jobs, sometimes opening new plants in the South and West where prevailing wages lagged behind those in the Upper Midwest. Well-paid UAW members lost their jobs and foreign-owned automakers located engine and assembly plants in “right-to-work” states. Meanwhile, global supply networks expanded into service industries and functions.³³


³⁰ “Red Tape and US Health Care,” *Bloomberg Businessweek*, April 2, 2018, p. 18.

³¹ The employment figure is for mid-2017 in this broad occupational group and the industry category Health Care and Social Assistance, from the BLS Occupational Employment Statistics Query System, online.

³² Margaret Kyle and Heidi Williams, “Is American Health Care Uniquely Inefficient? Evidence from Prescription Drugs,” *American Economic Review Papers & Proceedings*, Vol. 107, No. 5 (2017), pp. 486-490

³³ For example, Marcel P. Timmer, Abdul Azeez Erumban, Bart Los, Robert Stehrer, and Gaaitzen J. de Vries, “Slicing Up Global Value Chains,” *Journal of Economic Perspectives*, Vol. 28, No. 2 (2014), pp. 99-118.

Figure 3
Workers With More Years of Schooling Earn More



Source: Economic Policy Institute, Washington, DC, based on Current Population Survey.

Taken together, technological change, deregulation, and globalization spurred economic dynamism and created new wealth as well as many new jobs. At the same time, millions of other Americans lost out, and conservatives in thrall to doctrines of individual responsibility and market fundamentalism aggravated workers' suffering, minimizing or rejecting the obligation of society to share the costs of economic restructuring. Economists have done their part, if less directly, through their efforts to explain the effects of IT and automation on wages with models based on "skill-biased technological change." These models aim to explain the widening of the wage distribution since the 1970s, with greater earnings for those in the topmost tiers while those lower down experience stagnant or declining pay. The idea of skill-biased technological change takes off from the abundance of statistical data showing that workers with more years of schooling earn more, on average (see Figure 3). These models do not show causality, for two reasons (Box II-A). First, workplace skills cannot be measured in meaningful ways, so that equating skill with years of schooling amounts to a heroic assumption. Second, lacking measures

of skill, new technologies cannot be shown to demand more of it, especially since empirical studies in actual workplaces show many new technologies call for less skill (examples include deskilling as a result of automatic machine tools in place of manual, and, affecting everyone with a driver's license, antilock brakes in cars and trucks enabling the least skilled to stop quickly without loss of control even on rain-slicked pavement).

Box II-A

Skill-Biased Technological Change: Computers in the Workplace

Seeking to explain growing disparity in pay between workers in the topmost tiers of the distribution and those lower down, economists have constructed a substantial edifice based on the idea of skill-biased technological change.^a Their models postulate increases in labor market demand for high-level skills as a result of technological change, pushing up wages for workers with such skills and reducing demand for workers lacking such skills.

New technology has generally been taken to mean computers and robotics in the workplace.^b Statistics show that workers in the upper tiers of the wage distribution have more education, on average, and this is assumed to confer higher levels of skill, with years of schooling taken to be an adequate indicator of actual skills exercised in actual workplaces. The models also assume that employers pay workers according to their skills, in accordance with contributions to organizational objectives and the value of the output they produce.

Similar frameworks provide a basis for recent efforts to anticipate the effects of AI, expected in the future to encroach on higher-level cognitive skills formerly beyond the capabilities of computers, even including scientific research itself, and extending capabilities of robotics through enhanced perceptual capabilities (e.g., image processing and identification), mobility, and dexterity. Predictions follow of vanished jobs, with human workers crowded into narrower segments of the labor market, supply exceeding demand and wages falling.

The popularity of such models has pushed aside older debates about deskilling and upskilling—whether, when, and how new technologies alter actual patterns of skill demand across industries and occupations. Upskilling has simply been assumed, based on rising wage premia for those with more years in school. To be sure, economists, like sociologists, recognize that skills come in many varieties, differing among tasks, occupations, and employers—who can organize work in many ways. But the models have no way of accommodating this sort of complexity. Although data in the form of results on tests of literacy and numeracy can be collected in straightforward fashion, these skills are too generic to be linked with detailed occupational categories, of which there are many.^c More specialized skills are likely to be murky, unmeasurable, and unobservable, with exceptions as for chess players and baseball pitchers (along with others including hackers), and even in these cases the skills themselves cannot be reliably assessed, only the results of skillful play. And while statistics on years of education completed are widely available—and can be linked with occupations through surveys—schooling contributes far more to some sorts of skills than others. Even today, for instance, when the practice of engineering depends heavily on science and mathematics, more than 15 percent of those employed as engineers (excluding technicians, placed in a different occupational category) cannot claim a 4-year college degree, and many of those with a degree earned it in some nontechnical field.^d (How many hackers have computer science degrees?) Heterogeneity of these sorts cannot be easily incorporated in quantitative models, for reasons beginning with the absence of data, which is why modelers have generally taken “skill” to be a unitary attribute measured by years of schooling, disassociated from tasks and occupations as well as from subjects of study.

In hiring, then, employers rely on educational credentials not so much because they signal particular skills (with exceptions, as for accountants) but because they signal traits such as motivation, socialization, obedience, and the ability to learn.^e In most cases, moreover, employers do not know what or how employees, individually, contribute to organizational performance (group performance is easier to judge) and do not set wages and salaries on this basis but rather by job categories, with tiers within these categories set according to organizational norms (one reason why women and minorities do better in some firms than others).

Skill-biased technological change offers a parsimonious account: a monocausal explanation of a complicated phenomenon. Skill is the independent variable, wages the dependent variable. Almost certainly, the increases in wage

dispersion that have marked the US labor market since the 1970s have multiple causes, not just one. Models of skill-biased technological change mostly offer not so much an explanation as a plausibility argument. Since wages and education correlate, they feed views that the primary function of schooling is to prepare young people for the workplace. They also feed the assumption that with enough schooling young people will be adequately prepared for whatever the future may bring. Many people do of course go on to learn, to invent, innovate, and accommodate themselves to change over lengthy working lifetimes. What schooling has to do with this, beyond acquisition of basic skills (literacy, numeracy), is another matter entirely. What education should aspire to, and accomplishes imperfectly today, is to prepare people for the unknowable future, one that will bring much change that forecasters have not anticipated and could not have been expected to.

^a See, e.g., Chinhui Juhn, Kevin M. Murphy, and Brooks Pierce, “Wage Inequality and the Rise in Returns to Skill,” *Journal of Political Economy*, Vol. 101, No. 3 (1993), pp. 410-442; Yujia Liu and David B. Grusky, “The Payoff to Skill in the Third Industrial Revolution,” *American Journal of Sociology*, Vol. 118, No. 5 (2013), pp. 1330-1374. Less technical treatments include David H. Autor, “Why Are There Still So Many Jobs? The History and Future of Workplace Automation,” *Journal of Economic Perspectives*, Vol. 29, No. 3 (2015), pp. 3-30. For a critical account, see David Card and John E. DiNardo, “Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles,” *Journal of Labor Economics*, Vol. 20, No. 4 (2002), pp. 733-783.

^b David H. Autor, Lawrence F. Katz, and Alan B. Krueger, “Computing Inequality: Have Computers Changed the Labor Market?” *Quarterly Journal of Economics*, Vol. 113, No. 4 (1998), pp. 1169-1213; Timothy F. Bresnahan, “Computerisation and Wage Dispersion: An Analytical Reinterpretation,” *Economic Journal*, Vol. 109 (1999), pp. F390-F415; Daron Acemoglu and Pascual Restrepo, “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, Vol. 108, No. 6 (2018), pp. 1488-1542.

^c John A. Alic, “What We Mean When We Talk About Workforce Skills,” *Issues in Science and Technology*, Vol. 34, No. 3 (2018), pp. 30-36.

^d “Table 1.11 Educational Attainment for Workers 25 years and Older by Detailed Occupation, 2015-16,” online.

^e See, e.g., Michael Spence, “Signaling in Retrospect and the Informational Structure of Markets,” *American Economic Review*, Vol. 92, No. 3 (2002), pp. 434-459.

Still, these models have been taken by many to legitimize the widening of wage gaps and have led to policy prescriptions unlikely to reduce those gaps. By suggesting that what people learn in school is all that matters, many policymakers and the public have come to believe that more and better education would push up pay levels, increase mobility, and wash away other labor market ills. There is no persuasive evidence that any of this would follow (Box II-B), and the focus on education deflects attention from institutional problems in the labor market. These are broad and deep. For starters, the United States does little to support workers who find their jobs gone. Laid-off manufacturing workers (most of them male), whose plight can be linked to international trade, are entitled to extended unemployment benefits and job training—not much help, since if they manage to land a new position it usually pays less, sometimes only half as much (e.g., after moving from a unionized factory to work in retailing). For (often female) clerical workers laid off in dribs and drabs in financial services, retail, and offices, the United States offers virtually no assistance beyond six months of unemployment benefits (less in some

states). With displacement common and likely to rise because of AIs, this is not good enough. Reportedly, Amazon already has something over 100,000 robots working in its warehouses and distribution centers. The much-heralded internet-of-things—meaning “smart” interconnected hardware devices, whether household thermostats or autonomous vehicles or vibration sensors installed on jet engines to anticipate incipient failures in time for repairs—may have been overhyped, but still could affect the jobs of many craft and technical workers. Section IV will return to the question of robots on the march. First, a look at what will make them smarter.

Box II-B

Education: No Panacea

Public education has long served as a foundation for American prosperity. The schools have always had their critics, too. They struggle with conflicting demands, and today with budget constraints, particularly in school districts serving poorer segments of the population—those that stand to benefit the most from better educational services. Research over the past several decades, for example, shows that pre-kindergarten programs, delivered well, boost lifelong cognitive capacity and learning ability, hence later academic performance, and also job and career trajectories—even lifetime health.^a Strong arguments also support the extension of free public education past grade 12. Yet desirable as such steps may be, they cannot promise remedies for wage inequality, much less wealth inequality. The biggest reason is that inequality stems mostly from public policies and labor market institutions stacked against workers, factors largely unaffected by education levels.^b

Since the Reagan administration’s alarmist 1983 report *A Nation at Risk* we’ve been told repeatedly that America’s primary and secondary schools fail to prepare young people for work or for college.^c Employers picked up the refrain, complaining incessantly of skill shortages. Labor Secretary Alexander Acosta has stated that the United States “has more than six million open jobs, but some employers can’t find workers with the skills to fill them.”^d The National Association of Manufacturers claims “the skills gap challenge” is becoming “a full-blown workforce crisis.”^e When these complaints concern specialized skills, they can be less than realistic. One company, for example, responded to a state survey that it was “looking for experience in Atomic Force Microscopy.”^f On the other hand, employers rightfully expect decent basic skills, beginning with literacy and numeracy. Almost any employee, after all, may need to understand written instructions, if only safety manuals, perhaps compose emails, count and record the stock on warehouse shelves, make change at cash registers. Yet 35 years after *A Nation at Risk*—which by no means marked the beginning of pleas and proposals for reform—too many US schools still shortchange too many of their students. Beyond the basics, many employers claim the young people they hire lack “soft” skills, an amorphous designation that includes verbal communications and the ability to work cooperatively in groups. For example, a survey of “skill deficits identified by employers” in Northern Virginia found “communication was number 1, problem solving/critical thinking number 2, and relationship management number 3.”^g While these are related to basic skills—“relationship management” has to do with affability, persuasion, and otherwise keeping customers happy—they are not the same. And beyond these sorts of general categories, soft skills can be elaborated and extended almost indefinitely.

Indeed, skills whether “soft” or “hard” come in uncountable variety and, as discussed at many points in this report, are hard to observe, much less measure, and hence to teach effectively. In their hiring and promotion practices, many firms recognize these difficulties. Google, for example, once “recruited from elite schools like Stanford and MIT. But when Google examined its internal evidence, it found that grades, test scores, and a school’s pedigree weren’t a good predictor of job success. A significant number of executives had graduated from state schools or hadn’t completed college at all.”^h Still, job interviews being notoriously fallible, employers have little other than educational credentials to go on and treat these as signals of potential.ⁱ They assume, perhaps implicitly, that success in school points to desired traits: discipline (showing up); diligence (turning in assignments); agreeableness (please the teacher, please the boss); ability to learn; and social skills sufficient to stay out of too much trouble. In other words, employers hire, all else equal, on the assumption that job-seekers with more education are likely to be safer bets—whether for low-wage jobs in, say, the front office of a retail bank or

highly coveted positions in prestigious law firms, investment banks, and consulting firms.^j Suppose, then, that everyone in the workforce had a four-year college degree. What, exactly, would change? Would employers alter their hiring and promotion practices? Raise wages? Not likely. Rather, overeducation and underemployment, already common, would rise still more.

Workplace skills change rapidly, and schools have never pretended to do more than prepare young people, most of whom leave formal education behind in their twenties (if not before), for continuing learning after they leave. Many American workers manage to do so. Public policies provide little help. They should do more, a point to which we return in Section VI.

^a See, e.g., James J. Heckman, “Schools, Skills, and Synapses,” *Economic Inquiry*, Vol. 46, No. 3 (2008), pp. 289-324, which summarizes work conducted over many years by Heckman (a Nobel laureate in economics) and his colleagues on the persisting positive effects of early childhood education. Lengthier treatments include Tim Kautz, James J. Heckman, Ron Diris, Bas t. Weel, and Lex Borghans, *Fostering and Measuring Skills: Improving Cognitive and Non-cognitive Skills to Promote Lifetime Success*, OECD Education Working Papers No. 110 (Paris: OECD, 2014).

^b The remainder of this box draws on John A. Alic, “Beyond Schooling: Educating for the Unknowable Future,” *The Bridge*, Vol. 45, No. 3 (2015), pp. 29-36; and on Alic, “What We Mean When We Talk About Workforce Skills.”

^c *A Nation at Risk: The Imperative for Educational Reform* (Washington, DC: Government Printing Office, April 1983).

^d Quoted in Alyson Klein, “Betsy DeVos: Stop ‘Forcing’ Four-Year Degrees as Only Pathway to Success,” *Education Week*, November 13, 2017.

^e “NAM Manufacturer’s Outlook Survey, Third Quarter 2018,” October 5, 2018, online; quoted phrases from p. 2.

^f Alessia Leibert, “Are Skilled Workers Scarce? Evidence from Employer Surveys in Minnesota,” *Minnesota Employment Review*, May 2013, pp. 1-8; quotation from p. 5. As respondent noted, “It is hard to find people with Mainframe skills (older skills like COBOL that are no longer taught)” (p. 6).

^g Kenan Patrick Jarboe and Steve Olson, rapporteurs, *Adaptability of the US Engineering and Technical Workforce: Proceedings of a Workshop* (Washington, DC: National Academies Press, 2018), p. 48.

^h Jennifer Alsever, “How AI Is Changing Your Job Hunt,” *Fortune*, May 19, 2017, online.

ⁱ Jason Dana, “The Utter Uselessness of Job Interviews,” *New York Times*, April 8, 2017, online.

^j Lauren A. Rivera, “Hiring as Cultural Matching: The Case of Elite Professional Service Firms,” *American Sociological Review*, Vol. 77, No. 6 (2012), pp. 999-1022.

III. Artificial Intelligence

A visitor to our planet ... would read and hear all about wonderful 'mechanical brains' baffling their creators with prodigious intellectual performance. And he (or it) would be warned that these machines must be restrained, lest they overwhelm us by might, persuasion, or even by the revelation of truths too terrible to be borne.

~ Marvin Minsky, 1961¹

In the fall of 2017, DeepMind, a Google subsidiary, reported that its program AlphaGo Zero had bested all comers in the challenging board game known as Go. The previous AI champion, AlphaGo, also a DeepMind product, had trained on 100,000 games played by humans—a form of machine learning (ML) based on large sets of labeled data, in this case moves and countermoves. AlphaGo Zero adopted a different approach, dispensing with training data. As the “zero” implies, AlphaGo Zero learned from scratch, starting with no prior knowledge other than the rules of Go and playing 30 million games against itself.²

The term artificial intelligence has been used quite widely since the middle 1950s, and many AI applications have now been anticipated for half a century or more.³ These include computers as aids in medical diagnosis and health care administration; classroom “teaching machines”; and determination of molecular structures.⁴ Only the last of these applications could be considered fully realized today. Schoolchildren in wealthy countries use the web easily, but guided instruction remains a chimera. Frustrated motorists hoping automated vehicles will reduce congestion may wonder what happened to predictions of computer control of stoplights and just-in-time re-routings around traffic jams caused by accidents or a police cruiser stopped to issue a citation.

¹ Marvin Minsky, “Steps Toward Artificial Intelligence,” *Proceedings of the IRE*, Vol. 49, No. 1 (1961), pp. 8-30. In 1959, Minsky and John McCarthy had established the MIT Artificial Intelligence Laboratory.

² David Silver, et al., “Mastering the Game of Go Without Human Knowledge,” *Nature*, Vol. 550 (October 17, 2017), pp. 354-359. For an abbreviated glossary of AI-related terms, see “Artificial Intelligence, In So Many Words,” *Science*, Vol. 357, No. 6346 (2017), p. 19.

³ In 1955, several computer scientists, including John McCarthy and Marvin Minsky, wrote a proposal in search of funding for a “Summer Research Project on Artificial Intelligence” and organized a now-famous conference that took place in 1956. Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements* (New York: Cambridge University Press, 2010), p. 53.

⁴ W. F. Bauer, D. L. Gerlough, and J. W. Granholm, “Advanced Computer Applications,” *Proceedings of the IRE*, Vol. 49, No. 1 (1961), pp. 296-304. This article appeared in a special issue on “Advanced Computer Applications.”

The first step in understanding the implications of AIs for jobs is to grasp why and how their capabilities have grown so rapidly since around 2010. In principle, the basic task for almost any AI is simple to describe: to take in data and information, such as the positions of tokens in a game of Go, and to “think”—in this case to decide upon the next move. These sorts of processes are known as inferences. Science knows relatively little of the how and why of human inferences.⁵ Sometimes they follow years of effortful thinking, as in Darwin’s stumbling toward his theory of natural selection. Other inferences come suddenly, with little or no conscious thought, as in the shock of recognition when we glimpse a face from the past. Despite ML accomplishments such as that of AlphaGo Zero, which devised “Go strategies that human players haven’t invented in thousands of years,” it would be wrong to think of any AI as paralleling, in some meaningful way, either human intelligence or human learning.⁶

AI inputs can take almost any form that is reducible to digital bits and bytes. In software for language translation, the input is a digital representation of each letter, space, and punctuation mark. The AI seeks to make sense of the feed, by converting letters and spaces into sequences of words and assembling these into sentences. The hard part comes not in finding an equivalent for each word, but in arriving at a translation that makes sense, since languages differ in grammar, syntax, and nuances of common usage, and meaning may depend on context. For processing digital images the difficulties lie in recognition: telling a shaggy dog from a sheep, or a harmlessly empty trash bag captured by a camera on an AV from a full bag that might merit a sudden swerve.

From the 1950s into the 21st century, AI enthusiasts mostly overpromised and underdelivered.⁷ Although pattern recognition is a longstanding subfield in computer science, rapid advances began only around 2010. Declining costs for high-performance digital

⁵ See, e.g., John A. Bargh and Tanya L. Chartrand, “The Unbearable Automaticity of Being,” *American Psychologist*, Vol. 54, No. 7, (1999), pp. 462-479; Anthony G. Greenwald and Mahzarin R. Banaji, “The Implicit Revolution: Reconceiving the Relation Between Conscious and Unconscious,” *American Psychologist*, Vol. 72, No. 9 (2017), pp. 861-871.

⁶ The quotation is from Elizabeth Gibney, “Self-taught AI Is Best Yet at Strategy Game Go,” *Nature News*, October 18, 2017, online.

⁷ More than 60 years ago, Carnegie Mellon University’s Herbert Simon and Allen Newell wrote “there are now in the world machines that think, that learn, and that create.” Herbert A. Simon and Allen Newell, “Heuristic Problem Solving: The Next Advance in Operations Research,” *Operations Research*, Vol. 6, No. 1 (1958), pp. 1-10; quotation from p. 8. That would seem a bold statement even now.

hardware—the same sorts of hardware advances that underlie the latest smartphones—have made practical approaches to AI such as neural networks (see Box III-A), limited earlier by computing costs and capacity. These systems can take in, hold in memory, and manipulate almost unimaginable volumes of data and information, “big data.” Algorithms—computational procedures—then operate on this data, searching through it for patterns of one sort or another (e.g., in the images to be identified, sentences to be translated), extracting meaning, making decisions. AIs rely on probability and statistics for sense-making and decision-making. Their computational routines can be enormously complex, so much so as to be opaque even to those who built them. Still, whatever it is that AIs do, they follow computational routines embedded in software—coded rules. Humans, on the other hand, may know the “rules” intended to govern behavior—don’t talk in class, obey the speed limit—yet ignore or evade them.

Box III-A

Hardware and Software for Machine Learning

Checkers and chess, simpler games than Go, provided early testbeds for AI.^a The usual approach was to play out alternative sequences of moves and countermoves—something computers can do far faster and more accurately than people—and select the most promising. This is a brute force technique quite unlike the methods of good human chess players, who rely instead on heuristics, strategies based heavily on past experience and often emerging with little conscious awareness from intuitions and other tacit processes.^b Not until the late 1990s did a brute-force chess program, IBM’s Deep Blue, vanquish the best human chess masters.

Hardware advances—greater processing speed, more memory—made possible Deep Blue’s accomplishment, as well as others that came later. While early computers offered big gains in speed over the adding machines and electro-mechanical punched card equipment they supplanted, their performance soon came to seem primitive. Since about 1950, computing speed has increased by perhaps 100 trillion times. These gains resulted from a very large number of mostly small advances in hardware architecture and componentry—transistors, integrated circuits, disk storage—and also in programming languages and applications software. Since the 1970s and 1980s almost limitless computing power has become amazingly inexpensive. Semiconductor firms also developed specialized chips known as graphics processing units, originally for video games, that can run AI programs thousands of times faster than the general-purpose microprocessors used in PCs and servers. Performance metrics for hardware componentry such as hard disk drives have in some cases improved even faster than for ICs, with the costs for storing a gigabyte (a billion bytes) of digital data falling from around \$200,000 in 1980 to under 2 cents today.^c

Decade by decade, computing power became cheaper, putting capabilities once available only in supercomputers on a desktop. Software was a bottleneck, limited by slow and error-prone line-by-line coding and also by the complexities of designing large-scale programs and managing development, debugging, troubleshooting, and validation.^d While most metrics for software performance show very slow gains compared to hardware, increases in hardware performance help, particularly with ML programs, since once the basic code has been developed the system in essence reprograms itself.

Falling hardware costs brought renewed interest in neural networks, an approach to AI explored in the early years of computing and then largely abandoned.^e Neural nets can be visualized as a large number of interconnected nodes, or (electronic) neurons, at which signals arrive to be combined or otherwise altered before being passed on to other nodes. The simpler types of two-dimensional networks might be imagined as something like huge tennis nets. Greater computational resources made possible three-dimensional networks (speaking metaphorically), with vertical connections among hundreds of two-dimensional nets.

Although neurons in living brains—some 85 billion in the brains of humans and many billions more in some

animals—suggested the neural network label, this, if evocative, is otherwise misleading, since there is no known (bio) physical resemblance with living brains. From the beginning, computer scientists have insisted that the metaphor of computers as giant brains is false: the idea of AI is to emulate aspects of the mind, not the brain.^f For that reason alone, the idea of “artificial intelligence” seems misleading

Reduced to practice, tasks for ML systems based on neural networks can take many forms: image recognition (animals, human faces, fingerprints); selecting moves in chess or Go; flagging suspicious financial transactions; examining radiographic scans for signs of cancer. Neural networks improve performance on their assigned tasks by incrementally adjusting weighting parameters to bring outputs into closer correspondence with responses deemed correct: to recognize an electronically-scanned image as a dog rather than a sheep, to translate a sentence from Russian or Arabic that gets across the meaning. No one has to rewrite the software; it changes itself, by adjusting the statistical weighting factors. After millions of repetitions with labeled “training data” consisting of paired inputs and correct outputs, the system should be able to tell an image of a shorn sheep from a goat or a Weimaraner. These systems have registered remarkable successes, such as a “121-layer convolutional neural network that inputs a chest X-ray image and outputs the probability of pneumonia,” demonstrating performance superior to that of “average” human radiologists after training on “112,120 frontal-view chest X-ray images individually labeled with up to 14 different thoracic diseases, including pneumonia.”^g Still, these sorts of tasks present the AI with data from a structured environment: the system knows it is looking at a chest or face or a fingerprint, playing Go and not chess or poker.

