
GOVERNANCE IN AN **EMERGING** NEW WORLD

EMERGING TECHNOLOGY
AND THE U.S. ECONOMY

Erik Brynjolfsson, Dipayan Ghosh,
James F. Hollifield, Van Ton-Quinlivan



Convened by George P. Shultz
with James Cunningham, David Fedor, and James Timbie

GOVERNANCE IN AN
EMERGING
NEW WORLD

Convened by George P. Shultz
with James Cunningham, David Fedor, and James Timbie

Table of Contents

SPRING SERIES, ISSUE 619

Introduction..... **5**

How Will Machine Learning Transform the Labor Market?
Erik Brynjolfsson, Daniel Rock, and Prasanna Tambe..... **8**

The Commercialization of Decision-Making: Towards a Regulatory Framework to Address Machine Bias over the Internet
Dipayan Ghosh **19**

The Migration Challenge
James F. Hollifield **36**

Pathways to Economic Opportunity in the 21st Century: A Case Study on How the California Community Colleges Modernized to Deliver on Its Workforce Mission
Van Ton-Quinlivan **56**

Observations from the Roundtable..... **71**



A Letter from the Conveners

Sharp changes are afoot throughout the globe. Demographics are shifting, technology is advancing at unprecedented rates, and these changes are being felt everywhere.

How should we develop strategies to deal with this emerging new world? We can begin by understanding it.

First, there is the changing composition of the world population, which will have a profound impact on societies. Developed countries are experiencing falling fertility rates and increasing life expectancy. As working-age populations shrink and pensions and care costs for the elderly rise, it becomes harder for governments to afford other productive investments.

At the same time, high fertility rates in Africa and South Asia are causing both working-age and total populations to grow, but that growth outpaces economic performance. And alongside a changing climate, these parts of the world already face growing impacts from natural disasters, human and agricultural diseases, and other resource constraints.

Taken together, we are seeing a global movement of peoples matching the transformative movement of goods and of capital in recent decades—and encouraging a populist turn in world politics.

Second is automation and artificial intelligence. In the last century, machines performed as instructed, and that “third industrial revolution” completely changed patterns of work, notably in manufacturing. But machines can now be designed to learn from experience, by trial and error. Technology will improve productivity, but workplace disruption will accelerate—and will be felt not only by call center responders and truck drivers but also by accountants, by radiologists and lawyers, even by computer programmers.

All history displays this process of change. What is different today is the speed of change. In the early 20th century, American farm workers fell from half of the population to less than five percent alongside the mechanization of agriculture. Our K-12 education system helped to navigate this disruption by making sure the next generation could grow up capable of leaving the farm and becoming productive urban workers. With the speed of artificial intelligence, it's not just the children of displaced workers but the workers themselves who will need a fresh start.

Underlying the urgency of this task is the reality that there are now 7.6 million unfilled jobs in America. Filling them and transitioning workers displaced by advancing technology to new jobs will test both education (particularly K-12, where the United States continues to fall behind) and the flexibility of workers to pursue new occupations. Clearly, community colleges and similarly nimble institutions can help.

The third trend is fundamental change in the technological means of production, which allows goods to be produced near where they will be used and may unsettle the international order. More sophisticated use of robotics alongside human colleagues, plus additive manufacturing and unexpected changes in the distribution of energy supplies, have implications for our security and our economy as well as those of many other trade-oriented nations, which may face a new and unexpected form of deglobalization.

This ability to produce customized goods cheaply and in smaller quantities may, for example, lead to a gradual loss of cost-of-labor advantages. Today, 68 percent of Bangladeshi women work in sewing, and 4.5 million Vietnamese work in clothing production. Localized advanced manufacturing could block this traditional route to industrialization and economic development. Robots have been around for years, but robotics on a grand scale is just getting started: China today is the world's biggest buyer of robots but has only 68 per 10,000 workers; South Korea has 631.

These advances also diffuse military power. Ubiquitous sensors, inexpensive and autonomous drones, nanoexplosives, and cheaper access to space through microsatellites all empower smaller states and even individuals, closing the gap between incumbent powers like the United States and prospective challengers and giving potentially disruptive capabilities to non-state and terrorist actors. The proliferation of low-cost, high-performance weaponry enabled by advances in navigation and additive manufacturing diminishes the once-paramount powers of conventional military assets like aircraft carriers and fighter jets. This is a new global challenge, and it threatens to undermine U.S. global military dominance unless we can harness the new technologies to serve our own purposes. At the same time, the proliferation of nuclear weapons poses a serious global threat.