^a For a lucid discussion from the earliest years of electronic digital computing, see Claude E. Shannon, “Programming a Computer for Playing Chess,” *Philosophical Magazine*, Vol. 41, No. 314 (1950), pp. 256-275.

^b Valerie Thompson, “What Intuitions Are ... and Are Not,” *Psychology of Learning and Motivation*, Vol. 60 (2014), pp. 35-75. On heuristics for chess, see, e.g.,

^c Lucas Mearian, “CW@50: Data Storage Goes from \$1M to 2 cents per Gigabyte,” *Computerworld*, 2017, online.

^d See, e.g., Frederick P. Brooks, Jr., *The Design of Design: Essays from a Computer Scientist* (Upper Saddle River, NJ: Addison-Wesley, 2010).

^e Nils J. Nilsson, *The Quest for Artificial Intelligence: A History of Ideas and Achievements* (New York: Cambridge University Press, 2010). For an accessible and more recent account, see “Annex: Canonical Problems in Machine Learning,” pp. 120-121 in *Machine Learning: The Power and Promise of Computers That Learn by Example* (London: Royal Society, April 2017).

^f Joshua B. Tenenbaum, Charles Kemp, Thomas L. Griffiths and Noah D. Goodman, “How to Grow a Mind: Statistics, Structure, and Abstraction,” *Science*, Vol. 331, No. 6022 (2011), pp. 1279-1285; Brenden M. Lake, Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman, “Building Machines that Learn and Think Like People,” *Behavioral and Brain Sciences*, Vol. 40 (2017), pp. 1-72.

^g Pranav Rajpurkar, et al., “CheXNet: Radiologist-Level Pneumonia Detection on Chest X-Rays with Deep Learning,” *arXiv:1711.05225v2 [cs.CV]* (November 25, 2017), online; quotations from first page.

Put simply, AIs are machines and people are not, in the usual meanings of the term. Much human behavior has little to do with intelligence, as psychologists use the term (Box III-B). AIs will remain mechanical, mindless, so far as we know. Basically, what AIs do is automate decision-making within defined domains. They accomplish this drawing on technical disciplines including mathematical optimization, cybernetics and feedback control, computer science and engineering, and as already noted, probability and statistics. The domains for these accomplishments include those already mentioned—game playing; medical diagnosis; driving road vehicles (and piloting aircraft)—and many others. It is easy to extend the list to workplace

tasks susceptible to automation. For some time, IT systems have screened credit card transactions in search of questionable transactions; they do so based on algorithms that search for deviations from past purchasing patterns. Loan officers in many banks now call on AI systems for first-cut or even final-cut determinations of the creditworthiness of homeowners and small business proprietors (“automated credit scoring”). These and other sorts of AIs function more or less well within their intended domains. They are useless outside them: no chess-playing robot could fly a plane; that takes a different set of algorithms and knowledge base (rules of chess; allowable inputs to the flight controls of a 737 or a crop duster).

Box III-B

Decision-Making and Intelligence

In the early 1950s, in an article titled “Can Machines Think?,” Maurice Wilkes, lead designer on one of the earliest digital computers, asked, “Is the brain a machine?” He answered his own question by saying this was a matter “on which science has hardly begun to touch.”^a Today we know a great deal more about human and animal behavior, but not about intelligence in any very fundamental sense. For example, psychologists have been exploring the human mind for more than a century, but work on relating the results of intelligence tests to brain function has barely started.^b Computer scientists and engineers achieve reasonable simulations of “intelligent” behavior within restricted domains such as playing games. But machine learning and other approaches to AI make no attempt to replicate brain functions.

Unlike much else in life, games such as chess and Go have explicit rules, known to all and enforceable. This makes them exceptional. Psychologists and cognitive scientists know a good deal about human decision-making, including people’s propensities to follow written and unwritten rules and to bend or break these, as they frequently do. For this reason alone, game-playing does not go very far toward emulation of human intelligence and human behavior. The AV problem illustrates. There are plenty of rules for driving motor vehicles, some written into law and enforced mostly in the breach, others unwritten, the sort of thing driving instructors try to get across to learners. Some are more like common sense, others are more like behavioral norms—and the latter, like traffic laws, vary from place to place. Despite the possibilities for confusion, people mostly manage to drive about without too many accidents through some mix of rule-following and situational rule-bending, instinct, reflexes, and peripheral vision. AIs can in fact “learn” in such settings, if allowed. How well they will deal with anomalous or rarely encountered situations remains to be determined. No more than a human can an AV be allowed to come to a halt on an urban freeway if terminally confused.

“Intelligence” as measured by IQ and other tests of cognition is but a small part of human decision-making. Some tests stress speed, others accuracy; some emphasize recall of factual knowledge, others the ability to reason from facts to conclusions.^c Measures of intelligence capture some aspects of rationality, such as the extent to which a person’s beliefs about the world correspond to the world as empirically known. But intelligence tests make no attempt to measure other aspects of rationality, such as the extent to which behavioral choices lead to goals an individual claims to value.^d People, after all, sometimes do stupid things; reckless driving is irrational, since it can lead to arrest, accident, even death, yet is hardly restricted to low-IQ daredevils. On the other hand, no AV should drive recklessly, not even once in a lifetime. Some intelligence tests aim to capture other sorts of hard-to-measure attributes, such as the ability to plan for the future and to think and behave flexibly in unpredictable settings. These go well beyond the sorts of capabilities commonly associated with AI, at least currently. In some sense, after all, the idea of planning for an indeterminate future—one in which a Go-playing AI might be asked, out of the blue, to learn to control an AV—“does not compute.” Only if some large software package had been designed so that it could branch from one set of tasks to another (in effect comprising two or more programs linked together) might this be possible. Yet what would be the point? This is not to say that very general AI capabilities might not be possible in the future, or that some sort of universal AI is inconceivable. It is just that at this point no one has much idea of how such an AI might be built.

Some of the hurdles in AI have to do with inference and induction: moving from digital input data—an AV’s surroundings, the gyrations of the stock market, video from one of the military’s airborne drones—to actionable

conclusions. Hit the AV's brakes, buy Google and sell Facebook, abort the attack. Because we know little about how humans make such decisions, absent explicit and well-understood rules—such as statistically validated algorithms for stock trading or the decision rules the US military has adopted to govern drone attacks—much of AI development amounts to a sophisticated form of trial and error.

More serious hurdles have to do with endowing AI systems, unless restricted to narrow and well-structured domains such as game playing, with general knowledge of the world and how it works. Everyone, even an infant a few weeks old, has some familiarity with their surroundings and can think in ways that help figure out what is going on around them. Babies quickly learn to recognize their parents and siblings, to communicate desire for comforting or for food. As this suggests, much animal learning is experiential (and some is instinctual): humans pick up knowledge and skill from birth and throughout a lifetime. Most computer scientists have come to believe that there is no practical way to pack more than a small fraction of what people know about the world and how it works into digital systems, the more so because some particular bit of knowledge may be needed quickly, under unpredictable circumstances. It is easy enough to teach an AI system that time flies and so do flies; that depending on context, “cubs” refers to Chicago’s baseball team and not young animals; that “Ernie Banks” refers to one of the team’s famous players of an earlier generation and not some sort of earthwork. These are simple facts that can be stored in memory, of the sort that allowed IBM’s Watson to win against past Jeopardy champions in 2011. The problem is that such facts can be multiplied almost endlessly, and even if they have been stored in memory somewhere, how is the machine to know what it needs when it needs some piece of information? Almost any American adult will know at least some of the following: that Earth’s moon is singular; variably large in the night sky and showing a smiley face when full; has sometimes been said to be made of cheese (green cheese, at that); is the driving force of ocean tides; from time to time causes solar eclipses; was the goal of the Apollo project launched by President John F. Kennedy in 1961 and attained in 1969 when two American astronauts walked on its surface; that some people believe the landing was faked; that Michael Jackson made a dance he called the moonwalk famous. AI developers have struggled endlessly to figure out how to get everything a computer might at some time need to know into its memory (where do you stop?) and, the truly difficult part, how to categorize for call up when relevant. Humans can do much of this without deliberate thought or with casual effort. Sometimes they need to think about things—maybe talk them over with friends. But how might an AI system fitted with modules for image processing and speech recognition come to understand a Monty Python skit well enough to prepare a summary of the sort a high school student might write (without cribbing from the web) or, to take a more recent example, explain the title *The Shape of Water*?^e

^a M.A. Wilkes, “Can Machines Think?” *Proceedings of the IRE* Vol. 41, No. 10 (1953), pp. 1230-1234. Recent discussions of related topics include Christof Koch, “What is Consciousness?” *Nature*, Vol. 557 (May 10, 2018), pp. S5-S12.

^b Ian J. Deary, “Intelligence,” *Current Biology*, Vol. 23, No. 16 (2012), R674.

^c The more widely accepted measures, including IQ, correlate reasonably well with one another, giving rise to what is usually considered the most meaningful overall gage of cognitive ability, termed g for “general intelligence.” Richard E. Nisbett, Joshua Aronson, Clancy Blair, William Dickens, James Flynn, Diane F. Halpern, and Eric Turkheimer, “Intelligence: New Findings and Theoretical Developments,” *American Psychologist*, Vol. 67, No. 2 (2012), pp. 130-159.

^d Keith E. Stanovich, “The Comprehensive Assessment of Rational Thinking,” *Educational Psychologist*, Vol. 51, No. 1 (2016), pp. 23-34.

^e These and other examples are suggested by Ernest Davis and Gary Marcus, “Commonsense Reasoning and Commonsense Knowledge in Artificial Intelligence,” *Communications of the ACM*, Vol. 58, No. 9 (2015), pp. 92-103. The journal has placed a header above this article reading “AI has seen great advances of many kinds recently, but there is one critical area where progress has been extremely slow: ordinary commonsense [sic].” On AI and common sense also see Matthew Hutson, “Basic Instincts,” *Science*, Vol. 360, No. 6391 (2018), pp. 84-847.

Humans are general-purpose thinkers, problem solvers, decision makers. AIs are not, certainly not today. As machines, they remain specialized. It is not so much that a general purpose AI is inconceivable, it is that no one at this point has much idea how one might be built. The less than transparent aspects of problem-solving by humans, moreover, have thwarted

efforts to build domain-specific AIs such as expert systems. As explained in Box III-C, expert systems aim to capture the know-how of people skilled in tackling a given class of problems. Development stumbled because the experts often could not say how or why they proceeded as they did.

Box III-C

Expert Systems: Buoyant Hopes Deflated

An approach to AI known as expert systems attracted several billion dollars in R&D funding from the 1970s into the 1990s, in anticipation of applications in medicine, numerous industries, and the military. Much was learned technically, but the initial optimism proved misplaced.

Expert systems, also known as knowledge-based systems (KBSs), do not in any sense learn. Instead, they embed human know-how in directly programmed rules. These rules are intended to codify procedures followed by human experts. An everyday example: if your lawnmower engine won't start, check the gas tank, then the spark plug. KBS development, in other words, aims at decision support rather than full automation.

By interviewing and observing practitioners acknowledged by their peers to be especially capable—technicians and mechanics, geophysicists searching for exploitable oil deposits, physicians diagnosing some puzzling combination of symptoms—the know-how, skills, and thought processes of experts were to be uncovered and made the basis of a decision tree, as in the lawnmower example.^a Those with less skill and experience could then follow along, employing similar logic.

Defense Funding

The US Defense Department, which has a great deal of long-lived, costly, and complex equipment to maintain—warships, armored vehicles, aircraft (if the jet engine won't light off, check ...)—funded much KBS work.^b Other Pentagon initiatives had far more ambitious goals. The Strategic Computing Program, begun in the early 1980s by the Defense Advanced Research Projects Agency (DARPA) and running for a decade, spent around \$1 billion on AI-related R&D centered on expert systems and related technologies and applications.^c DARPA managers tailored three major themes to appeal to the military services: a “pilot's associate” to provide decision support for Air Force pilots in aerial combat, intended in part to deal with sensory overload; battle management software for the Navy to automate command and control of bristling shipboard arsenals and their accompanying networks of radars, sonars, and encrypted intra-fleet communication links; and robotic vehicles for the Army aimed at driverless supply trucks to bring up food, water, fuel, and munitions to troops in the field (R&D that built foundations for later AV design and development by non-defense firms such as Waymo and led to the DARPA-sponsored AV design contests that began in 2004).

The Bottleneck: Tacit Knowledge and Skill

Once viewed as highly promising, the expert systems approach to AI foundered for two chief reasons. First, the rules and procedures followed by human experts proved vexingly difficult to pin down. “Knowledge engineers” came to realize that high-level competency depended on instinct, heuristics, and tacit knowledge, much as in games such as chess. Experts sometimes solved problems following no discernible plans and procedures, or else these were so deeply submerged that even the most painstaking efforts failed to reveal them. Despite its aura of mystery, instinct is a perfectly respectable label among cognitive scientists, and heuristics even more so among computer scientists. They refer to mental processes that involve little or no conscious thought, also called tacit knowledge—things people know, or think they know, but cannot fully articulate. People use these sorts of decisional shortcuts all the time, but where they come from and how they are learned remain obscure.

Neural networks and machine learning skirt some of these obstacles by finding their own rules. But they cannot yet deal with the second set of problems encountered during expert systems development: building in very large stores of general knowledge to accompany rule-following behavior and ensure that the KBS does not go astray and reach incorrect, even absurd, conclusions—absurd to humans but perfectly “logical” to an AI that knows no better. These sorts of difficulties, associated with what people think of as common sense, afflict essentially all AIs outside tightly-constrained domains such as playing board games. As Edward Feigenbaum, a leading figure in computer science for decades, explained:

Eventually, this line of work hit a wall. It hit the wall of brittleness, when you go outside the area of specialization, the domain. The programs don't work right, because They don't know enough about the rest of the world, so they break. But people have vast stores of common sense with which they can paper over these breaks in the reasoning process.^d

In another interview, when asked "What are some of the big problems that need to be solved in AI?" Feigenbaum replied:

We need to have a way in which computers can read books on chemistry and learn chemistry. Or read books on physics and learn physics. Or biology. Or whatever. We just don't do that today.^e

Clinical AI illustrates: medicine knows "more than 100,000 diseases and conditions, thousands of guidelines, and millions of rules governing them. The rules ... lie buried inside a sea of facts that is constantly being modified by the clinical literature."^f Even if, say, the World Health Organization's International Classification of Diseases (listing more than 150,000 diseases and combinations currently) were available for an AI to "read," the system would need to understand much else besides. You don't have to be a physician to know that lungs and kidneys, but not brains and livers, come in pairs, that most people (but not all) have five fingers on each hand, and so on. For the AI, all this and much more must be spelled out; nothing that might matter in some particular situation can be overlooked.

^a Judith Reitman Olson and Henry H. Rueter, "Extracting Expertise from Experts: Methods for Knowledge Acquisition," *Expert Systems*, Vol. 4 (1987), pp. 152-168. For an example, diagnosis of quality issues in plastic injection molding such as voids and porosity, see John Kingston, "KBS Methodology as a Framework for Co-Operative Working," in M.A. Bramer and R.W. Milne, eds., *Research and Development in Expert Systems IX* (Cambridge, UK: Cambridge University Press, 1993), pp. 45-60. Parts of this box draw from John A. Alic, "Technical Knowledge and Experiential Learning: What People Know and Can Do," *Technology Analysis & Strategic Management*, Vol. 20 (2008), pp. 427-442.

^b J. Jeffrey Richardson, Robert A. Keller, Roy A. Maxion, Peter G. Poison, and Kenneth A. DeJong, *Artificial Intelligence in Maintenance: Synthesis of Technical Issues*, AFHRL-TR-85-7 (Brooks Air Force Base, TX: Air Force Human Resources Laboratory, October 1985).

^c Alex Roland and Philip Shiman, *Strategic Computing: DARPA and the Quest for Machine Intelligence, 1983-1993* (Cambridge, MA: MIT Press, 2002). For a summary of DARPA's recent planning for "third-wave" AI, see pp. 4-5 in "Statement by Dr. Steven Walker, Director, Defense Advanced Research Projects Agency," Before the Subcommittee on Emerging Threats and Capabilities, Armed Services Committee, US House of Representatives, March 14, 2018, online.

^d Robin M. Hogarth, "Intuition: A Challenge for Psychological Research on Decision Making," *Psychological Inquiry*, Vol. 21, No. 4 (2010), pp. 338-353.

^e "Interview of Edward Feigenbaum by John Plutte," CHM Reference number: X6409.2012, Computer History Museum, Mountain View, CA, March 20, 2012, online, p. 6.

^f "Oral History of Edward Feigenbaum by Nils Nilsson," CHM Reference number: X3896.2007, Computer History Museum, Mountain View, CA, June 20 and June 27, 2007, online, p. 64.

^g Floyd E. Bloom, "Science as a Way of Life: Perplexities of a Physician-Scientist," *Science*, Vol. 300, No. 5626 (2003), pp. 1680-1685; quotation from p. 1684.

The design of systems of almost any sort, including AI systems, hinges on some sort of objective function, a predetermined goal or goals around which the design is to be configured. The objective function is ready-made for games, straightforward to specify for experts systems or reading medical X-rays and MRI scans, for many applications of robotics, and reasonably clear-cut for AVs and some business processes, such as back-office processing of financial transactions. A bank can prepare a statement of objectives (for its internal staff or a consulting firm) for a system that will process, say, checks or credit card transactions based on minimum

costs, including those of finding and correcting errors. Something similar can probably be accomplished for a considerable range of loans, to both households and businesses, based on risks. Still, risk analysis depends on plentiful data for statistical analysis, and it is one thing, for example, to write a life insurance policy based on actuarial tables, quite another to estimate the probability that rising sea levels due to future climate change will swamp a waterfront golf course before a multimillion dollar loan has been paid off (with the declining popularity of golf adding to the uncertainties). There may no more to go on for investment banking deals. How does an AI or a human expert estimate the risks of bad behavior by a superstar CEO that may bring down the entire business? The extent to which AI will take over tasks and jobs, then, depends on the details of products to be delivered, and these come with many variations. Some tasks and jobs will be far better insulated from automation than others.

In many parts of the service economy, those in which production centers on “the interpretive model” laid out in *New Rules for a New Economy*, the initial objective is likely to be ill-defined—what the customer really wants—and this will restrict the scope of AI assistance.⁸ With exceptions for standardized products, as for example in much of retail banking and fast foods, attributes of service products emerge in the course of production, through a process of discussion and negotiation between service providers and customers. Even if the process begins with a template of some sort—a chef’s recipe, loan officer’s checklist, third-grade teacher’s lesson plan—the service as delivered takes shape during production (not before, as in goods-producing industries). (In the “engineering model” for manufactured goods, by contrast, products are designed and developed in advance of production by engineers, technicians, and managers; product attributes and processing steps are predetermined and prescribed in all details, leaving nothing of significance to the discretion of production workers.)

Clinical encounters in health care illustrate the interpretive model. A typical process might begin with dialog between physician and patient concerning the patient’s history and symptoms. Part of the history—the patient’s medical record or file—may be stored in a computer system, but there are likely to be gaps and errors in any such file. Error rates in computerized financial records, once scrubbed, are few; in medical records they are many and go mostly

⁸ Stephen A. Herzenberg, John A. Alic, and Howard Wial, *New Rules for a New Economy: Employment and Opportunity in Postindustrial America* (Ithaca, NY: Cornell University Press, 1998), pp. 85-94.

unidentified. Each encounter, moreover, is likely to generate new information, possibly in contradiction with old information and possibly introducing further errors. As the process continues, diagnostic tests may be called for, perhaps an MRI scan or even a genomic analysis. Clinical AI may help with some of this. If the patient's "presentation" is a puzzle, perhaps AI can also help with diagnosis and an initial plan of treatment (which may later be revised, depending on the patient's response). Plainly, devising and refining some sort of objective function for a general purpose health care AI would be enormously difficult. Current systems simply provide help that health care practitioners can call up if they wish, and many past attempts to provide support have floundered for reasons of awkwardness, inconvenience, or poor integration with existing work practices.⁹

Further, because of their complexity, the actual behavior of an AI system, what it does once installed and working, may or may not be what the developers wanted and expected. While people judge likelihoods all the time—the stoplight is going to change, it's not going to rain—they do so without conscious calculation. Humans formulate expectations more-or-less continuously and more-or-less unconsciously—that is, automatically—about what is likely to happen next. An AI, on the other hand, such as that charged with driving an AV (Box III-D), must be told or must learn through machine processes what to look for and how to respond. To be sure, parents tell their children to be careful because the traffic light may change, and experience teaches us to read the signs foretelling this afternoon's weather. Yet humans break rules all the time, while AIs follow the rules, whether these have been programmed directly by human coders or stem from ML processes. (ML rules evolve, but the AI will not break whatever rules govern its outputs at a given point in time). ML systems, for example, are "black boxes" and even if trained on identical data sets, their software will "learn" differently, perhaps in trivial ways, perhaps in significant. For this reason, identical inputs may not result in identical outputs (or they may; there is no way to know in advance). One AV taking in data from multiple sensors might, for example, swerve to avoid a trash bag and another AV, differing only in its ever-evolving learning history, might not (depending, among other things on how many trash bags each had seen and what had happened). In the general case it will be impossible to say why outputs or outcomes differ,

⁹ See, e.g., Derek W Meeks, et al., "An Analysis of Electronic Health Record-Related Patient Safety Concerns," *Journal of the American Medical Informatics Association*, Vol. 21 (2014), pp. 1053-1059.

Box III-D

Autonomous Vehicles: Sensors, Software, and Safety

AVs must “see” or otherwise sense their surroundings.^a They need situational awareness similar to that of humans who may catch a glimpse of a child or a bicycle or a dog on their right, a car exiting a driveway that might or might not stop, an erratically piloted vehicle with perhaps a drunk behind the wheel. Many of these situations are variations on a theme, and drivers learn to anticipate them. Others are statistically rare—not that many motorists have seen lightning strike a roadside tree or encountered a deer dashing past their front bumper. Also like a human driver, the AV must sense its surroundings at night, in rain, snow, or fog. It has to know where it is, where it is going, and how to get there. Still, unpredictable human behavior by drivers of conventional vehicles and by others in the vicinity may prove the most difficult technical challenge for AV control.

Human drivers rely heavily on sight and occasionally on hearing (horns, screeching tires). AVs will carry multiple sensors operating on a variety of physical principles—e.g., GPS signals, video (visible light), radar (reflected radio waves), laser-imaging radar or lidar (light detection and ranging with reflecting laser beams)—to build a picture of their surroundings. The sensor suite must provide reliable wide-area coverage regardless of conditions—fog, deep darkness or sunset glare, snow-covered roads.^b Available sensors vary in capabilities and costs. GPS signals vanish in tunnels and may be erratic among tall buildings in cities. Lidar operates on principles not unlike those of radar but gives a 3-dimensional picture rather than a blob-like image. But while radars can easily detect objects beyond the ranges necessary for safe operation of a road vehicle, lidars can only see relatively short distances. And although widely used in AV development and testing, lidars remain much too costly for production vehicles.^c

Redundancy is necessary for safe operation. A partially automated vehicle can kick control back to a human driver if it doesn’t know what to do or if the software crashes; ideally, the system will transfer control in time for the driver to pick up cues as to what is happening, decide how to respond, and do so appropriately.^d Humans can cope if the speedometer stops working; some might never notice. An AV must know how fast it is going. If, say, a speed sensor such as a radar pointing at the road surface no longer functions, GPS can still indicate speed, though less accurately. But what if the main electrical power supply shorts out? Such contingencies call for “soft” failure modes.

There must be redundancies in processing too. Software will compare and combine sensor signals through data fusion (as our brains and nervous systems do when we walk, maintaining balance with the aid of multiple inputs from eyes, inner ear, and nerves and muscles). Backup processing capability, perhaps an independent emergency computer, may be needed to bring an AV to a safe stop should the main system fail. Software bugs are especially insidious. Because full testing of any very large block of code is impossible, there being far too many end-to-end execution paths to be checked and verified under all possible operating conditions, errors sometimes lurk undetected for years. While perfect reliability is unattainable, engineers know many methods for enhancing reliability, and experience from military systems, space, and aviation will inform AV designs. As long ago as the 1950s, for example, military fighters and bombers incorporated multiple, redundant computer processors for navigation, weapons targeting, and missile guidance. In fly-by-wire aircraft, multiple communications channels “vote” on signals passed from autopilot to control surfaces, a split vote returning command to the flight crew (an idea evidently originated by John von Neumann in the early 1950s), and Airbus, for added assurance, has designed its fly-by-wire aircraft around two independent computer systems, programmed by different groups in different computing languages.^e

Under normal driving conditions, sensor inputs should be able to keep an AV in its lane at safe distances from other traffic and stationary obstacles. Navigation will depend on GPS location data referenced to some sort of established grid. Interstate highways, urban freeways, and major city streets can be accurately mapped with quick updates, for instance through data uploads from AVs that recognize a change such as a new pothole. Less detail will be available for rural roads, and AVs will have to rely more heavily on onboard sensors to stay in their lane and travel to a final destination between mapped waypoints such as crossroads.^f Isolated settings such as open pit mines, in which routings and obstacles can be readily located and controlled, are technically the easiest, and driverless robotic trucks already carry ore and waste up and down such pathways and roads. Intercity trucking on divided, limited access highways brings a substantially higher order of complexity, even if the contingencies are easier to anticipate and address than in the far messier environments of surface streets (where, for instance, tree branches sometimes hide stop signs and traffic signals may stop working).

Driving tasks as summarized above pose greater difficulties than, say, flying a jetliner with an autopilot. Even though aircraft maneuver in a third dimension, highways are crowded while the skies are largely empty, almost entirely so above a few thousand feet. Military and civilian controllers monitor traffic except for light planes, and even they are supposed to file flight plans (although drones bring new issues). Pilots are far better trained than motorists and are expected

to be attentive even when idle and bored while the autopilot does the flying. Even so, accidents happen. The crew of an Air France Airbus lost situational awareness and the plane fell into the South Atlantic in 2009 after the autopilot handed back control when iced-up air speed indicators gave conflicting signals. The pilots did not realize what was happening and put the plane into a stall, unrecognized until recovery was impossible.^g This was an extreme case. What it demonstrates is, first of all, human frailty. Airline pilots are professionals, certainly compared with truck and taxi drivers, much less young people with only a learner's permit. With more than a few motorists confused every time they step into an unfamiliar rental car, some standardization of safety-related AV features may come to be seen as a regulatory necessity.^h

Once AVs reach the new car market, they will share the roads indefinitely with conventional vehicles driven by people variously skilled and attentive, cautious and reckless. US-licensed drivers number around 220 million. Some are teenagers, others nonagenarians. Driver behavior is conditioned by social norms, unwritten and fuzzy, varying even between city and suburbs. AVs will have to make decisions at least as likely to be “correct,” meaning first of all safe, as those of the human drivers around them. That will require a great deal of testing, through which engineers and the firms trying to commercialize AV technologies can draw lessons learned, modify system behavior, and then test some more. Toyota's Gill Pratt believes something like a trillion miles of vehicle testing will be necessary to attain full autonomy.ⁱ Even that might not bring social acceptability. People will probably demand that AV performance exceed human performance by a substantial margin, since public opinion seems to judge machine error more harshly than human.

^a See, e.g., *Strategies to Advance Automated and Connected Vehicles: Briefing Document* (Washington, DC: National Academies Press, 2017). Many nontechnical discussions can be found in popular media. The widely accepted definitions of the engineering society SAE International (formerly the Society of Automotive Engineers) list six levels of driver assistance, from no automated features to full self-driving capability.

^b “From Here to Autonomy,” *Economist*, March 1, 2018, online; Shai Shalev-Shwartz, Shaked Shammah, and Amnon Shashua, “On a Formal Model of Safe and Scalable Self-driving Cars,” *arXiv:1708.06374v5 [cs.RO]* 15 Mar 2018, online.

^c Lidar costs are coming down. See, e.g., Katharine Bourzac, “Startup Baraja Promises Cheap, Reliable Lidar for Self-Driving Cars,” *IEEE Spectrum*, July 12, 2018, online.

^d James E. Tomayko, *Computers Take Flight: A History of NASA's Pioneering Digital Fly-By-Wire Project*, NASA SP-2000-4224 (Washington, DC: National Aeronautics and Space Administration, 2000).

^e Jeffrey Mervis, “Not So Fast,” *Science* Vol. 358, No. 6369 (2017), pp. 1370-1374; Ashley Nunes, Bryan Reimer, and Joseph F. Coughlin, “People Must Retain Control of Autonomous Vehicles,” *Nature*, Vol. 556 (April 12, 2018), pp. 169-171. For a closer look at some of the technical issues, see Holly E. B. Russell, et al., “Motor Learning Affects Car-to-Driver Handover in Automated Vehicles,” *Science Robotics*, Vol. 1 (2016), eaah5682.

^f Teddy Ort, Liam Paull, and Daniela Rus, “Autonomous Vehicle Navigation in Rural Environments without Detailed Prior Maps.” Paper presented at International Conference on Robotics and Automation, IEEE Robotics and Automation Society, Brisbane, May 21-25, 2018.

^g Jens Flottau and Robert Wall, “Painful Lessons: AF447 Crash Investigation Points to Inadequacy of Pilot Training,” *Aviation Week & Space Technology*, June 6, 2011, pp. 36-38.

^h *Automated Vehicles: Comprehensive Plan Could Help DOT Address Challenges*, GAO-18-132 (Washington, DC: Government Accountability Office, November 2017); Bill Canis, *Issues in Autonomous Vehicle Deployment*, R44940 (Washington, DC: Congressional Research Service, May 17, 2018).

ⁱ Bill Visnic, “Tapping the Brakes on Autonomy,” *Automotive Engineering*, February 2018, p. 4. Pratt is CEO of the Toyota Research Institute.

since the reasons lie so deeply buried in step-by-step learning histories extending over millions or billions of iterations. As the storied group of elite scientists known as JASON concludes, without equivocation: “the response of the network to all possible inputs is unknowable.”¹⁰ For such reasons, AIs will make mistakes—or what will be portrayed as mistakes—just as people do.

¹⁰ *Perspectives on Research in Artificial Intelligence and Artificial General Intelligence Relevant to DoD*, JSR-16-Task-003 (McLean, VA: JASON, January 2017), pp. 28-29. Elsewhere (p. 32) the authors point out that “People are creating systems that work, that may become commercial, but which have unknown liabilities downstream.”

Of course, new technologies are expected to have downsides. Neural networks can be spoofed, for example with “stickers on a stop sign, fooling a common type of image recognition AI into thinking it was a 45-mile-perhour speed limit sign.”¹¹ Authoritarian governments already use facial recognition AIs for tracking individuals in order to search out and suppress behavior that officials frown upon, and these same authorities may be indifferent to mistakes that would ignite mass social protests in a country like the United States.

In workplace applications, AIs will destroy some jobs, alter many others, and at the same time create new tasks and jobs. If unconstrained by political and institutional pressures, inequalities in wages and wealth will probably increase because many of the new jobs seem likely to be filled by people with relatively scarce, specialized and often tacit skills—people who can do tasks that computers (and robots) cannot. In the United States, the 1960 census included for the first time an occupational category for computer specialists; estimated employment totaled 13,000 people.¹² The Bureau of Labor Statistics now collects data on over a dozen computer occupations. Not all are well paid; they include support positions such as staffing help desks. But in total, this slice of the labor market already employs over four million people.¹³ Recent media accounts have highlighted a worldwide race for AI “talent” with companies chasing people they hope might seed billion-dollar lines of business.¹⁴ Million-dollar salary offers are said to have arrived for working-level engineers (not just executives) with the right sort of experience.¹⁵ Meanwhile, some mid-career computer scientists and engineers find themselves out of work alongside those displaced from factory jobs.

No matter the talent that floods into AI, finally, fears that these systems might take over the world seem overblown, or at least premature. Any encroachment on human agency would require, first, general purpose rather than domain-specific AI, and, second, agency on the part

¹¹ Matthew Hutson, “Hackers Easily Fool Artificial Intelligence,” *Science*, Vol. 361, No. 6399, p. 215.

¹² *Statistical Abstract of the United States 1973* (Washington, DC: Department of Commerce, Bureau of the Census, July 1973), Table 375, pp. 235 and 238.

¹³ 4.1 million according to May 2017 figures from the online BLS Occupational Employment Statistics Query System.

¹⁴ Typical examples: “Battle of the Brains: Google Leads in the Race to Dominate Artificial Intelligence,” *Economist*, December 7, 2017, online; Jeremy Kahn, “Just How Shallow is the Artificial Intelligence Talent Pool?” *Bloomberg News*, February 7, 2018, online; Tekla S. Perry, “Feeding Frenzy for AI Engineers Gets More Intense,” *IEEE Spectrum*, March 1, 2018, online.

¹⁵ David Cyranoski, “China Enters the Battle for AI Talent,” *Nature Briefing*, January 15, 2018, online.

of the system—the ability to set intentional goals without human intervention. To do so, the AI would have to generate repeated objective functions for itself, since goal-setting without intent would result in chaotic behavior, which AIs will be constructed to prevent, if only to avoid self-destruction. Arguably, all this demands something like consciousness. And since humans themselves do not understand consciousness and goal-setting, or how such attributes of living creatures came about except in quite general evolutionary terms, it is hard to see how such attributes might emerge in inanimate AI systems.

IV. Prediction: Are the Robots Coming, Again?

We see facts and figures marshaled in huge arrays that have failed somehow to include inputs from common sense or from human values.

~ Murray Gell-Mann, 1971¹

Many Americans look back with something like nostalgia to a time when high school graduates could get a secure, well-paid job in a factory, department store, or the telephone company with opportunities to advance within an internal labor market. That era is gone. These jobs will no more return than the self-reliant craft workers—hatters and shoemakers, cabinetmakers and ironworkers—of the antebellum era, or the mechanics and machinists who maintained the equipment in the factories that displaced these artisans. And it is well to remember that the past also included much grinding agricultural labor, grim textile mills, death-dealing coal mines, and pollution-belching blast furnaces. Artificial intelligence enables digital systems to take on tasks of greater cognitive complexity. According to the doomsayers, these could take over half or more of today's jobs and perhaps eclipse humans in other ways too. Soberer accounts find these depictions farfetched. Still, evolving forms of AI threaten not only workers in lower-paid occupations but promise to take over some of the responsibilities of those in well-paid occupations such as medicine and law. Certainly the nature of jobs will continue to change, probably at an accelerating pace, and the consequences seem likely to bring still more inequality in the absence of sensible policies to reverse the trends of the past five decades.

Think tanks, consulting firms, and academics have issued a seemingly endless stream of reports on AI and what it might bring.² The more helpful of these point to uncertainties and unknowns, which are many and large. The less helpful put forward quantitative predictions lacking validated causal mechanisms. Such mechanisms scarcely exist and technological forecasting remains more art than science. Quantitative methodologies rest on shaky foundations. Too often, equations full of Greek-letter mathematics and reams of statistics have been put

¹ Murray Gell-Mann, "How Scientists Can Really Help," *Physics Today*, Vol. 24, No. 5 (1971), pp. 23-25; quotation from p. 21. Gell-Mann received the Nobel Prize in physics in 1969.

² Useful discussions include: Jacques Bugin, et al., *Artificial Intelligence: The Next Digital Frontier* (Brussels: McKinsey Global Institute, June 2017); and *The Automation Readiness Index: Who Is Ready for the Coming Wave of Automation?* (London: Economist Intelligence Unit, 2018).

forward to hide less than realistic assumptions on which forecasts hinge. As Robert Pindyck argues in another but not irrelevant, context, “it is much too easy to use a model to generate, and thus seemingly validate, the results one wants.”³

There is at least as much to learn by viewing AI as a family of innovations that fit together piecemeal with others through much trial and error in some sort of sociotechnical system, for want of a better label. This messy process, underway currently for autonomous vehicles, is how most new technologies emerge and diffuse. Over the coming years, AI will become more deeply integrated into various forms of work systems. The impacts will differ by industry sector, by work system, and within these by task and occupation. Even were AI a stable technology, which it is not, there would be great heterogeneity. In this first report we have been able to do no more than point to a few examples. In the next stage of the project we intend to explore the workplace and employment impacts of AI in greater depth.

The Perils of Prediction

In 1987 President Ronald Reagan introduced his administration’s program for rapid commercialization of high-temperature superconductivity, which, he said, promised “a quantum leap in energy efficiency that would bring with it a host of benefits, not least among them a reduced dependence on foreign oil, a cleaner environment, and a stronger national economy.”⁴ Scientists continue to generate research results, but significant innovations, marked by commercialization, have yet to appear.

On the other hand, when incremental technological change continues over several decades, prediction becomes straightforward (Box IV-A). At some point yet to be determined, Moore’s Law for integrated circuits will break down. It is this sort of an event that defies anticipation: multiple predictions of the “end” of Moore’s Law have by now gone wrong. The

³ Robert S. Pindyck, “The Use and Misuse of Models for Climate Policy,” *Review of Environmental Economics and Policy*, Vol. 11, No. 1 (2017), pp. 100–114; quotation from p. 107. One of us spent many years on the editorial board of an academic journal in the field, *Technological Forecasting & Social Change*. The journal’s late founder, Hal Linstone, a mathematician turned systems analyst, was always skeptical of the lack of rigor in many common forecasting methods, criticisms shared with others in the first generation of quantification in the social sciences but not so many of those who followed.

⁴ One of us was present when President Reagan, flanked by three members of his cabinet and reading from a teleprompter, spoke of “a revolution of shattered para-didge-ims.” “Remarks at the Federal Conference on Commercial Applications of Superconductivity, July 28, 1987,” www.reagan.utexas.edu/archives/speeches/1987/072887a.htm.

same is true of many, though not all, initial discoveries. High-temperature superconductivity was a surprise: the two physicists who uncovered the phenomenon, in 1986, were following one of those hunches we have previously encountered. More commonly, innovations—the IC or the digital computer, “a new class of object the world had never seen before”—have been anticipated.⁵ This is because they have precursors, sometimes a great many.⁶ The Wright brothers were the first to fly but hardly the first to try. Predictions stumble over the details, those that turn out to matter, the how and the why. Once new technologies reach the marketplace—if they do—customer acceptance adds further uncertainties. Many products designed around ICs, such as pocket calculators, sold well, at least for a time. Others, including Apple’s personal digital

Box IV-A

Moore’s Law: Predictable Innovation

Integrated circuit density, the number of components per chip (e.g., semiconductor gates, the basic building block), have doubled every two years or so since the 1960s. In the middle of that decade, Gordon Moore, who in 1959-1960 helped develop Fairchild Semiconductor’s first IC chips, drew a straight line through four points spaced a year apart (component counts for 1962-1965, on a logarithmic scale) and extended it to 1975.^a The plot came to be known as Moore’s Law. An empirical regularity, not any sort of natural law, it still enables quite accurate predictions, something no one could have foreseen at the time, and its retrospective fidelity established provided microelectronics firms with planning targets. From around 100,000 components per chip in 1980, densities now reach into the tens of billions, while costs per chip have not changed that much.^b

In the years since, the stream of innovations responsible for the regularity in Moore’s Law plots represent quite fundamental advances in basic semiconductor technology, in design methods for large-scale chips, and in their fabrication. Even so, few of these innovations—not even industry-shattering shifts such as the transition from bipolar chips in the 1970s to MOS (metal oxide semiconductor), and then to a particular type of MOS, complementary metal oxide semiconductors, or CMOS—had much visibility outside the industry and its customers, firms that designed their own products around successive generations of ICs.^b Only fine-grained depictions of Moore’s Law reveal the bumps and wiggles that resulted from this multitude of innovations. From the outside, the process looked like one of continuous innovation. It appeared incremental, or nearly so—regular, hence predictable. Seen from within the microelectronics industry, keeping up meant success or failure in fierce international competition among firms concentrated in the United States, Japan, and South Korea. Still, by comparison the innovations to which ICs contributed, whether the internet or smartphones or machine learning (where special-purpose graphics processing chips were so important), have been far more difficult to anticipate because they are system-level innovations to which ICs contributed, in small ways and large.

^a Gordon E. Moore, “Cramming More Components Onto Integrated Circuits,” *Electronics*, Vol. 38, No. 8 (1965), pp. 114-117.

^b Ross Knox Basset, *To the Digital Age: Research Labs, Start-Up Companies, and the Rise of MOS Technology* (Baltimore: Johns Hopkins University Press, 2002).

⁵ The quotation is from p. 164 in Cuthbert C. Hurd, “Early IBM Computers: Edited Testimony,” *Annals of the History of Computing*, Vol. 3, No. 2 (1981), pp. 163-182.

⁶ George Basalla, *The Evolution of Technology* (Cambridge, UK: Cambridge University Press, 1988). Also see, among much else, Joel Mokyr, *The Lever of Riches: Technological Creativity and Economic Progress* (New York: Oxford University Press, 1990).

assistant, the Newton, and one of its early desktop computers, the Lisa, came to be counted as failures, despite advanced technical features.

In fact, the pathways traced by digital computing over many decades have by now confounded a lengthy string of predictions, and even a cursory look backward reveals flawed predictions for many other innovations in electronics.⁷ Color televisions reached stores in 1950; for the next 20-plus years their sales lagged those of black-and-white TVs.⁸ Although AT&T's much-hyped Picturephone never got a foothold in the marketplace, mobile phones, coming along later, sold year after year in volumes far in excess of essentially all forecasts.⁹ Picturephones, it turned out, did nothing that people actually wanted, while sales of color TVs had to wait for widespread color broadcasting. Past predictions for AI as technology have not held up well.¹⁰ Workplace applications bring in added sources of complexity and uncertainty.

AI in the Workplace

Work processes can take many different forms, even for closely similar products. Beyond the difficulties of anticipating technical advances in AI, grappling with workplace applications means considering, first, the extent to which AIs will be able to supplement or replace human capabilities in particular processes, and, second, the ways in which work systems will be restructured and reorganized on larger scales to exploit emerging AI capabilities. There is also a further set of considerations, still more difficult to assess: firms will redesign existing products and develop wholly products to take advantage of AI. Earlier generations of robotics and IT led to the same changes of the same general type, but such precedents mostly suggest what to be

⁷ Paul Ceruzzi, "An Unforeseen Revolution: Computers and Expectations, 1935-1985," in Joseph J. Corn, ed., *Imagining Tomorrow: History, Technology, and the American Future* (Cambridge, MA: MIT Press, 1986), pp. 188-201.

⁸ *The U.S. Consumer Electronics Industry* (Washington, DC: Department of Commerce, September 1975), p. 20.

⁹ Kenneth Lipartito, "Picturephone and the Information Age: The Social Meaning of Failure," *Technology and Culture*, Vol. 44, No. 1 (2003), pp. 50-81; on the expected growth of mobile telephony, see, e.g., "Cutting the Cord," *Economist*, October 7, 1999, online.

¹⁰ Stuart Armstrong and Kaj Sotala, "How We're Predicting AI—or Failing To," in Jan Romport, et al., eds., *Beyond AI: Artificial Dreams* (Pilsen: University of West Bohemia, 2012), pp. 52-75; Stuart Armstrong, Kaj Sotala, and Seán S. Ó hÉigeartaigh, "The Errors, Insights and Lessons of Famous AI Predictions—and What They Mean for the Future," *Journal of Experimental & Theoretical Artificial Intelligence*, Vol. 26, No. 3, (2014), pp. 317-342; Vincent C. Müller and Nick Bostrom, "Future Progress in Artificial Intelligence: A Survey of Expert Opinion," in Vincent C. Müller, ed., *Fundamental Issues of Artificial Intelligence* (Berlin: Springer, 2014), pp. 1/19-19/19; Katja Grace, et al., "When Will AI Exceed Human Performance? Evidence from AI Experts," arXiv:1705.08807v2 [cs. AI] (May 30, 2017).

alert for in looking ahead.

Some work systems are designed around the technology, as for assembly lines and internet retailing.¹¹ In other cases, technology is auxiliary, as it still is for a good deal of office work, whether the province of administrative assistants or high-paid executives. Hospital emergency rooms are filled both with technical systems and people, those who wheel gurneys and those who wield scalpels. AIs will be tailored to fit into work systems and production processes in their various forms, which will in turn evolve to take advantage of what AIs can do better or cheaper than people. Custodial workers still push brooms. Custodial robots may be on the way, but for now most such robots still navigate like Roombas—by bumping into things, gently—rather than AVs that can see where they’re going and plan a route. (Robots and Roombas, depending on their price points, may have proximity sensors, but not in the numbers of AVs.) How much autonomy future generations of custodial robots get, and how much job displacement they may cause, will depend on first costs and maintenance costs relative to prevailing wages for humans willing to do these sorts of jobs.¹²

The future of AI as a family of technologies will depend on finding good solutions to technical problems that may not yet be visible. Even so, forecasts of employment impacts hinge more on whether, when, and how AI capabilities will be cost-effective in supplementing human capabilities or taking over work altogether—across industries and occupations, economy-wide. Even over the short run, expert opinion on such matters will diverge, in small ways and large.¹³

For several decades, as we saw in Section II, economists have endeavored to explain growing dispersion in wages with models based on the idea of skill-biased technological change. Any model-based prediction will be only as good as the model and the input data.

¹¹ On work systems and their dynamics in an earlier period, when IT had penetrated deeply but before AI had any prominence, see Stephen A. Herzenberg, John A. Alic, and Howard Wial, *New Rules for a New Economy: Employment and Opportunity in Postindustrial America* (Ithaca, NY: Cornell University Press, 1998), pp. 37-81.

¹² For media accounts on automation of less structured tasks performed by lower-paid workers, see, e.g., Cade Metz, “In the Future, Warehouse Robots Will Learn on Their Own,” *New York Times*, September 10, 2017; Danny Vena, “Walmart Is Testing an AI-Powered Robot,” *Motley Fool*, January 16, 2018, online.

¹³ On the question of relating workplace technologies to workplace jobs and tasks, see Daron Acemoglu and David Autor, “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, Vol. 4b (Amsterdam: North-Holland, 2011), pp. 1043-1171; Tim Kautz, et al., *Fostering and Measuring Skills: Improving Cognitive and Non-cognitive Skills to Promote Lifetime Success*, OECD Education Working Papers No. 110 (Paris: OECD, 2014); David J. Deming, “The Growing Importance of Social Skills in the Labor Market,” *Quarterly Journal of Economics*, Vol. 132, No. 4 (2017), pp. 1593-1640.

Meteorologists, for instance, predict the paths of hurricanes with input data gathered from sensors that measure, among other things, temperature, humidity, and wind velocity in the atmosphere. These serve as inputs to models that represent physical processes in the atmosphere. Much as for machine learning AIs, the more data the better. Analyses of skill-biased technological change draw upon masses of data on wages and years of schooling. They fall down because schooling is an inadequate measure of skill. Forecasts of the likely impacts of AI on employment suffer from much the same problem, inappropriate data—and in this case not very much of it.

A common starting point has been information compiled by the US Department of Labor’s Bureau of Labor Statistics on the knowledge and skill content of occupations. BLS maintains the government-wide Standard Occupational Classification (SOC) system, which in its current version includes 867 occupations.¹⁴ In principle, the SOC provides a slot fitting the job of each person in a workforce now numbering over 160 million. BLS also maintains an online database called O*NET (for Occupational Information Network). Intended as a resource for the public (rather than a framework for maintaining consistency in the government’s statistics over the years), O*NET adds several hundred job titles and descriptions to the SOC list. Unlike the SOC, a tool of bureaucracy, O*NET aims to provide users—job seekers, educators, employment counselors—up-to-date information on what employers look for when hiring for listed jobs (education, knowledge and skills) and whether BLS expects the occupation to expand or contract over coming years.

Since O*NET is easy to access and includes summary listings of knowledge and skills by occupation—these drawn from sources including interviews and site visits by job analysts—forecasters have used it in attempting to predict which occupations will be susceptible to job displacement by AI and robotics. This calls for many assumptions and approximations. While some occupations, for instance—nursing, bookkeeping, barbering—have well defined boundaries, many do not: an employee classified as, say, a “computer support specialist” by one employer might at another firm be considered an administrative worker or a technician.¹⁵ But

¹⁴ The SOC was revised in 2018. *Standard Occupational Classification Manual: United States, 2018* (Washington, DC: Office of Management and Budget, 2018).

¹⁵ One way to appreciate the artificial nature of occupational definitions is to note that in earlier decades the government found a great many more. At the peak in 1910, when the economy was far smaller and less differentiated, federal agencies collected statistics on some 20,000 occupations. Margo Anderson Conk,

the greater difficulties originate in O*NET's classificatory scheme for skills and knowledge. This aligns loosely with standard distinctions between the procedural and declarative aspects of work—skills (what people do) and knowledge (what they know). While BLS may have seen little alternative to basing O*NET on the procedural/declarative distinction, it stems originally, not from anything so practically-oriented as job analysis, but from the branch of philosophy known as epistemology, commonly pursued at levels of abstractness well removed from the world of work. As we have already seen, work-related tasks and indeed everyday activities of all sorts, such as driving, mix knowledge and skill inextricably. Beyond this, O*NET's knowledge and skill taxonomies are quite general—they are intended, after all, for a lay audience including high school students—and in any case cannot be updated frequently in any detail since there are so many occupations and BLS's resources are limited.¹⁶

In its treatment of treatment of skills and knowledge, O*NET aims for simplicity and accessibility. Skills come in six categories (basic, complex problem-solving, resource management, social, systems, and technical), all but one of these subdivided, for a total of 35. The O*NET set of knowledge categories, while more expansive, has less coherence. Several of the categories are labeled according to discipline, as for instance “chemistry.” Others, such as “clerical,” have their basis in some broad occupational category; these are linked with knowledge requirements by descriptions that lead off with the word itself, the clerical description beginning “Knowledge of administrative and clerical procedures and systems such as word processing, managing files and records,”¹⁷

With this as a starting point, forecasts based on O*Net then typically rely on surveys or panels of experts to assess, for example, whether a particular task or job, based on its knowledge and skill content, is likely to be, say, 30 percent automatable within, say, two decades, or perhaps 80 percent automatable. This in itself is a formidable task, because even if there are not many

“Occupational Classification in the United States Census: 1870-1940,” *Journal of Interdisciplinary History*, Vol. 9, No. 1 (1978), pp. 111-130.

¹⁶ For more on O*NET and its shortcomings for labor market analysis, see Michael J. Handel, “The O*NET Content Model: Strengths and Limitations,” *Journal of Labour Market Research*, Vol. 49 (2016), pp. 157-176. On relating knowledge and skill to work, also see Lex Borghans, Francis Green, and Ken Mayhew, “Skills Measurement and Economic Analysis: An Introduction,” *Oxford Economic Papers*, Vol. 3 (2001), pp. 375-384; and Michael J. Handel, “What Do People Do at Work? A Profile of US Jobs from the Survey of Workplace Skills, Technology, and Management Practices (STAMP),” *Journal of Labour Market Research*, Vol. 49 (2016), pp. 177-197.

¹⁷ www.onetonline.org/. O*NETs knowledge descriptions for academic fields begin the same way, as “Knowledge of the chemical composition, structure, and properties of substances”

skills in O*NET, there are a lot of occupations. Since the experts cannot be expected to agree, some form of weighting scheme must also be imposed. And all this would of course become still more difficult given more fine-grained information on skills and knowledge than that extracted from O*NET. For example, in contrast to O*NET's 35 skills, and at something of an opposite extreme, a recent "big data" analysis has led to a "basemap" of 13,218 skills.¹⁸

Some of the projections based on methodologies as summarized above have received much attention over the past few years.¹⁹ Careful as the original investigators may have been, not all those who have picked up and publicized the findings have included caveats concerning their limitations. These are so many and so severe that it would be wrong to take the results as anything other than illustrative.

If anything, questions regarding what sorts of new products and processes AI might make possible are more daunting still. Early generations of IT led to wholesale restructuring in industries including financial services, airline ticketing and much of the hospitality industry (online reservations, Airbnb). Available IT products and services likewise underlie the slice of today's gig economy in which people receive payments for discrete tasks and jobs facilitated by telecommunications and the web—also called the platform economy because companies like Uber, Airbnb, and TaskRabbit act as "platforms" that bring buyers and sellers together and handle payments. Up to this point, AI has had little impact on platform work, but that is coming. Labor market forecasts have little to say about new ways of doing business such as these. After all, what is happening today in the platform economy remains uncertain. Estimates of the earnings of ride-hail drivers working for Uber and Lyft range widely, and these and other companies have been growing at such rates that data from last year, if collected, already seems ancient. Forecasters can say little other than that growth rates must at some point decline. By

¹⁸ Katy Börner, et al., "Skill Discrepancies Between Research, Education, and Jobs Reveal the Critical Need to Supply Soft Skills for the Data Economy," *Proceedings of the National Academy of Sciences*, Vol. 115, No. 50 (2018), pp. 12630-12637.

¹⁹ The best known such study is probably Carl Benedikt Frey and Michael A. Osborne, "How Susceptible are Jobs to Computerisation?" *Technological Forecasting & Social Change*, Vol. 114 (2017), pp. 254-280, an analysis that circulated for several years as a working paper before final publication. Also see Philipp Brandes and Roger Wattenhofer, "Opening the Frey/Osborne Black Box: Which Tasks of a Job are Susceptible to Computerization?" *eprint arXiv:1604.08823* (April 2016); and Melanie Arntz, Terry Gregory and Ulrich Zierahn, *The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis*, OECD Social, Employment and Migration Working Papers, No. 189 (Paris: OECD, June 16, 2016), both of which include critiques of Frey and Osborne's assumptions and methodology.

then, of course, there will be other companies trying to establish themselves in other lines of business.

Because so many forecasting models are built around questionable assumptions, and because so much of the data fed into these models is also questionable, there is no real likelihood of resolving differences among the forecasts. Indeed, given that it is essentially impossible to predict new forms of business that AI may spawn—anyone thinks they can do so is more likely to become an entrepreneur or an investor than a forecaster—policymakers cannot expect to have much to go on. Imponderables and contingencies, hard to recognize even in the moment, will affect or determine the pace and ultimate acceptance of future applications of AI and robotics, in the workplace and elsewhere. It is at least as useful to ask what sorts of general lessons past episodes of major technological change can offer.

Prediction: AI as Technological Innovation

The evolution of technologies, following a century of empirical and theoretical work by social and natural scientists, has become relatively well understood, certainly compared with processes of social change to which technological innovations contribute. Generalizations drawn from past episodes of major technological change suggest, first, that people tend to overestimate the effects of technological change in the short run and underestimate them in the long run. Once major innovations begin to mature—the steam engine and electrical power over the course of the 19th century, computing in the second half of the 20th century—powerful and unpredicted waves of technology-instigated change may sweep through the economy and society, upsetting longstanding patterns, including patterns of employment, as some see AI bringing.

Many years ago, Simon Kuznets, 1971 Nobel Laureate in economics, explored these sorts of dynamics, observing that “a major technological innovation requires a long period of sustained improvement, and many significant complementary innovations (some of them also major but derivative) before its ramified and significant effects ... are realized.”²⁰ The complementary innovations to which Kuznets referred may be widely dispersed across industry sectors and barely known beyond those directly involved as engineers, scientists, technicians, and managers. Box IV-B summarizes the internet example.

²⁰ Simon Kuznets, “Innovations and Adjustments in Economic Growth,” *Swedish Journal of Economics*, Vol. 74, No. 4 (1972), pp. 431-451.

Box IV-B

The Internet: An Unpredicted Socio-Technical-Economic Explosion

From a handful of big machines linked in the first ever wide-area computer network in the late 1960s, the net expanded to around 5,000 connections in 1985 and nearly 100,000 before the end of that decade. Growth then exploded, reaching 5 million interconnections in 1995 and over 50 million in another three years or so, on the way to uncountable totals.^a Only in the second half of the 1990s did the effects that would come, with their deep interpenetrations of societies and economies around the world, get widespread attention.

The internet metastasized to an indeterminately expansive web of no fixed identity through confluence among multiple independent streams of innovation: in software; fiber-optic telecommunications channels incorporating solid-state lasers; and hardware for digital systems. Key software innovations included packet switching, an idea that dates from the early 1960s, and browsers, which came much later—around 1990—and made the web user-friendly.^b Fiber optics has its own quite lengthy history, much of it within Corning, a firm known for innovation in glasses and glassmaking.^c Low-cost semiconductor diode lasers made fiber-optic cables practical for long-distance communications.^d Government contributed indirectly, by deregulating telecommunications in the 1980s: with AT&T's monopoly at an end, competing carriers seeking to build market share invested in far more fiber-optic trunk capacity than needed for voice communications, making excess capacity fortuitously available for speeding digital data around the burgeoning internet. Integrated circuit chips, following their Moore's Law trajectory, added a final ingredient—the raw digital processing capacity needed to manage huge flows of data and information. Together, these contributing technological streams made possible the internet. Second-best substitutes might perhaps have been found for one or another, but not for all.

^a These estimated figures are from a long-vanished website, www.cyberatlas.com, as reported there in the 1990s.

^b Paul Baran, *On Distributed Communications: V. History, Alternative Approaches, and Comparisons*, RM-3097-PR (Santa Monica, CA: RAND, August 1964); Janet Abbate, "Government, Business, and the Making of the Internet," *Business History Review*, Vol. 75, No. 1 (2001), pp. 147-176.

^c Gino Cattani, "Leveraging In-House R&D Competencies for a New Market: How Corning Pioneered Fibre Optics," *International Journal of Technology Management*, Vol. 44, Nos. 1/2 (2008), pp. 28-52.

^d M.E. Lines, "The Search for Very Low Loss Fiber-Optic Materials," *Science*, Vol. 226, No. 4675 (1984), pp. 663-668.

Systemic impacts emerge gradually, as potentially debilitating technical obstacles—which sometimes come in seemingly endless streams—are overcome or circumvented. There are always such obstacles, as in expert systems development, and technologies often appear to advance in fits and starts. Or they may stall, as for high-temperature superconductivity. In the early stages—AI in the 1960s, machine learning a dozen years ago—public awareness is slight. Past the point of market introduction, ongoing performance gains and cost reductions often spur mounting increases in scale, in something of a virtuous circle. This seems the usual case for innovations that come to be seen as transformative. Microelectronics entered such a stage in the 1970s. Early integrated circuits were costly and not very reliable. Costs came down, reliability improved, and ICs found their way even into low-priced consumer products such as watches and pocket calculators.

Over time, as entire families of interrelated technologies mature, the consequences mix disruption with benefits such as productivity growth (an imperfect indicator, since productivity

measures fail to capture much of what is transformative in innovations). Some of the effects emerge in the ups and downs of business firms, as for example Motorola, Nokia, and BlackBerry in mobile telephones, Macy's, Sears, and Blockbuster in retailing. Companies that shrink or disappear shed employees, leaving them adrift. Other firms thrive. Demand builds for new sorts of skills and know-how. Old occupations die out and statistical agencies create new categories, as for computer specialists.

Social ramifications may reach deeply into the everyday life of literally everyone. As sales of Model T Fords and other low-priced automobiles ballooned, parts suppliers hired new workers by the thousands, as did companies producing steel and tires, developing oil fields, operating refineries, selling gasoline and offering repair services. Taxpayers demanded more paved roads and governments responded. Cheap transportation and better roads enabled families that could afford detached suburban homes to move outward from cities. Others moved inward from rural America, as agricultural innovations, including mechanization (farmers sometimes adapted Model Ts as tractors or barnyard power sources) that forced many farm families into wage work in cities and towns. The trucking industry expanded to supply wholesale distribution and retail outlets scattered through low-density suburbs. Shopping centers, ring roads, urban expressways, and the Interstate Highway System followed. So did injuries and deaths in traffic accidents and air pollution. Who in 1914, when Model T prices dropped below \$500, might have foreseen these outcomes—outcomes Box IV-C places at Level 3, that of complex sociotechnical systems? These are the sorts of questions now being asked about AI.

Today, for instance, some observers tout AVs as a route to reduced accident rates and less congestion. Utopian visions promise city streets friendly to pedestrians and bicycles, freed from horns and heedless drivers, with lower energy consumption through ride-sharing and optimal routings. Dystopian visions suggest hordes of empty AVs circulating cities in search of customers willing to pay an AI-calculated charge set a penny above break-even and parked during off-peak hours in endless rows like bicycles piled up in Beijing. It helps to debate alternative futures, but nobody at this point, any more than customers buying Model Ts in 1914, can know how people, communities, and society at large will adapt to AVs, how patterns of vehicle usage may change, and what the long-run consequences might be. Box IV-D sketches some oft-overlooked considerations likely to affect AV penetration.

Box IV-C

Technological Innovation and Social Change

Technological innovation can be separated without too much violence to analytical rigor from societal innovations which technologies affect, which affect technologies, and in which technologies are embedded.^a Whether the innovation is technological, political, or behavioral, the common feature is this: Whatever is considered new is not just an idea but a change of some sort that has diffused and taken root, that has salience within the economy and society. Twitter, a social innovation enabled by technology, succeeded, but no one could know initially whether users would accept and adapt to 140-character messaging. Twitter spurred other, follow-on innovations, as essentially all successful innovations do. Social scientists have learned a great deal about the diffusion of such innovations—how free public education, for example, leads young people to stay in school longer, learn more and earn more, with consequences including a greater likelihood that they will contribute to further innovation.

At the same time, we understand the processes of invention, innovation, and marketplace introduction, the seeds of Schumpeterian “creative destruction” associated with particular inventors, entrepreneurs, and firms, far better than the cascading effects, cumulating and interacting over decades, that follow, where individual actors and cause-effect relationships tend to be submerged in large-scale social dynamics.

At the simplest level (Level 1—see the table at the end of this box), we can think of technologies as things that work and perform functions, some of them vital and others matters of convenience.^b The microprocessors in our laptops run text editing software and web browsers. An Airbus flies us from New York to London with remarkable safety. Smartphones adjust automatically to London time and make it easy to reconfirm hotel reservations. Level 1 technologies evolve through ongoing interactions between those who develop the technology and those who use it. Demand is difficult to gauge because potential customers cannot easily or accurately envision what the innovation will mean for them personally. People had to try Twitter before accepting it. Organizations had to experiment with digital computers in the 1950s and 1960s to learn how they might contribute to productivity.

Technologies, at the same time, function as parts of more complex networks—Level 2. Microprocessors and the devices designed around them enable us to tie into the Level 2 internet and AVs will at some point communicate with the internet, trading data with one another almost instantaneously. Generally speaking, societies have learned to design, build, and operate Level 2 networks. When failures do occur, the causes can be diagnosed and future failures of a similar sort prevented, in principle if not always in practice. Standard time zones famously came about in response to confusion causing train wrecks. Failures do still occur: trains and jetliners crash, as do electrical grids. Yet organizational competence—keeping refrigerated trains and trucks carrying fresh foods running on time—has always been a key to using technology effectively, and in principle the organizations that manage Level 2 networks have access to the knowledge and skills needed to ensure that Level 1 technological components function reliably. Learning is possible because the causes of failure can be detected, indeed may be obvious, and the system can organize itself around performance improvement.

This is not true at Level 3. Here complexity becomes pervasive and outcomes difficult or impossible to predict.^c As parts of digital systems, IC chips have contributed, inarguably if murkily, to jobless growth and wage inequality in the United States and elsewhere, and also to system-level failure modes, as when an undetected software bug triggered a cascading sequence of events leading to a major power blackout over the northeastern United States and Canada in 2003 (an event that could perhaps be placed somewhere between Levels 2 and 3). Because full testing of any very large block of code is impossible, errors sometimes lurk undetected for years. The 2003 error will not recur, but other bugs no doubt lurk, and they will crop up in AVs too.

Level 3 problems cannot be fully analyzed and resolved on their own terms. Societies do not know how to intervene to achieve predictable outcomes: the variables are too many, their interconnections too uncertain for cause-effect relationships to be derived. This has been the case with technological effects in the labor market, suggesting the need for creative, anticipatory policy thinking, rather than after-the-fact reactions as in the past.

^a This box draws from John Alic and Daniel Sarewitz, “Energy-Climate Innovation: Simple as One, Two, Three,” Working Paper, Consortium for Science, Policy & Outcomes, Arizona State University, Tempe, AZ and Washington, DC, March 2015.

^b For more on this three-level structure, see Braden Allenby and Daniel Sarewitz, *The Techno-Human Condition* (Cambridge, MA: MIT Press, 2011), pp. 31-85.

Technologies, Networks, and Systems
<i>Level 1: Technologies</i>
• Progressive incremental performance improvement over many decades (Moore’s Law). • Increased technical complexity in a largely closed, engineered, managed system with transparent performance metrics (microprocessors; smartphones). • Uncertainty reduced with accumulating technical knowledge and operating experience (air travel web sites).
<i>Level 2: Technological Networks</i>
• Core system reliability coexists with ancillary dysfunction (malware; hackable, spoofable social media; instability-inducing financial trading bots). • Continually increasing system complexity in a partly open, difficult-to-manage sociotechnical system (error-prone, discriminatory facial recognition; duplicitous chatbots; automated disinformation). • Attempted management of uncertainty managed through institutional arrangements (social media privacy regulations; disclosure agreements).
<i>Level 3: Complex Sociotechnical Systems</i>
• System boundaries dissolve, system dynamics vary unpredictably (autonomous and conventional vehicles mix; killer robots). • Uncertainty grows with system complexity (manipulated elections; truth rot).
<i>Source:</i> Adapted from John Alic and Daniel Sarewitz, “Energy-Climate Innovation: Simple as One, Two, Three,” Working Paper, Consortium for Science, Policy & Outcomes, Arizona State University, Tempe, AZ and Washington, DC, March 2015, p. 13.

^c Complexity is a fraught term, suggestive but hard to define. Many years ago, Herbert Simon, a pioneer in AI and much else, had this to say:

Roughly, by a complex system I mean one made up of a large number of parts that interact in a nonsimple way. In such systems, the whole is more than the sum of the parts, not in an ultimate, metaphysical sense, but in the important pragmatic sense that, given the properties of the parts and the laws of their interaction, it is not a trivial matter to infer the properties of the whole.

Although others have expressed similar sentiments, often suggesting modifications or amplifications, Simon’s statement has not been meaningfully superseded. Herbert A. Simon, “The Architecture of Complexity,” *Proceedings of the American Philosophical Society*, Vol. 106, No. 6, 1962, pp. 467-482; quotation on p. 468.

Implications

Two generalizations hold with few exceptions for major technological change. First, many technologies develop slowly. Early performance may disappoint, as it has so often for AI: machine learning was a common term in the 1960s, long before the renewed expectations of recent years. Second, if and when big impacts appear, they tend to surprise in their force and extent, as illustrated by the internet, disrupting entire industries, destroying some jobs and sparking demand for new skills in new occupational categories.

Disruption caused by technological change, and the economic and socioeconomic forces with which technological change interacts, are nothing new. Taken together, they constitute

powerful sets of forces that reshape peoples' lives and all sorts of human organizations, doing so in ways that are hard to discern in the moment, yet over time can be transformative, as illustrated by the struggles of military strategists and heads of state to grapple with nuclear weapons after 1945 or, more mundanely, the struggles of shopping malls to grapple with sagging customer traffic and the rise of Amazon today. Earlier periods have seen massive disruption: the Industrial Revolution hit Great Britain so hard that the Luddites hit back; productivity growth in US agriculture during the first half of the 20th century saw farm employment plummet, impoverishing rural areas that offered few alternative sources of employment; the wholesale replacement of millions of office workers by computers between, roughly, 1960 and the early

Box IV-D

Autonomous Vehicles: Social and Economic Considerations

Nearly 270 million motor vehicles, mostly cars and light trucks bought for personal use, travel US roads and highways.^a With new car prices rising and vehicles more reliable than in the past, people hold onto older cars and trucks longer: the average age of the fleet, now around a dozen years, has been increasing for decades. Advanced driver assistance features, culminating in autonomy, will push new car prices still higher. Out-of-pocket expenditures (i.e., exclusive of R&D) on AV hardware and software for current-generation test and development vehicles reportedly adds in the range of \$100,000 to the cost of the base car or van.^b Incremental costs will come down with production scale and ongoing advances in currently expensive equipment such as lidars. Still, AV capability will probably always carry a price premium, perhaps quite large. Now-common features such as smart cruise controls (capable of automatically maintaining following distance) and lane-keeping showed up first on expensive luxury cars. Automakers count on option packages carrying big mark-ups to boost profit margins, a pattern expected to continue. Assuming it does, high prices will depress sales of AVs destined for personal use, although not sales of commercial vehicles if purchasers expect savings on driver wage bills to outweigh increases in first cost and maintenance.

Given higher sticker prices, the personal vehicle fleet may age further. Whether or not it does, AVs will share the road with older conventional vehicles for decades to come. Once in customers' hands, moreover, AVs will probably experience in-service failures with growing frequency as the years pass; the standard expectation for reliability is a "bathtub" curve, with early, burn-in failures followed by a period of stability and then a final upturn in failure rate as components wear out or otherwise deteriorate and useful life approaches an end. Troubleshooting AVs will, quite predictably, baffle even the best-trained, best-equipped technicians at times, a common occurrence today with older luxury cars packed with electronics (late-model vehicles may incorporate more than 100 microprocessors for functions ranging from power windows to powertrain controls). Repair manuals tend to be confusing, incomplete, and error-ridden. Replacement parts for low-volume models may be unavailable after a few years. Resale values drop because, even if repairs are possible, the bills mount quickly. Meanwhile, older yet simpler and cheaper cars and trucks run on seemingly forever, at least in the absence of road salt and rust.

Prognostications concerning AV futures rarely mention these well-known aspects of the personal vehicle fleet, the dealer, service, and parts industries, and markets for new and used cars. Factors such as these—maintenance and repair most of all—will affect the acceptance of AVs. Poor early experiences could affect the market for years afterwards.

^a The legal category known as light trucks includes pickup trucks, vans, and sport-utility vehicles. Statistics on vehicle registrations, fleet age, and licensed drivers in this box come from online sites of the US Department of Transportation.

^b Paul Lienert, "Self-driving Costs Could Drop 90 percent by 2025, Delphi CEO Says," *Reuters*, December 4, 2017, online.

2000s, by which time the internet was everywhere and employer failures in organizational “reengineering” had largely ended, a result of lessons learned and diffused.

No one could have anticipated the transformative effects of the internet with confidence. The many past failures of AI prediction provide no reason to expect that future predictions of AI impacts will be better. Imponderables, hard to recognize much in advance, will affect or determine the pace and ultimate acceptance of future applications of AI and robotics, including in the workplace. Much as business process reengineering brought disappointment and disillusion before employers learned what to do and what not to do—in support of their interests, not those of workers—AI will probably see many early failures in workplace applications: what comes afterwards is what will matter. As with organizational reengineering, work systems will be redesigned to fit the capabilities of IT rather than people, absent effective policies (which businesses will resist, as they resist nearly all proposals that threaten their freedom of action) and many workers will lose out.

Several European countries have sought to incorporate working-level employees in their reorganizational efforts, with mixed outcomes. The difficulties are hardly surprising, as the more notorious failures to set up and maintain user-friendly web sites should indicate. As paying customers, on the other hand, people have far greater ability to guide and steer innovation. The marketplace rejected the much-hyped Segway, judged, like the Apple Lisa, to be a poor value. When computer prices came down, as they did for later desktop and laptop machines churned out like commodity products, consumers flocked to stores bearing names such as Circuit City and CompUSA, already lost to most memories.

For AI, policy-makers should discount quantitative predictions and be appropriately skeptical of fine-grained qualitative predictions too. In most cases, as with technological innovation narrowly, workplace impacts will continue to take place with starts and stops, successes and failures. Public policies should aim at preparation for an uncertain future, emphasizing workforce adaptability and learning. Even then, many or most Americans may have to adapt to whatever comes along, willy-nilly. Those best able to do so will prosper. They will learn and advance. Those unable to adapt will suffer, like so many displaced factory and office workers have. Policymakers in the United States have seldom taken concrete steps to prepare for labor market disruption. This time they should.

V. The Plight of the American Worker

[D]eaths of despair come from a long-standing process of cumulative disadvantage for those with less than a college degree. The story is rooted in the labor market

~ Anne Case and Angus Deaton, 2017¹

Developing policies aimed at an AI economy that works for all requires fully appreciating the context in which new technologies will be adopted. The United States has experienced a dramatic increase in economic inequality over the past four to five decades. The causes, which include the divergence between hourly compensation and productivity growth noted in Section I, lie in the economy and in the labor market, and in the political dynamics that shape and condition these markets. As the “labor share” of productivity growth declined, wages stagnated for most of the workforce—see Figure 4. Many Americans, depending on occupation, industry, and other factors, earn less today than their counterparts three or four decades ago.


Figure 5 shows the massive shift in wealth from the great bulk of the population to those at the very top that began during the presidency of Ronald Reagan. For decades, from the 1930s into the 1970s, those in the lower 90% of the distribution gained steadily relative to the richest Americans. Then came a dramatic shift, with huge gains for those at the very top—the upper 0.1%. Greater wealth has come principally from growing income, especially capital income. Since about 1980, the annual earnings of those in the bottom half of the income distribution have hardly changed. Annual per capita income for this group—nearly 120 million adult American residents—averaged only about \$16,000 (pretax, in 2014 dollars) then and now, and the group’s share of national income has fallen from about 20% in 1980 to 12%-13% in recent years.² With static incomes at the bottom and rapidly rising incomes at the top, the richest Americans, the top 0.1% of the distribution, now hold as much wealth as the entire bottom 90%; in the 1970s they claimed only about one-third as much.

¹ “Mortality and Morbidity in the 21st Century,” *Brookings Papers on Economic Activity* (Spring 2017), pp. 397-476; quotation from p. 398. Deaton received the Nobel Prize in economics in 2015.

² Thomas Piketty, Emmanuel Saez, and Gabriel Zucman, “Distributional National Accounts: Methods and Estimates for the United States,” *Quarterly Journal of Economics*, Vol. 133, No. 2 (2018), p. 553-609. Transfer payments such as food stamps and the Earned Income Tax Credit raise post-tax income for the bottom 50% but do not change between-group shares of national income meaningfully.

Figure 4


Wage Inequality: Hourly Pay Has Stagnated for Most US Workers



Source: Keystone Research Center based on Current Population Survey data supplied by Economic Policy Institute, Washington, DC.


Figure 5

The Top 0.1% Have Gotten Wealthier at the Expense of Other Americans



Source: Emmanuel Saez and Gabriel Zucman, "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *Quarterly Journal of Economics*, Vol. 131, No. 2 (2016), pp. 519–578.

Wages after the 1970s fell most dramatically for the once-privileged non-college men that lost high-wage jobs in Midwest and Northeast manufacturing plants. Some of these men never reattached fully to the job market. Growing shares of the next two generations of non-college men, sometimes living in city neighborhoods or rural regions offering little but fast food or other poverty-wage service jobs, also became detached from the formal job market. In the mid-1950s, 85% of white male adults had a job. This share is now only about two-thirds (Figure 6). For black men, the share has been consistently lower.³



Downward mobility has been accompanied by an unprecedented decline in life expectancy among non-college whites, women as well as men. Deaths per 100,000 people nearly doubled (from 177 to 335) for 30-34-year-old non-Hispanic whites from 1999 to 2015 and increased for every five-year age cohort from 25 to 64, leapfrogging (by a lot) rates among black Americans (which kept declining).⁴ Rising numbers of “deaths of despair” (due to drugs

³ Restricting the sample to prime age men 25-54 with a high-school degree or less, the employment-to-population ratio is two thirds for blacks (65.9%) and 79.2% for whites. Keystone Research Center analysis based on the Current Population Survey for 2017.

⁴ Case and Deaton, “Mortality and Morbidity.” The quotations later in the paragraph are from pp. 398-399.

including opioids, alcohol, and suicide) more than offset mortality declines due to advances in treatment for cancer and heart disease. Case and Deaton hypothesize that rising mortality reflects a cumulative, intergenerational disadvantage for those with less than a college degree that “involves many aspects of life, including marriage, child rearing, and religion” while being “rooted in the labor market.” The plight of non-college workers has now begun to spread, with the labor market consigning a growing fraction of 4-year college graduates to low-wage jobs.⁵

Other indicators also show the “land of opportunity” struggling to live up to its self-conception. A recent comparison of human capital, defined to include (only) health and education outcomes, found the United States to have regressed compared to other economies dropping from sixth position in 1990 (between the Netherlands and France) to twenty-sixth in 2016 (between Australia and the Czech Republic).⁶ In the US labor market, wages for women continue to lag well behind those for men doing similar jobs. Upward mobility is down for almost everyone: younger generations of Americans are doing poorly compared to those in their parents’ and grandparents’ generations.⁷ The United States now has lower intergenerational mobility than most Western European countries.⁸ Education, long seen as a means to upward mobility in America, increasingly reproduces socioeconomic stratification across generations. With greater inequality, gaps have grown between have and have-not communities in funding for primary and secondary education, which comes in part from local sources, and schools have segregated based on race (again) and income. High-quality early childhood education, now known to be a highly desirable foundation for later learning (Section II, Box II-B), is likewise most available to those who need it least: except for a small number of low-income families, the upper-middle class and

5 *Low-Wage Workers: Poverty and Use of Selected Federal Social Safety Net Programs Persist among Working Families*, GAO-17-677 (Washington, DC: Government Accountability Office, November 2017). GAO found the share of workers with college degrees who earned from \$12 to \$16 per hour rose from 16% in 1995 to 22% in 2016.

6 Stephen S. Lim, et al., “Measuring Human capital: A Systematic Analysis of 195 Countries and Territories, 1990-2016,” *Lancet*, Vol. 392 (2018), pp. 1217-1234. While these sorts of comparisons are necessarily problematic, subject to many sorts of possible errors and uncertainties, the overall results suggest that while the United States, the gains have been small compared to other wealthy economies (Finland, Iceland, and Denmark topped the rankings in both years, the order unchanged).

7 Raj Chetty, David Grusky, Maximilian Hell, Nathaniel Hendren, Robert Manduca, and Jimmy Narang, “The Fading American Dream: Trends in Absolute Income Mobility Since 1940,” *Science*, Vol. 356, No. 6336 (2017), pp. 398-406.

8 Jason DeParle, “Harder for Americans to Rise from Lower Rungs,” *New York Times*, January 4, 2012; Miles Corak, “Economic Mobility,” *Pathways: State of the Union—The Poverty and Inequality Report* (Special Issue 2016), pp. 51-57.

rich are the only ones who can afford more than custodial child care. The rising cost of public colleges, finally, has led to a plunge in the share of students who come from low- and moderate-income families. Except for a fortunate few scholarship recipients that fuel the meritocratic self-delusions of the rich, below-average students from affluent families now have as good a chance of graduating from college as high-achieving students from low-income families. The effects of educational policy choices extend far beyond the individuals and their families directly impacted, since economy-wide productivity growth and innovation reflect the contributions of all those in the workforce.

What Explains the Plight of the American Worker?

The trends summarized above stem from reassertion since the 1970s of the longstanding domination of business interests and the wealthy over American public policy and institutions. These groups have only rarely found their sway over policies at local, state, and federal levels checked by offsetting forces. For a time during the first decades of the 20th century, farmers—and farm-labor coalitions—supported progressive policies. That period ended with the rise of agribusiness. In the 1930s New Deal legislation gave fleeting endorsement to organized labor. With the slow decline of unions since the 1950s, there has been no truly broad-based group able to stand against corporate interests and the wealthy—something that conservatives understand, hence their determination to strip labor unions even more fully of the ability to influence political choices.

Models of skill-biased technological change fail to explain increasing wage dispersion, for reasons already explored (Section II). The feebleness of efforts by economists to explain the striking increases in inequality over recent decades points to institutional explanations (Box V-A). These are rooted in disparities in political power between ordinary Americans and those who enjoy wealth and what it buys. Put simply and perhaps a bit crudely, money is a megaphone—and a major reason why “mass-based interest groups and average citizens have little or no independent influence” on policy as compared with “economic elites and organized groups representing business interests.”⁹

9 Martin Gilens and Benjamin I. Page, “Testing Theories of American Politics: Elites, Interest Groups, and Average Citizens,” *Perspectives on Politics*, Vol. 12, No. 3 (2014), pp. 564-581; quotations from p. 565. Their findings rest on much statistical data and analysis.

Box V-A Institutions and Inequality

If skill-biased technological change explains a small part of rising inequality, what might provide a more complete explanation? Levy and Temin point to the importance of institutions—interlocking sets of public policies and norms that condition negotiations between employers and labor unions.^a In the system built in the 1930s and 1940s, unions anchored in manufacturing used decentralized collective bargaining to lift real wages at the national rate of productivity growth. By the 1950s a decentralized (by European standards) system of bargaining emerged, with agreements at leading manufacturers in leading industries (e.g., General Motors) setting the “pattern” for wage and benefit improvements then emulated by other firms and industries. Supported by periodic federal increases in the minimum wage that maintained its value relative to manufacturing wages overall—and buttressed by Keynesian macroeconomic policies and persistent low unemployment after World War II—the entire wage structure rose roughly at the rate of inflation plus productivity growth. Macroeconomic policy and the safety net, including Social Security and unemployment insurance, both part of the Social Security Act of 1935, boosted mass consumption and sustained economic expansion. Government from the National War Labor Board to President John F. Kennedy weighed in for equity—supporting fair wages for workers to keep production humming as the United States churned out aircraft, ships, and tanks in unprecedented volume during World War II, and in the 1960s opposing outsized increases in executive compensation.

This system was never close to complete. It denied the benefits and safeguards of federal law to large segments of the population, most egregiously blacks in the Jim Crow South: “sponsors fashioned key bills to avoid disturbing the region’s racial civilization by employing two main policy instruments: the exclusion of agricultural and domestic labor, the principal occupational categories of blacks, from legislation, including...the Wagner Act, Social Security, and the Fair Labor Standards Act; and decentralized administration,” the latter in effect licensing discriminatory enforcement by local white officials.^b The result: “a public-private welfare regime ... based on the exclusion of those who most needed economic assistance.... islands of security within the economy, with high waters all around.”^c After World War II, factory migration to the non-union South widened some of the cracks in the system. At the same time, other changes, such as the gradual acceptance of blacks into more skilled occupations, made it more equal. What proved to be a fragile postwar political coalition finally fractured over the Vietnam War and racial unrest in America’s cities. By about 1980 the New Deal system had collapsed completely. The proximate economic causes—inflation fueled by two oil shocks, interest rate hikes that led to an overvalued dollar, soaring imports, and a deep recession—provided cover for a conservative resurgence by ideologues, business interests, and the wealthiest Americans, a group that statistical research shows to be far more conservative than the bulk of the population.^d Inequality rose, much as before the Depression and New Deal.

A small research literature has examined the subsequent contribution of institutional variables to rising US wage and income inequality. Western and Rosenfeld explore both the direct effect of unions on inequality through wage bargaining and the indirect effects through influence on nonunion wages, public policies, and “wage norms.”^e Dividing the private-sector economy into 18 industries and four regions (72 combinations), they find that the decline of organized labor explains a fifth to a third of the growth in inequality over the period 1973-2007, comparable to the growing stratification of wages by education.

Fortin and Lemieux examine the impact of the falling real value of the minimum wage, union decline, and economic deregulation on rising wage inequality.^f They conclude that about a third of the increase in male and female wage inequality in the 1980s can be attributed to these variables. Deunionization significantly increased inequality for men whereas for women the fall in the minimum wage mattered more. Economic deregulation had a smaller impact on inequality overall, despite steep wage declines in newly deregulated industries such as trucking.

The summary above points to two difficulties that plague economic research on institutions and inequality. First, institutional variables—e.g., deregulation and union density in the trucking industry—do not operate independently but have intertwined impacts. Because of this, efforts to disaggregate the effects of institutional factors, while useful and informative, cannot be taken literally. Federal Reserve Chairman Paul Volcker’s decision to raise interest rates starting in 1979 illustrates the interconnections: this increased the value of the dollar, worsening financial problems at major auto assemblers, leading to an abandonment of the three-decade labor-management agreement to raise wages annually at inflation plus productivity growth, triggering “me-too” concession bargaining in other industries, even including supermarkets, which halved entry-level wages in a few years even though these employers face no competition from outside the region.^g That halving in turn drove supermarket entry-level wages low enough that the failure to raise the minimum wage since 1979 has contributed, along with union decline, to a further 35 years of low pay.

Second, empirical estimates of economy-wide impacts of institutions on inequality suffer from a lack of fine-grained data. The Western and Rosenfeld division of the economy into 72 industry-region combinations, for example, is a valiant attempt to drill down to a scale that more fully reveals the influence of labor unions. But these industries and regions are vast relative to the actual scale of many union and collective bargaining impacts. Within inherently local industries, including most services (e.g., supermarkets) and construction, union influence on inequality operates at the city and metropolitan/regional level, with some policy variables—e.g., minimum wages, prevailing wage laws that impact specific industries, other labor standards—determined partly by local and partly by state governments. With sufficiently localized data it is sometimes possible to show that institutions matter a lot. For example, Jeff Waddoups uses a unique data set to estimate that, if the hotel and casino industry in Reno had the same union density as in Las Vegas, the median wage in Reno would be 24% higher in occupations that are “highly unionized” in Las Vegas, with better non-wage benefits too.^h In the absence of such data, however, the impact of institutional variables may wash out, resulting in underestimates of their actual effects on more localized scales—and in turn feeding pessimism about the potential of institutional changes to restore prosperity. Yet what might it mean for inequality if many inherently local industries had union densities similar to the Las Vegas hotel and casino industry?

Levy and Temin conclude that explanations for income disparities, including skill-biased technical change, operate within a “broader institutional story.” They believe that “Only a reorientation of government policy can restore the general prosperity of the postwar boom, can recreate a more equitable distribution of productivity gains where a rising tide lifts all boats... Only time will tell if more economic distress is needed to change policy yet again.”ⁱ

^a Frank Levy and Peter Temin, “Inequality and Institutions in 20th Century America,” MIT Industrial Performance Center Working Paper Series, MIT-IPC-07-002, Cambridge, MA, revised June 27, 2007.

^b Ira Katznelson, Kim Geiger and Daniel Kryder, “Limiting Liberalism: The Southern Veto in Congress, 1933-1950,” *Political Science Quarterly*, Vol. 108, No. 2 (1993), pp. 283-306; quotation from p. 297. The authors note (same page) that 62% of black workers in 11 southern states worked in agriculture or as domestics in 1930.

^c Jennifer Klein, “The Politics of Economic Security: Employee Benefits and the Privatization of New Deal Liberalism,” *Journal of Policy History*, Vol. 16, No. 1 (2004), pp. 34-65; quotations from pp. 34 and 58. Although the Fair Labor Standards Act has been amended a number of times since initial passage, black women are still the most likely to find themselves among the working poor (as defined by the Labor Department). See Chart 2, p. 3, in *A Profile of the Working Poor, 2016*, Report 1074 (Washington, DC: Bureau of Labor Statistics, July 2018).

^d On the views of the topmost strata of wealth-holders, see Benjamin I. Page, Larry M. Bartels, and Jason Seawright, “Democracy and the Policy Preferences of Wealthy Americans,” *Perspectives on Politics*, Vol. 1, No. 1 (2013), pp. 51-73.

^e Bruce Western and Jake Rosenfeld “Unions, Norms, and the Rise in U.S. Wage Inequality,” *American Sociological Review*, Vol. 76, No. 4 (2011), pp. 513-537.

^f Nicole M. Fortin and Thomas Lemieux, “Institutional Changes and Rising inequality: Is There a Linkage?” *Journal of Economic Perspectives*, Vol. 11, No. 2 (1997), pp. 75-96.


^g For a case study of the supermarket industry in Pittsburgh and documentation of the decline in wages from the late 1970s to the mid-1980s, see Stephen Herzenberg, Diana Polson, and Mark Price, “The Right Choice for Giant Eagle and Western Pennsylvania: A Partnership with Workers That Improves People’s Everyday Lives and Well-Being,” Keystone Research Center, Harrisburg, PA, February 2018.

^h C. Jeffrey Waddoups, “Unions and Wages in Nevada’s Hotel-Casino industry,” *Journal of Labor Research*, Vol. 21, No. 2 (2000) pp. 345-361.

ⁱ Levy and Temin, “Inequality and Institutions,” pp. 43-45.

US employment began to shift away from the highly-unionized manufacturing and construction sectors in the 1950s and private-sector union density began to decline (Figure 7). At the same time, public-sector union membership jumped as teachers and other government employees took advantage after first Wisconsin (in 1959) and then other states granted these workers collective bargaining rights. By the 1970s, private-sector union density was falling even

Figure 7
Private and Public Sector Union Density



Source: Economic Policy Institute, Washington, DC, based on Current Population Survey and historical data from the Labor Research Association.

in manufacturing and construction. The seeds of the decline had been planted in the interwar years, well before threats of import competition, when production of textiles and apparel began drifting to low-wage southern states hostile to organized labor. Other manufacturing sectors followed. When the steel industry began to restructure around nonunion mini-mills in the 1960s, for example, many of these firms built their plants in the South or in similarly accommodating border states. So did foreign-owned automakers and parts suppliers that moved to the United States after about 1980.

Other seeds of decay in the New Deal’s institutional arrangements had been embedded in the political compromises of the original legislation. As noted in Box V-A, conservative Republicans and Southern Democrats in Congress blocked some New Deal proposals completely and compromised others. It is not an oversimplification to suggest that deeply rooted attitudes and ideological tropes associated with both groups—veneration of property rights and states’ rights, “freedom” for individuals and business enterprise—continue to shape the regressive policies that make the United States such an outlier among high-income nations.¹⁰ These themes

¹⁰ Not just historians but economists too have found such characteristics to persist over lengthy time periods, influencing patterns of income inequality as well as innovation and industrial development more broadly. Enrico Spolaore and Romain Wacziarg, “How Deep Are the Roots of Economic Development?” *Journal of Economic Literature*, Vol. 51, No. 2 (2013), pp. 325-369; Alberto Alesina and Paolo Giuliano, “Culture and Institutions,” *Journal of Economic Literature*, Vol. 53, No. 4 (2015), pp. 898-944.

have long histories. “States rights” had been a rallying cry in defense of slavery before the Civil War and a bulwark of white supremacy afterward. For conservatives, individual freedom has been forced to coexist with expectations of deference to hierarchical authority, that of local grandees and the privileged. Political alliances not unlike those of the 1930s underlie the more recent push for antiunion policies and “right-to-work” laws.

By the mid-1980s, private union membership had so collapsed that employers, the right-wing movement, and the growing network of free-market think tanks could train their legal, policy, and political sights on public sector unions. By the late 1980s, some observers had concluded that US-style industrial unionism—organizing and bargaining at the level of the worksite or firm, with wages, job rights, and grievance procedures tied to narrow job classifications—was dead. The US legal system, a conservative force generally and the more so with regard to labor rights (Box V-B), took another step towards euthanasia for collective bargaining in the June 2018 5-4 Supreme Court *Janus* decision. *Janus* makes the country’s public sector “right-to-work” on first-amendment, freedom-of-speech grounds. Nonmembers can now free ride with impunity; they cannot be required to pay anything to the union even though it bargains for them, boosting their pay and benefits, and represents them on the job, protecting them from arbitrary and unlawful actions. The immediate goal of *Janus* backers was to reduce the resources available to public sector unions. A longer-term goal is to strip away their members, reducing public union density to the levels of the 1950s. Conservatives may next seek to make the entire private sector right-to-work through the courts, again on free-speech grounds, to complete the evisceration of labor unions and turn the clock back to the pre-New Deal era when it comes to worker rights.

Box V-B
Legal Institutions and the American Worker

The authors of the US Constitution left many of its provisions open to interpretation, giving the third branch of government great discretion as the final authority on the meaning of much legislation and the legality of executive actions. The courts have used this discretion to shape labor relations in quite fundamental ways. Until the 1930s, judges frequently overturned state legislation intended to establish and enforce worker rights and labor standards, leaving in place “judge-made law” (or case law).^a Case law does reflect precedents and precedents can be revised or overturned, but this may take decades. The legal system, then, counts as a form of institution, with unwritten aspects. Judges have their views, sometimes ideologically determined, they may or may not be aware of these, and it is naïve to think of the courts as some sort of neutral forum for parsing precedents and the Constitution.

From the early years of rapid industrialization in the latter part of the 19th century, the nation’s laws and policies and their legal interpretations have favored the interests of employers over those of workers. Employers, with few

exceptions, have been unremittingly hostile to unions and both state and federal courts have generally supported them. In hindsight, many of their rulings astound. Judges nullified state laws requiring wages to be paid in timely fashion and in legal tender rather than in company scrip. They overturned early workplace health and safety laws on the basis that employees had accepted employers' existing practices in taking a job, whether or not they were aware of these and whether or not there was a written contract. Not infrequently, judges derided such laws as "paternalistic" interference with the "natural" workings of the economy. They routinely upheld injunctions against collective actions such as strikes, finding these to be illegal conspiracies and gifted antiunion firms—which is to say nearly all private employers—with something close to total impunity in meeting strikers with violence delivered by public or private police forces.^b

Such rulings have their basis in English common law, the original basis for US jurisprudence. ("Common law countries," generally one-time English colonies, generally rely less heavily on explicit laws and regulations and more heavily on private contracts and judicial interpretations than "civil law countries" such as France and its own former dominions.) For centuries, English law took the view that a worker's labor constituted his or her property and treated employment as a contract between equals, no matter that one might be master and the other servant.^c Perhaps needless to say, it was the masters—the aristocracy, the owners of vast estates, later joined by wealthy merchants and industrialists—who appointed the magistrates and shaped the legal precedents and understandings implanted in colonial America. With independence, the new nation enshrined private property in its Constitution, here going beyond England (which has no written constitution): "In no country in the world is the love of property more active and more anxious than in the United States; nowhere does the majority display less inclination for those principles which threaten to alter, in whatever manner, the laws of property."^d Reverence of private property, taken to include one's labor, reinforced the common law inheritance that had crossed the Atlantic with the colonists. US law continues to endorse "the legal fiction that employer and employee [stand] as equals in the bargaining process" and the courts treat employment contracts, implicit or explicit, as "private arrangement(s) between two individuals—a seller and a buyer of a service—not amenable to legislative intervention."^e A symmetrical principle, the sanctity of private contracts, upholds US doctrine of "employment at will," the right of any employer to fire any employee at any time, no reason needed (with exceptions chiefly for specifically prohibited forms of discrimination, as on grounds of race or gender—relatively recent and easy to evade).^f

With judges finding ample grounds to strike down, as violating freedom of contract, laws passed by state legislatures to shield workers from abusive practices, reforms inched along until the assertion of federal authority that came with the New Deal. Early in the 20th century, for example, industrial accidents killed something like 35,000 Americans annually, and injured another half-million.^g Survivors had hardly any recourse: if sued, employers argued that if workers knew the dangers and didn't like them they should have quit, their lawyers often claiming contributory negligence as well. Even so, some big companies felt at risk, since they could never be certain that a rogue judge and jury might not find for a plaintiff and award hefty damages. With the support of business associations including the National Association of Manufacturers (NAM), state legislatures began to write laws establishing compulsory insurance programs—workman's compensation—and these the courts began to accept.^h

President Franklin D. Roosevelt's efforts to lift the nation out of Depression plunged the federal government into matters of labor law and enforcement earlier left mostly to the states. Businesses were almost unanimously opposed, challenging newly passed federal legislation with strategies honed earlier in state courts. Employers coupled their legal assaults with public relations campaigns that, among other goals, sought to rebrand "private enterprise," with its suggestion of private interests, as "free enterprise."ⁱ As one commentator put it, NAM sought to sell "free enterprise the way Proctor and Gamble sold soap."^j At the same time, business interests and their conservative allies drummed out messages of the sort heard for decades before and since, depicting labor unions as harbingers of socialism and communism, corrupt and un-American. Nonetheless, with passage of the 1935 Wagner Act (National Labor Relations Act, NLRA)—immediately challenged but finally affirmed by the Supreme Court in 1937—union density tripled.^k

World War II brought a simulacrum of labor peace in the name of outproducing the Axis. It did not last. Supported by Republicans in Congress, employers began to claw back lost ground. Even while the fighting continued, Arkansas and Florida passed "right-to-work" laws in 1944, resurrecting an antiunion strategy from early in the century. In so doing, they repudiated President Roosevelt's call earlier that same year for a "second Bill of Rights"; indeed, they inverted his meaning, since Roosevelt had led off his list with "The right to a useful and remunerative job"^l The Republics took over both houses of Congress following the 1946 election and in 1947 passed the Taft-Hartley Act over President Harry S. Truman's veto, a set of amendments that weakened the NLRA significantly.^m By the end of that year, another 16 states had passed right-to-work laws..

Observers of Congress sometimes say it takes 10 years or more to lay groundwork and pass legislation. After this, the courts have their say: Supreme Court Justice Antonin Scalia once portrayed the main job of that body as "just figuring out" what laws mean and went on to add that "It takes about a decade to get all the kinks out of a new piece of legislation."ⁿ We do

not point to these aspects of US governance to suggest resignation but to emphasize that making the US economy work for all Americans will necessarily be a long-term undertaking, one requiring a broad and stable political coalition able to overcome those successful for so long in promoting the narrow self-interest of the tiny minority on the topmost rungs of the distributions of income, wealth, and corporate power.

^a William E. Forbath, *Law and the Shaping of the American Labor Movement* (Cambridge, MA: Harvard University Press, 1991); Kate Andrias, “The Fortification of Inequality: Constitutional Doctrine and the Political Economy,” *Indiana Law Journal*, Vol. 93 (2018), pp. 5-28.

^b See, e.g., Sanford M. Jacoby, “American Exceptionalism Revisited: The Importance of Management,” in Sanford M. Jacoby, ed., *Masters to Managers: Historical and Comparative Perspectives on American Employers* (New York: Columbia University Press, 1991), pp. 173-200.

^c See, e.g., pp. 3 and 160-161 in Alan Fox, *History and Heritage: The Social Origins of the British Industrial Relations System* (London: George Allen & Unwin, 1985).

^d Alexis de Tocqueville, *Democracy in America*, trans. Henry Reeve (State College, PA: Penn State Electronic Classics, 2002), p. 714.

^e Quotations, respectively, from p. 66 in Melvin I. Urofsky, “State Courts and Protective Legislation during the Progressive Era: A Reevaluation,” *Journal of American History*, Vol. 72, No. 1 (1985), pp. 63-91; and p. 1512 in Katherine van Wezel Stone, “The Post-War Paradigm in American Labor Law,” *Yale Law Journal*, Vol. 90, No. 7 (1981), pp. 1509-1580.

^f “The Tennessee Supreme Court articulated the employment at will doctrine in 1884 This doctrine has been, and still is, a basic premise undergirding American labor law.” Thus “The law, by giving total dominance to the employer, endows the employer with the divine right to rule the working lives of its subject employees.” Clyde W. Summers, “Employment at Will in the United States: The Divine Right of Employers,” *University of Pennsylvania Journal of Labor Law and Policy*, Vol. 3, No. 1 (2000), pp. 65-86; quotations from p. 65.

^g James Weinstein, “Big Business and the Origins of Workmen’s Compensation,” *Labor History*, Vol. 8, No. 2 (1967), 156-174.

^h Price V. Fishback and Shawn Everett Kantor, “The Adoption of Workers’ Compensation in the United States, 1900-1930,” *Journal of Law & Economics*, Vol. 41, No. 2 (1998), pp. 305-342.

ⁱ Robert Griffith, “The Selling of America: The Advertising Council and American Politics, 1942-1960,” *Business History Review*, Vol. 57, No. 3 (1983), pp. 388-412.

^j Richard S. Tedlow “The National Association of Manufacturers and Public Relations during the New Deal,” *Business History Review*, Vol. 50, No. 1 (1976), pp. 25-45; quotation from p. 33. Elizabeth A. Fones-Wolf, *Selling Free Enterprise: The Business Assault on Labor and Liberalism, 1945-1960* (Urbana and Chicago: University of Illinois Press, 1994), p. 25, reports that NAM’s public relations budget increased by more than 20 times from 1934 to 1937.

^k Richard B. Freeman, “Spurts in Union Growth: Defining Moments and Social Processes,” in Michael D. Bordo, Claudia Goldin and Eugene N. White, eds., *The Defining Moment: The Great Depression and the American Economy in the Twentieth Century* (Chicago: University of Chicago Press, 1998), pp. 265-296.

^l Franklin D. Roosevelt, “State of the Union Message to Congress, January 11, 1944.” Franklin D. Roosevelt Presidential Library and Museum, National Archives, Hyde Park, NY. www.fdrlibrary.marist.edu/archives/pdfs/state_union.pdf, p. 7. The second item in the president’s list in effect defined “remunerative” as “The right to earn enough to provide adequate food and clothing and recreation.”

^m Sean Farhang and Ira Katznelson, “The Southern Imposition: Congress and Labor in the New Deal and Fair Deal,” *Studies in American Political Development*, Vol. 19 (2005), pp. 1-39.

ⁿ Bob Cohn, “Scalia: Our Political System Is ‘Designed for’ Gridlock,” *Atlantic*, October 6, 2011.

One possible upside: progressives awakened to the possibility of a vanishing labor movement may become more creative and courageous in their support of forms of organizing such as multi-employer and sectoral bargaining, towards which, as noted in Box V-B, the United States has been hostile. The threat posed by robotics, AI and the gig economy (Box V-C) provides a strong new rationale for supporting the growth of industry-wide forms of worker voice, initially at local and state levels, matters that we address in the next and last section of this report.

In the long sweep of time from the first Gilded Age into the new Gilded Age we're now trying to escape, labor was gravely disadvantaged relative to capital. There have been ups and downs, but the period in which worker interests surged relative to those of business and industry was relatively short, roughly the mid-1930s into the 1960s. AI and the voices expressing concern that its benefits will skew to already advantaged elites provide a fresh opportunity to reassess where we're going as a society. It is easy enough to see that AI will create further openings for disadvantaging employees. When computers replaced the old mechanical time clocks, employers began to tweak software, invisibly, to round off hours of work in ways that shortchanged hourly workers.¹¹ AI will bring many more opportunities for exploiting the vulnerable. But AI also promises new tools for righting existing wrongs. Over the next decade or two we can reconsider the policies of the past—an era of self-conscious “anti-social engineering” that resulted in redistribution to the very top. Section VI looks in more detail at the shape future policy approaches might take.

¹¹ Elizabeth C. Tippet, “How Employers Profit from Digital Wage Theft Under the FLSA,” *American Business Law Journal*, Vol. 55, No. 2 (2018), pp. 315-401.

Box V-C

The Gig Economy: 21st Century Casual Work

Classic images of the old gig economy featured musicians, often middle-aged black men photographed amidst swirls of cigarette smoke, perhaps playing for tips and a cut of barroom receipts (and perhaps worried that they'd be cheated at the end of the night by an owner backed up by police officers moonlighting as bouncers). The tag is now applied to ride-hail drivers, freelance website designers, and others platform workers, some of whom pick up spare cash irregularly, supplementing other sources of income, and some off whom live primarily of such work.^a

Casual work has always been common, much of it goes unreported, and while the share is probably rising there is no way to know how fast. Newer forms of gig and platform work coexist with hiring halls (stevedores, construction laborers), seasonal employment in agriculture, tourism, and retailing, substitute teachers in the schools and adjunct faculty in higher education, along with postdoctoral research positions in the sciences. Surveys by Census and the Bureau of Labor Statistics (BLS) cover traditional categories such as “regular” part-time employment and work found through temporary help agencies.^b Beyond these categories, boundaries in the less formal parts of the labor market are fuzzy, and some of the nether regions have hardly been explored. Spotty and inconsistent data lead to further uncertainties. While analysts generally agree that nonstandard employment has been increasing, probably fairly slowly, much uncertainty surrounds the available estimates. Tax records, for example, indicate that many more households receive income from self-employment than do Census/BLS surveys.^c Little is known about the earnings of self-employed gig workers, unlike temporary and contract employees who get paychecks, if erratically. Gig workers generally pay expenses out of pocket—e.g., insurance, fuel, and maintenance for ride-hail and package delivery drivers and independent truckers, along with vehicle depreciation. Not all keep good records and some keep none, so that net earnings cannot be estimated with any accuracy.^d

While some analysts would limit the label of gig worker to those who find jobs online or through smartphone apps, subcontractors and those who depend on other forms of casual employment face many of the same risks and contingencies: no work when they need it most; no benefits (although the Freelancers Union and its over 375,000 members are trying to change that); clients who pay late or only in part or not at all. Just as mobile telephony and the internet accelerated the expansion of the gig economy, so will AI—and it will probably bring new risks and uncertainties. Advances in online bidding, for example, may push down compensation for many. Forthcoming AI applications are likely to pull more tasks once reserved for “regular” workers in formal enterprises into the platform world. Entirely new products, services, and platforms seem bound to arise. AI will also enable new forms of computer crime, perhaps expanding gray areas between plainly illegal and plainly legitimate activities and perhaps attracting gig workers willing to venture into such areas.

Today, anyone who depends on the gig economy for their livelihood must contrive their own social safety net. Yet if some freelance software engineers and biohackers plan ahead, not all do. And how many freelance musicians or independent truckers have retirement saving plans? Conservatives who would further subvert and flense the Affordable Care Act (Obamacare) have no apparent answer to these questions, nor do libertarians who laud the gig economy as a route to economic freedom.

^a Sarah A. Donovan, David H. Bradley, and Jon O. Shimabukuro, *What Does the Gig Economy Mean for Workers?* R44364 (Washington, DC: Congressional Research Service, February 5, 2016) provides a concise introduction.

^b The most recent BLS data (for October 2018) find about 26 million Americans working part time (according to the official definition, 34 hours per week or less), mostly by choice (and about 4.2 million involuntarily); <https://www.bls.gov/news.release/empsit.t08.htm>. BLS reports show employment by temporary help service agencies falling sharply to about 1.75 million in the aftermath of the 2008 recession and rebounding since to something over 3 million. <https://data.bls.gov/timeseries/CES6056132001>. Also see *Contingent and Alternative Employment Arrangements – May 2017*, USDL-18-0942 (Washington, DC: Bureau of Labor Statistics, June 7, 2018).

^c Lawrence F. Katz and Alan B. Krueger, “Understanding Trends in Alternative Work Arrangements in the United States,” Working Paper 25425, Cambridge, MA, National Bureau of Economic Research, January 2019. Annette Bernhardt, “Making Sense of the New Government Data on Contingent Work,” *Noteworthy: the Journal Blog*, June 10, 2018, while pointing to the lack of statistical data on subcontracting calls the phenomenon a more important change in US work patterns than the gig economy itself. See also Annette Bernhardt, Rosemary Batt, Susan Houseman, and Eileen Appelbaum, *Domestic Outsourcing in the United States: A Research Agenda to Assess Trends and Effects on Job Quality* (Washington, DC: Center for Economic and Policy Research, March 2016).

^d Eric Newcomer, “The Hard Math of the Gig Economy,” *Bloomberg Technology*, November 2, 2018, online.

VI. Policy

The times in which we work and live are changing dramatically. The workers of our parents' generation typically had one job, one skill, one career, often with one company that provided health care and a pension.

This changed world can be a time of great opportunity for all Americans to earn a better living, support your family and have a rewarding career. And government must take your side.

~ George W. Bush, acceptance speech at the Republican Party Convention following his nomination for a second term, September 2004 ¹

The economic changes President Bush highlighted march on, yet little has been done. New policies to address the underlying fundamentals, beginning with the imbalance in political and economic power between employers and employees, will be needed to create an AI economy for all. Table 1 presents policies in eight categories with illustrations of measures that could fall in each. Our goal at this stage is not to be comprehensive but to provide enough specifics so that readers can understand each category and evaluate its potential.

The first category highlights the importance of “narrative.” Leaders and spokespersons for business and finance, and their acolytes, have worked for decades to valorize a singular image of a “free market economy.” In the vision put forward, in the absence of the oppressive hand of government, innovation and entrepreneurship generate prosperity for all. There is much that is compelling in this picture. Technological innovation has made the United States the richest nation the world has yet seen, unsurpassed in military power and a magnet for migrants looking for a better life. Entrepreneurial vigor and a light regulatory hand lead to rapid commercialization of new technologies and widespread diffusion. But the picture is also highly selective. Relentlessly rising inequality since the 1970s demonstrates to anyone willing to consider the evidence that “trickle down” does not work. Market fundamentalists have largely written government out of the story except as an obstacle to free enterprise, whereas in fact the public sector has been a powerfully productive force in building foundations for high-technology innovation, through the military especially, as we have already highlighted for AI. The obsessive

¹ “Transcript: President Bush’s Acceptance Speech,” *Washington Post*, September 2, 2004.

focus on a heroic few inventors, entrepreneurs, and tycoons ignores the countless and vital contributions of the great mass of workers. And while US markets have indeed been loosely regulated compared with those in many other countries, interested parties have constructed and managed many of these markets, whether for agricultural commodities, financial services, health care, or pickup trucks. Business enterprises work assiduously to persuade government officials to tilt the scales in their favor with respect to domestic rivals, foreign industries, and of course workers. Abetted by conservatives who question even the basic legitimacy of government, they keep up their attacks on regulation generally—of labor standards, environment protection, health and safety. In recent years too, the rhetoric of individualism and free market libertarianism, sometimes verging on crude Social Darwinism, has furthered atomization and social conflict, weakening the potential for collective action in support of the interests of ordinary Americans and their families.

Even if assaults on the very idea of forward-looking regulation diminish, it will be difficult as AI spreads to advance public policies strong enough to deliver broad societal benefits. A good starting point for advancing a new narrative would be a more grounded paradigm of economics—which we label “economics that works for all”—less hostile than the neoclassical framework to policy intervention. As the collective authors of the International Panel on Social Progress conclude: “There is a general need for the market to be re-embedded in society if social progress is not to be halted or even reversed.”²

The second category, “trust-busting 2.0,” highlights the raw economic power of modern service-producing corporations (Microsoft, Google, Facebook, Amazon, Walmart), many of them innovators in information technology and major investors in AI. Does popular reaction to violations of privacy or the myriad tentacles of influence reaching into communities and private lives everywhere create an opening for an analogue to Teddy Roosevelt’s trust-busting

² *Rethinking Society for the 21st Century, Vol. 2 Summary* (Cambridge, UK: Cambridge University Press, 2018); quotation from p. 6 of the online summary. In this statement the authors were explicitly reacting against Milton Friedman’s *Capitalism and Freedom*, first published by the University of Chicago Press in 1962. Friedman’s rhetoric mirrors that of Friedrich A. von Hayek, especially *The Road to Serfdom* (Chicago, University of Chicago Press, 1945), which in turn drew on the prolix work of Ludwig Von Mises and others who, in arguing against pre-World War II fascism and Soviet-style economic planning, saw in almost anything governments might do a prelude to disaster. Read today in light of the flood of empirically-based studies of how markets actually function, these polemics, arguing against perceived threats long past, seem closer to political philosophizing than economics, the continuing influence of Chicago school economists notwithstanding.

Table 1. Major Policy Categories for Achieving an AI Economy That Works for All

<i>Category</i>	<i>Objectives</i>	<i>Examples</i>
1. Advance the narrative that we can have an economy that works for all.	Push back on free market ideology and technological fatalism; create space for more forceful policy.	• Promulgate “economics that works for all.” • Build regional “just economy” collaboratives • Promote and publicize measures that track progress towards a more egalitarian economy.
2. Promote trust-busting 2.0.	Reassert government’s right, indeed obligation, to regulate business enterprises for the public good; rein in the worst abuses of corporate power over consumers and workers.	• Strengthen antitrust enforcement at the Federal Trade Commission and Department of Justice under existing laws, which give much discretion. • Base enforcement on recent findings by economists of the effects of concentration, including evidence of restrictions on competition in labor markets. • Prohibit predatory pricing that is anticompetitive and depresses wages at smaller employers. • Strengthen universal service rules/net neutrality.
3. Promote stakeholder over shareholder capitalism when implementing AI.	Give communities and workers a stronger voice as AI technologies reshape the economy and society.	• Strengthen and enforce advanced notice of layoff laws. • Require technological impact statements and community benefits when AI has big projected effects. • Expand support for worker ownership and other forms of democratic capitalism.
4. Promote work-relevant skills and continuous learning to build a more resilient workforce and a more innovative economy.	Reduce fear of displacement and its economic, social, and psychological costs; boost worker satisfaction/self-respect and economic performance.	• Provide universal access to robust income maintenance and re-employment supports for displaced workers. • Invest in programs that integrate classroom and work-based learning. • Reinvent grades 11-14 for the AI economy and society. • Increase investment in training of employed workers, especially those without four-year degrees.
5. Rebalance the scales between employees and employers.	Reduce inequality, strengthen the middle class, and restore opportunity, mobility, and democracy.	• Enforce existing labor standards. • Raise and modernize labor standards. • Universalize benefits (health care, retirement) to protect all workers, including contingent workers. • Reduce employer power in contracts with individual employees (e.g., limit non-compete agreements). • Promote area-wide unions.
6. Create enough paid, family-sustaining jobs to go around.	Avoid mass joblessness and labor market exits; maintain dignity of and from work; improve quality of life for workaholics.	• Establish a right to a job and make that right real (e.g., via Civilian Conservation Corps 2.0). • Reduce work time (e.g., via mandatory vacations, a shorter work week, learning sabbaticals).
7. Share AI productivity bonanza through taxes and social insurance (e.g., Universal Basic Income, UBI).	Reduce after-tax inequality; increase intergenerational mobility; establish AI as a model for social management of innovation.	• Make income taxes more progressive. • Establish wealth taxes. • Enact a financial transactions tax. • Experiment with UBI.
8. Promote responsive democracy.	Reverse drift toward oligarchy; build political power to implement policies in other categories.	• Reduce the impact of money in politics. • End gerrymandering. • Make it easier, not harder, to vote.

of the early 1900s? Railroads and manufacturing firms in industries such as steel, oil, and meatpacking had exploded in size since the 1870s and many of their practices alarmed farmers, smaller businesses, workers, and the public at large. Like Roosevelt, modern-day trust busters could accept that giant corporations bring economic benefits through scale and resources while reasserting our right to bring the activities and behavior of these businesses—creatures of the state, after all—into better alignment with society’s interests. Today, expansive service enterprises, now including many non-profit health care providers, exert outsized influence over consumers and the labor market. The combined market share of the two biggest companies in many industries has grown since the early 2000s, often to 50% or more.³ Firms with 10,000 or more employees now account for 28% of US employment, up from 24% in 1995.⁴ Adding in franchises, misclassified independent contractors, low-wage subcontractors, and other gig workers would increase the share of the workforce over which these quasi-monopolies exert overweening influence. Economic research demonstrates that employers with monopsony power (that are, that is, the only organization that hires particular categories of labor in a given region, such as warehouse workers or nurses) systematically drive down wages and benefits.⁵ But imagine if policy and/or workers could flip a switch to change that: the scale of, say, Amazon, could be used to establish a new balance between corporations and rest of society as happened at GM in the 1930s. Recent administrations have underutilized (to varying extent) US antitrust law, which dates from the Sherman Act (1890) and the Clayton and Federal Trade Commission Acts (both 1914). These laws have never been regarded as strong and have hardly been modernized since passage. Still, recent enforcement has been lax and the Justice Department and Federal Trade Commission could deploy them more effectively to support workers and consumers rather than ignoring abuses or waffling on remedies.⁶

³ David Leonhardt, “The Monopolization of America,” *New York Times*, November 25, 2018.

⁴ David Leonhardt, “The Charts That Show How Big Business Is Winning,” *New York Times*, June 17, 2018, based on data from the Census Bureau’s Longitudinal Business Database for 1977-2014.

⁵ *Labor Market Monopsony: Trends, Consequences, and Policy Responses* (Washington, DC: Council of Economic Advisors, October 2016); Suresh Naidu, Eric A. Posner, and Glen E. Weyl, “Antitrust Remedies for Labor Market Power,” *Harvard Law Review*, forthcoming; online at papers.ssrn.com/sol3/papers.cfm?abstract_id=3129221.

⁶ See, e.g., Jonathan B. Baker, Jonathan Sallet, and Fiona Scott Morton, “Unlocking Antitrust Enforcement,” *Yale Law Journal*, Vol. 127 (2018), pp. 1916-1920, introducing a collection of nine papers on antitrust in today’s economy.

Category 3, the promotion of stakeholder capitalism, would seek to strengthen corporations' obligations to consider community and worker interests, especially when it comes to innovations in robotics and AI. Strengthening and enforcing provisions for advanced notice of layoffs and requiring some kind of "technological impact statement" (analogous to environmental impact statements) might be first steps. Proposals for technological impact statements would draw predictable attacks as brakes on innovation to the detriment of economic growth and competitiveness. That does not invalidate the idea. After all, policy interventions of many sorts, such as telecommunications standards and airline and automobile safety regulations, shape technological outcomes today. Another approach would be a general embrace of stakeholder capitalism. The "Accountable Capitalism Act," introduced in the US Senate in August 2018, borrows from Germany the idea of employees electing at least 40% of the directors on company boards. It would also require the largest companies (\$1 billion or more in annual revenue) to obtain a Federal Charter obliging them to consider the interests of all corporate stakeholders, with revocation of the charter possible for repeated and egregious illegal conduct.⁷ Promoting worker ownership and other forms of democratic capitalism could also help, less directly, offset the power of profit-hungry technology behemoths.

Category 4 includes the starting point for many when it comes to helping individuals cope with technological change: worker training and retraining. The United States as a society does a pretty good job of keeping up with changing skill demands, employer complaints over "skills gaps" notwithstanding. That ongoing learning takes place as well as it does is thanks mostly to the initiative of ordinary Americans who manage, despite sometimes long odds, to learn what they need to know on their own, with whatever help they can scrounge up from relatives, friends, and local resources. It is a massive improvisation, in other words. Employers provide little help, having reduced already meager investments in training of non-college workers. This all takes place today not that differently from when Henry Ford was teaching himself first to build autos and then to build an industrial empire, as illustrated recently when Carnegie Mellon dropout George Hotz developed his own self-driving car.⁸ Continuous learning of needed new skills is thanks a bit to government, and especially to the much-maligned public education system for the

⁷ The text of the bill is online at <https://www.congress.gov/115/bills/s3348/BILLS-115s3348is.pdf>.

⁸ Ashlee Vance, "The First Person to Hack the iPhone Built a Self-Driving Car. In His Garage," *Bloomberg Business*, March 3, 2016.

simple reason that it receives an order of magnitude more money than the “workforce” programs on paper most directly responsible for reskilling workers—albeit still not enough support. K-12 schools do a decent job with the resources available to them, which vary widely among local school districts. Most community colleges have been underfunded from their beginnings after World War II. More recently, under the sway of free-market conservatives and libertarians, governments have been cutting support for four-year public colleges and universities too.

Since skill demands change rapidly, and the pace is more likely to accelerate than decelerate, we should find better ways of keeping up, less dependent on forcing people to figure things out on their own. The United States must do better especially for those in the many occupations for which employers do not expect or require a college degree. Nondegree workers, except for the unionized building trades, rarely have credentials that capture what they know and can do, or job-matching institutions that help them find their next job. The challenge is how to do better, not only to open better career opportunities, but because innovation and productivity growth reflect the contributions of *all* workers, for reasons summarized in Box VI-A. We have no panacea for skill enhancement and continuous learning but discuss possible directions in Table 2 below and the text accompanying that table.

This brings us to category 5, rebalancing the scales between employers and workers to address the chief cause of inequality. As Section V made plain, the difficulty is not in imagining institutional and policy remedies but in implementing them. The obstacles are formidable. They include the heritage of individualism. Long ago on the frontier, families did have to stand on their own two feet, so to speak. Now “individual freedom” has become a central strand in conservative ideologies that underpin some of America’s most regressive policies. It is by intent, to divide and rule. In their campaigns against collective bargaining, businesses have always appealed to individualism—as a banner carried by labor advocates in an 1880 parade recognized: “Each for himself is the bosses’ plea/Union for all will make you free.”⁹ Social scientists too understand that “Business was, and it remains ... the vital center of individualism both in creed and practice.”¹⁰ Of course corporate leaders and conservative politicians also expect

⁹ See p. 504 in David Montgomery, “To Study the People: The American Working Class,” *Labor History*, Vol. 21, No. 4 (1980), pp. 485-512.

¹⁰ Bernard Barber, “Some Problems in the Sociology of the Professions,” *Daedalus*, Vol. 92, No. 4 (1963), pp. 669-688; quotation from p. 670.

Box VI-A

Collective Innovation: Fountainhead of Prosperity

Technological innovation drives increases in productivity, per capita income, and living standards.^a The processes are primarily incremental, resulting in large numbers of mainly small improvements. Most of these have little visibility except to those directly involved. Over time, this multitude of incremental innovations yields large cumulative gains. Since so many of the sources are imperceptible, analysts often represent them by experience curves (Moore's Law for integrated circuits is an example).^b Studies beginning with aircraft production in the 1930s trace many of the gains to hands-on production workers.^c

Occasional bursts of radical innovation, the IC or digital computer, understandably attract more attention. Yet these episodes too almost invariably reflect the work of many people. Arguments still persist over who deserves what sort of credit for early developments in computing. Independent projects at Texas Instruments and Fairchild Semiconductor led to the first ICs, months apart. Jet propulsion developed in parallel over the 1930s and early 1940s in Britain and Germany despite tight military secrecy. Collective discovery, invention, and innovation, as in these examples often more or less simultaneous, is the rule, not the exception.

Few radical innovations, moreover, would merit the label absent years of ongoing incremental advances, the necessary preludes to acceptable performance and reasonable costs. When the first jet planes flew during World War II, they could barely stay aloft for 10 minutes before running short of fuel—this at a time when American P-51 Mustangs fitted with piston engines and disposable drop tanks escorted Allied bombers on round trips of nearly 2000 miles. Militaries continued to pump money into jets because, if the early difficulties could be overcome, they promised overwhelming warfighting advantages. The first ICs, similarly, were unreliable and unaffordable except by military and space agencies. The US Department of Defense and National Aeronautics and Space Agency purchased them in volume because vacuum tubes failed even more frequently in missile guidance systems and spacecraft under the harsh conditions of rocket launch. Scale and learning-in-production boosted reliability and brought down costs to the point that commercial applications began to open.

Even now, in this age of “big science,” productivity growth builds on skills and know-how distributed across the workforce from bottom to top. Not much laboratory science would get done without the unsung labor of technicians who make and maintain equipment. While opportunities in engineering occupations for those with limited formal education have diminished since World War II, never has a college degree been a hard-and-fast requirement: around 300,000 of the 1.7 million Americans who work in jobs classed as engineering (i.e., excluding technicians) cannot claim any sort of four-year college degree, and many of those who do have degrees earned them in nontechnical fields.^d Information technology provides further illustrations of informal acquisition of the technical skills on which innovation depends. Only about 10 percent of those in IT occupations in the mid-1970s, a decade or so after universities began to offer computer-related courses and curricula, held a terminal degree in some computer-related field, and even today not even half of those in technical IT jobs have degrees in computer science or computer engineering.^e It could hardly be otherwise, since the technology has changed so rapidly, as indicated by the frequency with which programming languages have come and gone. Today, Python seems to be the most popular language. Who had heard of it a decade ago? Who today remembers Algol?

Worker learning processes, then, remain central for economic growth and wealth creation even though they may be underappreciated. Put simply, it takes a community to innovate. This is close to self-evident in much of advanced manufacturing, for example specialty chemicals and fabrication of jet engine components. In such cases, close product-process interdependencies had led many firms to co-locate R&D and production facilities. In other US industries, such as auto manufacturing, the significance of “blue collar” innovation had to be learned, or relearned, from Japan in the 1980s and 1990s.^f

Similar stories can be told about almost any industry and occupation. In much of the service sector, employees work directly with customers, seeking to understand their needs and desires in order to craft appropriate outputs, whether a sit-down restaurant meal or a long-term retirement investment plan. Some of this is routine, even if involving freighted decisions (balancing expected risks and returns in a retirement plan). In other cases, service production may amount to a kind of real-time discovery process, as when emergency medical technicians assess the condition of accident victims to properly stabilize them for transport.

As a little thought will suggest, nearly the entire range of variations in service delivery crop up in health care. Major innovations sometimes originate in clinical practice. Whether they stem from laboratory research or the clinic, innovations, major and minor alike, must be incorporated into work practices before any substantial number of people can

benefit. In the course of diffusion and learning, performance gains accrue and new discoveries may follow—continuous innovation. Our own field research explored a past case in which teams of nurses charged with implementing a complicated hospital IT system supplied by the vendor in the form of templates, had to customize software, sometimes on a hospital-wide basis and sometimes for particular departments (e.g., intensive care, the laboratory). Ongoing modification through onsite learning followed. This is characteristic of such systems, which never work perfectly to begin with and must also be periodically modified in line with technological and organizational changes, tasks often entrusted to front-line employees—those who know the work.

Continuous innovation depends on both peoples' skills and the work setting.^g Deep engagement by individual employees, whether in a hospital or an auto plant, leads not just to process improvements but product improvements as well; gains in quality and reliability follow from design modifications incorporating the ideas of front-line workers. In the end, there is little in technology that is truly new at a given point in time. It is just that there is much that is extraordinarily hard to accomplish. This is why ongoing learning and innovation—"everyone an innovator"—is so important. Today, too many workers, because their skills are hard to identify, much less measure, and because they lack credentials such as college degrees, get little recognition and few rewards for their contributions. As a result, they may not contribute all that they could.

^a As explained in Dale W. Jorgenson, Mun S. Ho, and Kevin J. Stiroh, "A Retrospective Look at the U.S. Productivity Growth Resurgence," *Journal of Economic Perspectives*, Vol. 22, No. 1 (2008), pp. 3-24, "total factor productivity is defined as the output per unit of both capital and labor inputs and primarily reflects innovations in both products and processes" (p. 8).

^b Béla Nagy, J. Doyne Farmer, Quan M. Bui, and Jessika E. Transcik, "Statistical Basis for Predicting Technological Progress," *PLoS ONE*, Vol. 8, No. 8 (2013), e52669.

^c John M. Dutton, Annie Thomas, and John E. Butler, "The History of Progress Functions as a Managerial Technology," *Business History Review*, Vol. 58, No. 2 (1984), pp. 204-233; Igal Hendel and Yossi Spiegel, "Small Steps for Workers, a Giant Leap for Productivity," *American Economic Journal: Applied Economics*, Vol. 6, No. 1 (2014), pp. 73-90.

^d Table 1.11 Educational Attainment for Workers 25 years and Older by Detailed Occupation, 2015-16, www.bls.gov/emp/ep_table_111.htm; *Science & Engineering Indicators 2018* (Arlington, VA: National Science Board/National Science Foundation, 2018), Table 3-4, p. 3-30.

^e *Selected Characteristics of Persons in Science or Engineering: 1974*, Current Population Reports, Special Series P-23, No. 53 (Washington, DC: Department of Commerce, Bureau of the Census, July 1975), Table B, p. 5; *Science & Engineering Indicators 2018*, Table 3-4, p. 3-30.

^f Jeffrey K. Liker, John E. Ettlie, and John Creighton Campbell, eds., *Engineered in Japan: Japanese Technology-Management Practices* (New York, Oxford University Press, 1995).

^g For an example, coronary bypass surgery, that finds considerable variation among hospitals, see Robert S. Huckman and Gary P. Pisano, "The Firm Specificity of Individual Performance: Evidence from Cardiac Surgery," *Management Science*, Vol. 52, No. 4 (2006), pp. 473-488. As these authors note, the surgeon "is only one member of a larger surgical team that includes anesthesiologists, nurses, perfusionists [who among their tasks run heart-lung machines], and other technicians" (p. 476).

these ostensibly free individuals to obey authority figures such as themselves. A reconstructed progressive narrative could highlight these contradictions in our political economy.

Since the 1940s, states have led the way in solidifying the primacy of employers over employees, most obviously by passing right-to-work laws. Looking forward, state and local governments could become venues for positive innovations. In fact, they already are: examples discussed later include increases in state (and local) minimum wages, earned sick days, and paid family and medical leave.

Category 6, create enough work to go around, includes both the idea of a public employment guarantee and that of shorter work hours. Renewed fears of technological displacement have already led to the reemergence of arguments for a “right to employment” and proposals for implementing such a right.¹¹ Legal scholars identify a constitutional basis in the non-pecuniary value of work, “including respect, dignity, self-realization, and self-respect.”¹² Data reviewed in Section V on declining employment-to-population ratios among men and rising “deaths of despair” strengthens the case. From a practical point of view, public programs to provide paid work for all who want it would promote more equitable sharing of the benefits of AI. There would be fewer jobless and more workers earning a wage. Pay and benefit standards set in public employment programs would place a floor beneath private sector jobs and could be set to rise in step with productivity. Some champions of public employment rather than universal basic income (discussed below) perceive it as a better fit in the United States, given widespread beliefs that people of working age, leaving aside heirs to fortune, should expect to earn their living. Political figures including several possible future presidential candidates have come out in favor of a job guarantee and advocates argue that if properly crafted the costs would be modest.¹³ There are more than enough unfilled national needs for something like a reimagined Civilian Conservation Corps (CCC) to address, including environmental remediation, infrastructure, and adaptation to climate change. With the proliferation of (little-compensated) creative classes in the age of YouTube, Instagram, and other platforms, CCC 2.0 could also be coupled with a modernized Federal Arts Project so that artists don’t actually starve.

The other example policy in category 6 is shorter work time. Conceptually, this provides a means for sharing whatever paid work robots and AI leave in their wake. In heuristic models, Dean Baker shows that productivity growth sustained over a 30-year period at 2% or above can

¹¹ Several US groups are working on job guarantee proposals. See, for example, Mark Paul, William Darity, Jr., and Darrick Hamilton, “The Federal Job Guarantee—A Policy to Achieve Permanent Full Employment,” Center on Budget and Policy Priorities, March 9, 2018; L. Randall Wray et al., “Public Service Employment: A Path to Full Employment,” April 2018; and the website of the Jobs for All coalition (<https://njfac.org/>), the driving force behind proposed federal legislation, HR 1000 (<https://www.congress.gov/bill/115th-congress/house-bill/1000>) and a New York City bill.

¹² R. George Wright, “Towards a Constitutional Right to Employment,” *Seattle University Law Review*, Vol. 38 (2014), pp. 63-89; quotation on p. 88. For discussion of the checkered history of past proposals going back to the 1930s, see William P. Quigley, “The Right to Work and Earn a Living Wage: A Proposed Constitutional Amendment,” *City University of New York Law Review*, Vol. 2, No. 2 (1998), pp. 139-182.

¹³ Harvey estimates the cost at 3% of the wage income of a local area and less than that at the federal level.

create mouth-watering futures, provided inequality does not rise to still higher levels. By the end of the 30 years, the bottom 90% could enjoy incomes nearly twice as high, or much shorter work time with 50% gains in income.¹⁴

Our seventh category includes policies to distribute the benefits of AI-induced productivity growth more broadly—through the tax system or through social programs. More steeply graduated income and wealth taxes would provide resources for training, public job creation, income maintenance, wage insurance, and other social supports. A Universal Basic Income (UBI), the most visible “big idea” to enter policy discussion as a response to AI anxiety, would channel unconditional cash payments (universal rather than means-tested) to individuals (not households). A guaranteed minimum income would give Americans the freedom to reject unattractive or exploitative employment and focus instead on whatever brings them satisfaction—family, the arts or other hobbies, watching cricket or fútbol. Payments for everyone would spare judgment calls about who is deserving or able-bodied and legitimize fulfilling ways of spending time besides working for a paycheck. Because universal a UBI could be expensive: Van Parijs and Vanderborght propose a benefit set equal to one-quarter of gross domestic product (GDP) per capita, about \$14,000 per person per year in the United States.¹⁵ On the other hand, a UBI would supplement or replace other social supports, with offsetting savings.

Our final category is restoring responsive democracy responsive that is to ordinary Americans and reversing the nation’s slither toward oligarchy, a political system responsive mostly to corporations and the wealthy. This might seem a bit distant from debates over how to manage or mitigate AI’s impact. It is not. For example, AI promises new tools for those bent on subverting democracy, as through automated generation of blizzards of online falsehoods, millions or billions per day, each tailored to the recipient’s internet profile. Beyond such concerns, it will take a better functioning democracy to achieve many of the policies we outline.

¹⁴ Dean Baker, “Can Productivity Growth Be Used to Reduce Working Time and Improve the Standard of Living of the 99 Percent? The Future of Work in the 21st Century,” Washington, DC, Economic Analysis Research Network (EARN), 2014. For a summary of Baker’s analysis, see also Stephen Herzenberg and Laura Dresser, “Beyond Anti-Social Engineering—The Future of Work in the States: Policy an Institutional Choices Toward an Economy That Works for All,” Economic Analysis Research Network, Washington, DC, December 2014.

¹⁵ Philippe Van Parijs and Yannick Vanderborght, *Basic Income: A Radical Proposal for a Free Society and a Sane Economy* (Cambridge, MA: Harvard University Press, 2017), p. 11.

Promote Task-Relevant Skills and Worker Resilience

Our earlier discussion of category 4 sketched the need for support structures to help Americans, young and old, navigate a world of rapidly changing skills demands. Table 2 fleshes out prospective policies. The starting point is miserly support for workers changing jobs: the United States spent only 0.27% of GDP on labor market programs in 2016, less than one quarter the (unweighted) average of 29 mostly high-income countries belonging to the Organization for Economic Cooperation and Development.¹⁶ Only Mexico falls below the United States, Denmark at the top spends 12 times what we do, and US spending fell by nearly two-fifths from 2000 to 2016 (in part because some state governments cut unemployment benefits, doubling down on the “you’re on your own” mentality of conservatives). At the same time, worries over AI have renewed calls for “active labor market policies” to support dislocated workers.¹⁷ Such policies contrast with so-called passive measures such as unemployment insurance (UI), intended simply to temporarily cushion joblessness. Active labor market policies, instead, aim to help people identify, prepare for, and move into new jobs and perhaps new occupations.¹⁸ For some, this can mean a better paying and more satisfying career trajectory. Past US retraining programs too often provided little more than rudimentary instruction for locally available jobs, even if these might already be in decline, coupled perhaps with resume-writing and practice interviews. More effective programs might begin with individualized knowledge and skill inventories, going on to build atop existing competencies. Successful retraining will, for most people, take time, since the objective is to build a basis from which workers can continue learning and, ideally, go on to advance on their own. For such reasons, those enrolled in extended training need income support beyond 26 weeks (or less) of UI.

The Information Technology and Innovation Foundation argues that workers displaced by technology, like those displaced by international trade, merit adjustment assistance, calling for the United States to emulate “flexicurity” (a blend of “flexibility” and “security”) as in

¹⁶ data.oecd.org/social/exp/public-spending-on-labour-markets.htm.

¹⁷ *Artificial Intelligence, Automation, and the Economy* (Washington DC: Executive Office the President, December 2016), pp. 33-37.

¹⁸ For an overview of experience with such programs, see David Card, Jochen Kluve, and Andrea Weber, “What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations,” *Journal of the European Economic Association*, Vol. 15, No. 3 (2018), pp. 894-931.

Table 2. Policies to Promote Work-Relevant Skills and Worker Resilience	
<i>Policy</i>	<i>Comments</i>
Provide universal access to robust reemployment supports (starting with long-term training), and more generous income replacement through unemployment insurance (UI+).	The United States spends little on income maintenance, retraining, and reemployment (e.g., wage insurance) for dislocated workers. AI should stimulate debate about the need for more generous and universal supports. Trade adjustment assistance has not been very effective, but much can be learned from four decades of experience.
Invest in programs that integrate classroom and work-based learning for young people. Reinvent grades 11-14 for those not bound for four-year college. Diffuse lessons from education programs that effectively develop skills and resilience.	The federal government and a growing number of states have rediscovered apprenticeship as an entry point to good jobs, and internships and cooperative programs also enjoy wide support. Education and training programs of all types for youth and young adults should build on proven best-practices including career exposure, applied learning, cohort approaches, and cultures of high standards. Examples include YouthBuild and the US military, which prepares young people, most with only a high school education and little more than basic skills, in a remarkable range of competencies. The analytical question is how to design all such programs for broader scale and coverage, greater appeal, and sustainability.
Increase investment in training of employed workers to lower the likelihood and costs of technological displacement.	US employers invest little in training (with exceptions for fast-tracked college graduates), nor does government, reducing opportunities for “anticipatory” skill acquisition to minimize future layoffs due to technology. A national best-practice standard, awards program, or training tax/mandate could increase private-sector investment. Workers’ representatives could be given a voice in the use of training funds and worker-oriented career counselors (“learning representatives”) could help people choose programs aligned with their aptitudes and desires.

Scandinavia and Singapore.¹⁹ These approaches accept the displacement that comes with technological change, seeing new technologies and business innovations as vital spurs to economic dynamism. In exchange, displaced workers get generous support through active labor market policies. While there is no sign yet of similar enthusiasm in the United States, a full-blown flexicurity system would cost less than a meager UBI, which should help build support as AI filters into more workplaces. Existing levels of inequality promise to be a bigger obstacle. Non-college workers who lose good jobs, perhaps in manufacturing or utilities—and laid-off middle managers too, many of them college graduates—often suffer big losses in income. In more egalitarian Scandinavia, by contrast, the next job is like to pay about as much as the old, perhaps more. American workers (and their unions) would seem unlikely to embrace flexicurity on the basis of generous labor-market adjustment assistance alone. Any such policy would have

¹⁹ Robert D. Atkinson, *How to Reform Worker-Training and Adjustment Policies for an Era of Technological Change* (Washington, DC: Information Technology and Innovation Foundation, February 2018).

to be coupled with measures that boost wages and benefits in the jobs available to displaced workers.

The rest of Table 2 explores policies to strengthen worker resilience *before* dislocation—through reforms to education and training. How do we build a workforce able to learn and relearn work-relevant competencies on a continuing basis over careers spanning four or five decades? A century ago, the US public education system seemed a beacon for the rest of the world. No longer, as attested by international comparisons covering practical as well as academic skills.²⁰ Other nations, realizing the importance for productivity growth and innovation of an educated and capable citizenry—well endowed with human capital, in the jargon of economists—followed in our footsteps and then began to pass us by.²¹ The reasons are many. Among them, in the United States “both vocational education with its overly specific training and academic education with its tendency to become too abstract and decontextualized have abandoned the requirements of work.”²² Partly in response, the Obama and now the Trump administration have sought to scale up and broaden apprenticeship.²³ So have state governments. Much of the interest in revitalizing apprenticeship centers on its earning potential. As important, apprenticeship, along with other programs that combine work and classroom such as internships and cooperative education, impart skills and know-how that other forms of education too often shortchange. While “learning by doing” has never been limited to craft skills in construction (plumbers, electricians) and manufacturing (machinists, toolmakers), the similarities between artisanal training and “professional” training go mostly unacknowledged. The general model in both is the same: learners couple academic preparation with experiential skill acquisition,

²⁰ *Skills Matter: Further Results From the Survey of Adult Skills* (Paris: OECD, 2016); *Learning: To Realize Education's Promise* (Washington, DC: World Bank, 2018).

²¹ But see, for one of many cautionary notes on international comparisons, Judith D. Singer and Henry I. Braun, “Testing International Education Assessments,” *Science*, Vol. 360, No. 6384 (2018), pp. 38-40.

²² Frank Achtnhagen and W. Norton Grubb, “Vocational and Occupational Education: Pedagogical Complexity, Institutional Diversity,” in Virginia Richardson, ed., *Handbook of Research on Teaching*, Fourth Edition, (Washington, DC: American Educational Research Association, 2001), pp. 604-639; quotation from p. 625.

²³ Richard V. Reeves and Katherine Guyot, “Trump Gets Something Right: Apprenticeships and Social Mobility,” *Brookings Social Mobility Memos*, Washington, DC, August 28, 2017, online. For a review of research on the return on investment to apprenticeship, see Robert Lerman, “Restoring Opportunity by Expanding Apprenticeship,” in Irwin Kirsch and Henry Braun, eds., *The Dynamics of Opportunity in America* (New York: Springer and Educational Testing Service, 2016). For a profile of apprenticeship and training funds in unionized construction in Pennsylvania, see Stephen Herzenberg, Diana Polson, and Mark Price, “Construction Apprenticeship and Training in Pennsylvania,” Capital Area Labor-Management Council, Harrisburg, PA, June 2018.

working alongside others, peers, colleagues further along in skill development, and teacher-mentors of established competence. Novices learn from experts, whether expert welders or expert surgeons. The experts pass along tacit skills and know-how of the sort largely missing in classroom settings. This is how artists, musicians, and actors learn, aircraft pilots and police officers, and how scientists cap off their postgraduate training. There is ample opportunity to expand apprentice-like learning to occupations that have yet to benefit. AI, moreover, despite earlier disappointments, promises much-needed help; flight simulators, which allow pilots to train without risking passengers' lives, now incorporate artificial reality and similar technologies are finding their way into other occupational training.

The second row of Table 2 also highlights education reforms that—like apprenticeships—more self-consciously seek the sweet spot between academic coursework and task- and job-specific training, while also hinting at the possibility of a more radical transformation of the last two years of high school and first two years of college. Schooling serves multiple functions. Americans from the beginning viewed a literate and informed citizenry as basic for democracy and a bulwark against autocracy. In the present discussion, our concerns begin with education as preparation for a resilient workforce, one in which Americans by the millions are ready and able to learn new skills as technology and the economy evolve. The most successful preparatory programs expose young people in positive ways to the world of work, an essential prerequisite in an age in which many know little about actual workplaces and workplace tasks, their perceptions filtered through popular media and entertainment, perhaps further shaped by stories told by from adults more likely to gripe than offer useful guidance. Guttman Community College in New York City seeks to overcome such obstacles by fostering peer learning among cohorts of students beginning with a summer “bridge program” that leads into a common first-year curriculum. This common core exposes each student to five alternative pathways—business, information technology, human services, urban studies, and “liberal arts.” Only then, when they have a basis for informed choice, do students pick a major.²⁴ YouthBuild USA, an intensive year-long program for at-risk, out-of-school 16-24 year-olds available in many parts of the country

²⁴ Ranked the best community college in New York State by BestColleges.com in 2018 due primarily to its three-year graduation rate, which exceeds the national norm three-fold, Guttman aims to implement best practices nationwide as distilled in Thomas R. Bailey, Shanna Smith Jaggars, and Davis Jenkins, *Redesigning America's Community Colleges: A Clearer Path to Student Success* (Cambridge, MA: Harvard University Press, 2015).

supports participants in earning a high-school diploma or GED, acquiring occupational skills (originally on community construction projects such as building low-income housing, but now in health care and information technology too), and through structured life-skills training and mentoring. YouthBuild, like Guttman Community College, relies on a cohort approach and a culture of positive reinforcement and peer support.²⁵

The success of the US military in building the skills of service women and men on tasks ranging from troubleshooting complicated electronics systems to cyberwarfare to emergency medicine provides a further model. Like apprenticeship, military training integrates classroom learning and applications. Like the cohort approaches at YouthBuild and Guttman it incorporates discipline, team effort, and shared responsibility to implant and reinforce high standards of ethics and job performance.²⁶

The last row of Table 2 addresses policies to reverse underinvestment in training of employed workers.²⁷ As already noted in our discussion of flexicurity, greater access to training, and training programs of greater effectiveness, would help incumbent workers avoid layoff and if layoff should come pave the way to a new job, perhaps a better job. At some US employers, unions negotiate contributions to funds that give employed workers access to training. These funds go well beyond construction apprenticeship; for example, 15 regional labor-management health care training funds now come together into the Healthcare Careers Advancement

²⁵ YouthBuild has produced estimated returns on investment (social benefits divided by program costs) ranging from seven to 22 times. Mark A. Cohen and Alex R. Piquero, "Benefits and Costs of a Targeted Intervention Program for Youthful Offenders: The YouthBuild USA Offender Project," *Journal of Benefit-Cost Analysis*, Vol. 6, No. 3 (2015), pp. 603-627.

²⁶ Jennifer Kavanagh, *Determinants of Productivity for Military Personnel A Review of Findings on the Contribution of Experience, Training, and Aptitude to Military Performance*, TR-193 (Santa Monica, CA: RAND, 2005); Ryan Keltz, Meredith Kleykamp, and David R. Segal, "The Military and the Transition to Adulthood," *The Future of Children*, Vol. 20, No. 1 (2010), pp. 181-207.

²⁷ Much as for the gig economy, lack of statistical data on training hampers analysis. Not since the 1990s has the Bureau of Labor Statistics had funding for surveys of employer-provided training. A 1995 BLS survey found that far fewer workers with a high-school education or less received formal training than those with four or more years (60% compared with 90%; college-educated workers also benefited from more hours of training). Harley Frazis, Maury Gittleman, and Mary Joyce, "Results of the 1995 Survey of Employer-Provided Training," *Monthly Labor Review*, June 1998, pp. 3-13. Census Bureau reports since that time show declines through 2008 from these earlier and already low levels of employer-sponsored training but provide less detail. See *Economic Report of the President 2015* (Washington, DC: White House, February 2015), pp. 145-147. Reports by industry groups offer rosier accounts of employer programs but mean little because of self-selection bias; good-practice employers eager to highlight how much they invest in their workers are more likely to complete voluntary surveys.

program.²⁸ The United Kingdom’s “learning representatives” offer another model. With paid time off to serve in this role, learning representatives serve as career coaches for front-line unionized employees, providing sounding boards and advice.²⁹ As AI enters more workplaces, demand for this sort of guidance will only rise. In the United States, a “right to training” coupled with a mandate that employers invest a small percentage of payroll could substantially expand training and career counseling. Such a policy would have to be crafted to ensure that employers did not steer the benefits disproportionately to white-collar managers and technical professionals.

Rebalance the Scales Between Workers and Employers

Table 3 provides additional detail on policies to offset current US legal and regulatory biases that favor employers and disadvantage workers. A growing number of states have become more proactive in enforcing labor standards, as did the Labor Department in Washington under Presidents George W. Bush and Barack Obama. Rather than waiting for complaints from workers and their representatives when employers pay less than the minimum wage, deny overtime, or disregard health and safety standards—complaints that may never come if workers fear retaliation or believe the effort futile—proactive enforcement taps historical data, which is abundant, and informants embedded in industries and workers’ social networks to identify firms and industries most likely to violate the law.³⁰ Violations cluster—i.e., in sectors such as day care, nursing homes, janitorial and security services, construction, hotels/motels, and restaurants. Employers competing in these markets, some of them highly cost-sensitive—many business services, such as office cleaning, go to the low bidder, the only factor considered—pay no more than they must to find employees willing to do such work. Turnover is high and workers may not know their rights; many employers hire immigrants, in part because they may be afraid to complain even if their papers are in perfect order. Proactive enforcement targets recidivist employers, a far more efficient use of limited resources (the Labor Department has hundreds of inspectors to monitor millions of employers). Partnerships between government, unions, non-profit worker advocacy groups (“worker centers”), and law-abiding employers and

²⁸ <https://hcapinc.org/>.

²⁹ www.unionlearn.org.uk/union-learning-reps-ulrs.

³⁰ David Weil, *Improving Workplace Conditions Through Strategic Enforcement: A Report the Wage and Hour Division* (Boston, MA: Boston University, May 2010); David Weil, *The Fissured Workplace: Why Work Became So Bad for So Many and What Can Be Done to Improve It* (Cambridge, MA: Harvard University Press, 2017).

their associations have also been shown to improve compliance with labor standards, sometimes dramatically.³¹ The Obama Administration had begun to require bigger firms to exert pressure over their contractors (e.g., for services to buildings) and franchisees (e.g., fast food outlets) to abide by labor standards. The Trump administration has been reversing strategic enforcement and related labor market policies.³²

Table 3. Policies to Rebalance the Scales Between Workers and Employers	
<i>Policy</i>	<i>Comments</i>
Enforce labor standards.	States and the Obama administration have demonstrated the effectiveness of proactive (as opposed to complaint-driven) enforcement. Collaboration with non-government groups also boosts compliance.
Raise and modernize labor standards.	A majority of states and many localities have legislated minimum wages above the federal level, indexed to inflation in some cases. A number of states require employers to offer paid sick days, medical leave, and family leave. Laws requiring work schedules with greater predictability, stability, and flexibility have also begun to spread.
Rein in employer abuses in contracts with workers.	Ban use of “non-compete agreements” for low-wage workers. Ban mandatory arbitration and collective action waivers that deny employees access to courts.
Allow workers to join or form a union free from intimidation.	Many employers aggressively fight any and all unionization efforts, exploiting the broad latitude available under existing US law or simply ignoring the law, which is easy to do and nearly riskless. Enforcement should be stiffer, surer, and timelier. Since announcement of union organizing campaigns opens the door to employer intimidation, workers could be allowed to unionize by signing cards rather than voting in an election.
Promote area-wide unions of service industry employees (e.g., in health care, hotels, restaurants, and retail).	Area-wide unions in non-mobile service industries, those that must remain near their customers, can turn poverty wages into middle-class jobs. They could also negotiate industrywide over new technology, advanced notice of layoffs, and retraining and re-employment rights.
Strengthen the collective rights of independent contractors and caregivers who deliver services in the home.	Those who work as independent contractors by choice, or because employers insist on classifying them as such, and those who provide home-based child care, elder care, or care to the disabled, could be given expanded union rights or antitrust exemptions allowing them to set minimum charges for standard services, avoiding ruinous price competition.
Promote industry-wide wage, benefit, and other labor standards.	In Europe and in New York state, mechanisms exist to establish industry-wide wage and benefit standards that also apply to nonunion employers. Such mechanisms could also extend to industry-wide training and reemployment and to provisions for adoption of new technology.

31 Seema N. Patel and Catherine L. Fisk, “California Co-Enforcement Initiatives that Facilitate Worker Organizing.” Paper prepared for Harvard Law School Symposium, Could Experiments at the State and Local Levels Expand Collective Bargaining and Workers’ Collective Action? Cambridge, MA, September 19, 2017; Janice Fine, “New Approaches to Enforcing Labor Standards: How Co-enforcement Partnerships between Government and Civil Society Are Showing the Way Forward,” *University of Chicago Legal Forum*, Vol. 2017, Article 7 (2018), pp. 143-176.

32 Sharon Block, “State and Local Enforcement: Stepping up and Filling in on Workers’ Rights,” *On Labor*, October 25, 2018, online.

States and localities have also led the way in raising and modernizing labor standards, filling the vacuum left by inaction in Washington. Congress, in the 1938 Fair Labor Standards Act, intended the minimum wage to be a “living wage.”³³ It is far from that today. Forty-two localities and 29 states have now set minimum wages higher than the federal floor, and 18 states index to inflation or will soon.³⁴ Seven states provide for paid family and medical leave, and 10 states and 34 localities for paid sick days.³⁵ Three states and seven cities have fair scheduling laws intended to limit employer practices such as “just-in-time scheduling”—calling in workers with little notice or requiring them to remain “on call” and available at the drop of a hat.³⁶ Given that scheduling algorithms have been a standard application in operations research for decades, there seems no reason why more employers could not offer fair scheduling, and AI might even help.

Moving to the third row of Table 3, states should also curb employer abuses in individual contracts with employees, which have been expanding simply because businesses have found they can get away with them. Among the most egregious, non-compete agreements, long permissible under the laws of many states in contracts with well-paid employees privy to trade secrets and proprietary information, have spread to low-wage sectors such as fast foods. These agreements bar a worker from taking a job with another firm in the same industry, sometimes for several years. That may be acceptable for corporate planners or AI engineers. But for 20-year olds serving pizzas? The motive for employers is to hold down wages by preventing job-hopping. Transparent this may be, yet some state legislatures have accepted legal language supplied by business lobbyists with barely a blink.³⁷ More than a few employers compound the offense by presenting newly hired employees with contract forms after they’ve accepted the job and started work—sign or go home.

³³ Amanda Rose Kapur, “A Faulty Federal Standard: A Call for a Federal Minimum Wage that is Actually ‘Fair’ Under the Fair Labor Standards Act,” *University of Miami Business Law Review*, Vol. 25, No. 3 (2017), pp. 137-170.

³⁴ On state minimum wage measures, see David H. Bradley, *State Minimum Wages: An Overview*, R43792 (Washington, DC: Congressional Research Service, February 28, 2018). On local minimum wages, see www.epi.org/minimum-wage-tracker/#/min_wage/.

³⁵ <http://familyvaluesatwork.org/states>.

³⁶ Julia Wolfe, Janelle Jones, and David Cooper, “‘Fair Workweek’ Laws Help More Than 1.8 Million Workers,” Economic Policy Institute, Washington, DC, July 19, 2018.

³⁷ *Non-compete Contracts: Economic Effects and Policy Implications* (Washington, DC: Department of the Treasury, March 2016); Matt Marx, “Reforming Non-Competes to Support Workers,” Hamilton Project Policy Proposal 2018-04, Brookings Institution, Washington, DC, February 2018.

Since the 1980s mandatory arbitration agreements, often including class and collective action waivers, have also spread. They now cover an estimated 60 million Americans, over half the nonunion private-sector workforce.³⁸ Through these agreements, which may be buried in employee handbooks (and obscured by legalese), employers force complaints over matters of statutory law, including wage theft, racial or gender discrimination, sexual harassment, and deficient health and safety practices, into forums they effectively control. With the assent of the courts, businesses have replaced “law” with “nonlaw.”³⁹ Not only do they avoid any but the most minimal consequences in the event a ruling goes against them, confidential proceedings and sealed findings enable employers, and the offenders that remain on their payrolls, to continue past behaviors with little fear of exposure. Nonetheless, in May 2018 the Supreme Court “held that companies are free to include clauses in employment contracts that force workers to use arbitration rather than courts to enforce their rights—a decision with implications in areas ranging from pay disputes to workplace sexual harassment.”⁴⁰ Beyond disadvantaging employees in particular cases, the legal decisions of the past few decades aid business and conservative agendas more generally, since closed-door proceedings and bars on collective action, by hiding employer practices, impede efforts to generate through publicity the political pressures needed for reform.

Many firms also misclassify as independent contractors workers who meet the legal tests for “regular” employees covered by minimum wage laws and other protections and benefits, including overtime pay, UI and workman’s compensation, sick leave, access to retirement plans, and employer contributions to Social Security and Medicare. By slashing their wage bills

³⁸ Alexander J.S. Colvin, *The Growing Use of Mandatory Arbitration: Access to the Courts Is Now Barred for More Than 60 Million American Workers* (Washington, DC: Economic Policy Institute, April 6, 2018). Similar clauses appear in the purchase agreements and service contracts that consumers by the millions check off with hardly a thought.

³⁹ Katherine Van Wezel Stone, “Mandatory Arbitration of Individual Employment Rights: The Yellow Dog Contract of the 1990s,” *Denver University Law Review*, Vol. 73, No. 4 (1996), pp. 1017-1050; quotations from p. 1018. Stone traces the creeping replacement of statutory law by “private law” to choices made for expediency under the production pressures of World War II.

⁴⁰ Jeff John Roberts, “The Supreme Court: Coming to a Cubicle Near You,” *Fortune*, August 2018, pp. 9-10; quotation from p. 10. The article goes on to quote a supporter of mandatory arbitration who called the high court’s ruling “great news for those of us who believe in freedom of contract and who think the remedy for an unsatisfactory job is to quit and get another”—as if this would be an inconsequential undertaking for typical workers.

through misclassification, employers such as Federal Express seek to undercut their rivals (most UPS workers are covered by union contracts). When sued, FedEx argues that the legal standards do not apply and draws out the proceedings for years: a class-action suit originally filed against the firm in 2004 and finally approaching resolution as of early 2019, will, if approved, award misclassified employees payments averaging a bit less than \$200 each.⁴¹

Now we come to what is perhaps the most potent way to rebalance the scales, and, for that reason, among the most difficult politically: strengthening and reinvigorating workers' collective rights. Proposals that would give workers a bit more leverage within the statutory framework for union organizing and collective bargaining established in the 1930s have been commonplace. That framework has decayed, with employers nibbling away ceaselessly and on occasion managing to carve out big chunks. Meanwhile, the nation's economy has moved away from the industrial structure of the time, keyed to manufacturing but now dominated by services. Row four of Table 3 points to the obvious need for penalties that bite when employers fire pro-union workers during organizing campaigns. This has been a favorite management tactic for more than a century. While the Wagner Act bars employers from firing workers for union activity, the penalties are easy to dodge, delay, or, being miniscule, shrug off.⁴² Plenty of lawyers, industrial psychologists, and consulting firms stand ready to help deploy the many other weapons in the "union avoidance" arsenal.⁴³ The National Labor Relations Board, supposedly a neutral arbiter, has been hamstrung by partisan appointments even as complaints have doubled and redoubled.⁴⁴ The volume of complaints would be still higher if filing were not seen so widely as futile gestures. Yet for 40 years it has proved impossible to get the 60 votes needed to end debate in the US Senate and vote on proposals for strengthening union organizing rights.

⁴¹ Dave Stafford, "FedEx Drivers Get Tentative \$13.3M ERISA Settlement," *Indiana Lawyer*, September 7, 2018, online.

⁴² In an estimated one-in-four (26%) union election campaigns in the 2000s workers were illegally fired, up from 16% in the late 1990s. John Schmitt and Ben Zipperer, "Dropping the Ax: Illegal Firings During Union Election Campaigns, 1951-2007," Center for Economic and Policy Research, Washington, DC, March 2009.

⁴³ On the mercenary assistance available to managements bent on keeping out unions or driving out those they have inherited, see, e.g., Michael J. Zickar, "Using Personality Inventories to Identify Thugs and Agitators: Applied Psychology's Contribution to the War against Labor," *Journal of Vocational Behavior*, Vol. 59 (2001), pp. 149-164; and John Logan, "The Union Avoidance Industry in the United States," *British Journal of Industrial Relations*, Vol. 44, No. 4 (2006), pp. 651-675.

⁴⁴ Joan Flynn, "A Quiet Revolution at the Labor Board: The Transformation of the NLRB, 1935-2000," *Ohio State Law Journal*, Vol. 61 (2000), pp. 1361-1455.

A second approach to fortifying workers' collective rights would strengthen a paradigm of unionism that meshes, as much as any one model of unionism can, with the economy of today and the future. Informed by "justice-for-janitors" and other organizing models new and old, a strong case can be made for area-wide organizing and bargaining, what we have elsewhere called "new unions for a new economy."⁴⁵ In the United States, the bulk of service-sector employers are "non-mobile." Unlike producers of goods that can be held in inventory, stored, and shipped over hundreds or thousands of miles, these enterprises cannot realistically flee or threaten to flee to low-wage states or abroad. They have no choice but to locate near their customers. Non-mobile service businesses account for fully half of all US service sector jobs, and a still larger share of lower-wage service jobs.⁴⁶ They employ health care and long-term care workers, early childhood educators, restaurant, hotel, and other retail workers, taxi drivers and now ride-share and delivery drivers. Many of these jobs call for quite tangible skills, even if these go largely unrecognized and unrewarded.⁴⁷ Others who work for non-mobile employers in smaller occupational groups with specialized skills include contingent faculty in higher education and research scientists working for universities and more than 700 federal laboratories.

In the past, the institutional environment in the United States has been seen as uniquely hostile among industrialized economies to multi-employer and sectoral bargaining. Unions (outside construction) can do little to compel groups of employers to negotiate a master agreement and employers can exit multi-employer bargaining at will. Unlike much of Western Europe, Canada, and Mexico, US laws do not facilitate extension of the economic terms of union contracts to nonunion employers.⁴⁸ Yet states and localities have begun to demonstrate

⁴⁵ Stephen A. Herzenberg, John A. Alic, and Howard Wial, "New Unions for a New Economy," *The New Democrat*, March 1, 1998; <http://keystoneresearch.org/sites/default/files/New-Unions-for-New-Economy.pdf>.

⁴⁶ Conservatively, non-mobile enterprises account for 49% of all service jobs. This figure includes retail trade, urban transit, taxis and limousine services, education, health care, leisure and hospitality (including restaurants and fast food), services to buildings and dwellings, landscaping, personal and laundry services, car washes and auto repair, and the postal service. Adding in other inherently local industries—construction, public administration, and wholesale trade (but excluding utilities, banking and other financial services, almost all professional and technical services, and all transportation other than urban, as well as all mining, agriculture, and manufacturing—brings employment in non-mobile sectors to 63% of the US total. Both percentages calculated from 2016 BLS tabulations for employment by detailed industry, online at www.bls.gov/cps/cpsaat18.htm.

⁴⁷ Dorothy Sue Cobble, *Dishing It Out: Waitresses and Their Unions in the Twentieth Century* (Urbana and Chicago: University of Illinois Press, 1991) explores underappreciated skills in an historical study of craft-like waitress unions.

⁴⁸ Stephen A. Herzenberg, John A. Alic, and Howard Wial, *New Rules for a New Economy: Employment and Opportunity in Postindustrial America* (Ithaca, NY: Cornell University Press, 1998), pp. 162-64, briefly

that, with the right legal innovations area-wide sectoral unions can grow quickly. This has already happened in home care and home-based child care over the past two decades. States and localities, not the federal government, regulate the bargaining rights of these workers (the Wagner Act exempted them, along with employees of state and local public-sector employees). Once progressive states established procedures for state-wide (or county-wide) elections through legislation or executive order, hundreds of thousands of caregivers joined unions.⁴⁹

During area-wide bargaining, unions negotiate with multiple employees or employer associations on behalf of all represented workers in the region covered. That region might be a downtown district or city (hotels and casinos in Las Vegas, hospitals in Manhattan and the boroughs, janitorial contractors or security guards in Philadelphia). It could also be a county or an entire state (as with home care workers and unionized home-based child care providers in some parts of country). Regional bargaining could lift up the bottom third of the US job market and, taking a page from the craft union playbook, give more workers access to multi-employer training funds. Raising wages would cut down on turnover, which runs over 100% annually in some low-wage service industries, resulting in a more stable workforce. Employees who built their skills through experience would provide better service to customers, something that sensible business owners and managers would welcome once wages had been removed from the local competitive landscape. More workers might see a career ahead instead of just “a job.”

“Social bargaining”—pulling in the public as an interested party alongside workers and employers, representative of the community—complements legal strategies for multi-employer organizing. As illustrated by the Justice for Janitors campaign, the logic takes advantage of local and regional political forces. Bargaining at one or another site, a particular factory or warehouse or hospital—the common model in the United States—rarely attracts much interest among the general public, in the absence of high drama such a strike. Negotiations between some large group of workers spread across the region, on the other hand, promise greater attention. There

summarizes the lack of support for multi-employer organizing and bargaining in the United States and statutory changes that could make it more supportive. Also see Mark Barenberg, *Widening the Scope of Worker Organizing: Legal Reforms to Facilitate Multi-Employer Organizing, Bargaining, and Striking* (New York, NY: Roosevelt Institute, October 7, 2015); and David Madland, *Wage Boards for American Workers Industry-Level Collective Bargaining for All Workers* (Washington, DC: Center for American Progress, April 2018).

⁴⁹ About 600,000 US home care workers belong to unions, some 30% of total employment in the industry. Leigh Anne Schriever, *The Home Health Care Industry's Organizing Nightmare* (Washington, DC: The Century Foundation, August 18, 2015).

may be employers in the sector in many neighborhoods (e.g., retail outlets, child care providers). Almost anyone in the community will know some of those involved or affected, on either side of the negotiations or both, as family, friend, or acquaintance. Human interest and the possibility of juicy stories exposing employer abuses or “business models” based on systematic exploitation of vulnerable residents will attract the local media. Political actors will be drawn in too, hoping, whether in or out of office, to appeal to voters; even those who don’t like the tune may jump on the bandwagon. Grabbing the attention of the public and their representatives is the first step in building political support at municipal and state levels for reforms.⁵⁰

Economics That Works for All

Most of those in the business community and conservative movement, even if they disagree deeply about other matters, will fight even baby steps to help workers. For this reason alone, new laws with the potential to bring private sector union density back to 25% or higher will need large-scale backing. Voters saw the original Wagner Act as a means to better prospects for people like themselves, and voters will have to see proposals for new laws and revamped institutional structures as steps that will improve their lives and those of others in their communities, not as abstract measures to restore equity in the nation’s economy. Given the constraints of federalism (and, currently, gerrymandering), little will happen in Washington absent public support sufficiently broad and deep to bring along not just progressive politicians but the ambivalent, uncommitted, and cautious. That is a one of the tasks for a new narrative. In the 1930s, the progressive narrative centered on both the manifestly unfair and exploitative treatment of workers and the New Deal’s promise for lifting the nation out of depression and onto a path of wage-led growth. Looking ahead today, any new narrative should also include the economy-wide benefits of pro-worker policies, including contributions to innovation and productivity growth.

There are times when the “pessimism of the intellect” bears down a little too hard. Doomsday may not be quite as clear and present as when Neville Shute wrote the post-apocalyptic novel, *On the Beach*, in which people lived out their final days while waiting for deadly radiation to arrive in Australia. But the fear that the 1940s to the early 1970s will

⁵⁰ For more on social bargaining, see especially Kate Andrias, “The New Labor Law,” *Yale Law Journal*, Vol. 126 (2016), pp. 2-100.

prove to have been an anomaly hangs heavy. Even the editors of the *Economist*, a journal staunchly supportive of free markets since the first issues appeared in the 1840s, now fret over the intersection between concentrated wealth and the political power of the rich, warning of a “vicious cycle of rising inequality.”⁵¹ Daron Acemoglu’s and James Robinson’s book *Why Nations Fail*—identifying the fatal combination of “extractive” economic and political systems that exist primarily to enrich and empower a tiny slice of the population—also cuts uncomfortably close.⁵² After all, the rich not only influence political decisions directly, through the money they steer to politicians, they have avenues for amplifying their voice ranging up to ownership and control of media empires.

The Heinz Endowments has called on business and civic leaders to—well—lead. It has asked corporate partners to acknowledge that “An increasingly unstable, inequitable, anti-democratic society is ultimately a formula for collapse, not profitability.” For corporations as well as for the rest of us, charting a different path “...isn’t a charity—it’s an investment in your own future.”⁵³

For too long, the United States—and through its leading position internationally, the world—has been in the sway of an economic framework that is empirically barren, intellectually indefensible, and, too often, morally bankrupt. That framework was born in reactions by European academics and intellectuals in the 1930s to corporatist fascism in Nazi Germany and, especially, the monstrous excesses of Soviet economic planning. This same framework now serves to rationalize the far from altruistic claims of those who appeal to individual “freedom” and “free” markets while tarring governments as if anything and everything they might do, leaving aside internal policing, border security, fighting external wars, and enforcement of property rights, risks sliding into the worst excesses of the past. The true agenda is further concentration of wealth and corporate power.

⁵¹ “As inequality grows, so does the political influence of the rich: concentrated wealth leads to concentrated power.” *Economist*, July 21, 2018. Keystone Research Center made the same argument in strikingly similar terms, with special attention to Pennsylvania, in Stephen Herzenberg and Jon White, “Democracy in Pennsylvania,” Harrisburg, July 2, 2018; <https://www.keystoneresearch.org/publications/research/democracy-pennsylvania-2018>.

⁵² Daron Acemoglu and James A. Robinson, *Why Nations Fail: The Origins of Power, Prosperity, and Poverty* (New York, NY: Crown, 2012). Henning Hillmann, “Economic Institutions and the State: Insights from Economic History,” *Annual Review of Sociology*, Vol. 39 (2013), pp. 251-273, provides a broad survey of related literature.

⁵³ “Remarks by Grant Oliphant,” President, The Heinz Endowments, at Metro 21: Smart Cities Institute Launch, Pittsburgh, PA, March 2, 2018, www.heinz.org/UserFiles/File/GO%20speech%20Metro21%20Smart%20Cities%203_2_18.pdf, p. 4.

The full weight of the evidence, within and beyond the United States, makes clear that marketplace competition can operate in constructive or destructive ways. We have wide leeway to shape competition constructively and to influence the distribution of its benefits without undermining innovation, which, as economics also tells us, is the driving force in wealth creation. Indeed, simple rules that shape competition in more constructive directions (away from squeezing workers and despoiling the environment) and more progressive taxation (supporting investment in education and learning for the many and in scientific understanding) will, quite predictably, generate wealth through long-run productivity growth.

Economics says nothing meaningful about how wealth is to be shared. Those are political choices. As a society and a country, we have a right to reject the idea that public policy should be guided by “letting the market decide” and “deregulation” and “shrinking government.” These are not ends in themselves. They are slogans promulgated, sometimes overtly and sometimes through stealth, in efforts to reshape society by self-interested individuals and the advocates they buy. We have a right to determine public policies based on the values of ordinary working Americans and on the kind of country, state, and region we, and they, want to create. We need to assert that right in the next 40 years, forthrightly and with confidence.