Finally, the information and communications revolution is making governance more difficult everywhere. An analogue is the introduction of the printing press: as the price of that technology declined by 99 percent, the volume grew exponentially. But that process took ten times longer in the 15th, 16th, and 17th Centuries than we see today. Information is everywhere—some of it accurate, some inaccurate, such that entire categories of news or intelligence appear less trustworthy. The “population” of Facebook now exceeds the population of the largest nation-state. We have access to ceaseless and instantaneous communication with everybody, anybody, at any time. These tools can be used to enlighten, but they can also be used to distort, intimidate, divide, and oppress.

On the one hand, autocrats increasingly are empowered by this electronic revolution, enabled to manipulate technologies to solidify their rule in ways far beyond their fondest dreams in times past. On the other hand, individuals can now reach others with similar concerns around the world. People can easily discover what is going on, organize around it, and take collective action.

At present, many countries seek to govern over diversity by attempting to suppress it, which exacerbates the problem by reducing trust in institutions. Elsewhere we see governments unable to lead, trapped in short-term reactions to the vocal interests that most effectively capture democratic infrastructures. Both approaches are untenable. The problem of governing over diversity has taken on new dimensions.

The good news is that the United States is remarkably well-positioned to ride this wave of change if we are careful and deliberate about it. As an immigrant nation, we have always had to govern over diversity. Meanwhile, other countries will face these common challenges in their own way, shaped by their own capabilities and vulnerabilities. Many of the world's strongest nations today—our allies and others—will struggle more than we will. The greater our understanding of other countries' situations, the stronger our foundation for constructive international engagement.

This is why we have embarked on this new project on Governance in an Emerging New World. Our friend Senator Sam Nunn has said that we need to strike a balance between optimism about what we can do with technology and realism about technology's dark side. So we aim to understand these changes and inform strategies that both address the challenges and take advantage of the opportunities afforded by these transformations.

To do so, we are convening a series of meetings and calling for papers to examine how these technological, demographic, and societal changes are affecting the United States (our democracy, our economy, and our national security) and countries and regions around the world, including Russia, China, Latin America, Africa, the Middle East, and Europe.

In past volumes, we have considered the impact of these global transformations on major countries and regions around the world and on international security and democratic processes. In this edition, we will look closer to home, to the future of the U.S. economy in a rapidly changing world.

Artificial intelligence, advanced manufacturing, particularly 3D printing, and the other technologies of the “4th Industrial Revolution” appear poised to transform the world economy. Global supply chains may shift as we learn to produce goods closer to where they are consumed, and AI-enabled technologies may change the relationship between workers and machines, thereby redefining the workplace. At the same time, we are witnessing dramatic workforce demographic shifts.

Although we cannot forecast the exact character of these transformations, history teaches us that they will likely be painful. It also teaches us that they will create new opportunities for prosperity and human flourishing. And, fortunately, the United States is well-positioned to take advantage of that potential.

The papers in this volume consider what advancing technologies and demographic transitions will mean for the U.S. economy and society, and they propose ways for the United States to manage those impacts and take advantage of new opportunities to ensure a growing, productive population. Importantly, rather than simply appreciating the problem, each of our contributors is focused on things that can be done at different levels of government and enterprise, given the social license for it.

The volume opens with an assessment of the future labor market by Erik Brynjolfsson, from MIT's Initiative on the Digital Economy, using a novel dataset from the professional social network LinkedIn. Brynjolfsson identifies skills from today's

jobs that may be performed in the near future by machine learning, finding that machines can do at least some tasks of almost every job, suggesting we will see substantial redesign of work and significant reskilling across the economy. Some occupations, such as retail, are likely to be disproportionately affected. He also considers the degree to which an individual employee's skills and education have value to his or her employer—the more valuable to the employer, the larger the incentive for it to directly invest in its own employees' up-skilling. To us, Brynjolfsson and his co-authors' findings show the value of enabling bottom-up, interest-based frameworks in approaching this transformational but long-term and ultimately individually-tailored challenge. This is an area where the United States can excel.

Following this, we hear from Dipayan Ghosh, a fellow at Harvard Kennedy School and former White House advisor. Ghosh warns that the consumer internet industry's AI-enabled business practices will continue to prioritize revenue over fairness—absent development of a regulatory framework focused in part on improving transparency in firms' practices, and to protect against discriminatory and anti-democratic tendencies. His concerns reflect the growing centrality of today's emerging technologies across a broad swath of society and commerce and the novel governance challenges that they pose.

Turning from technology to people, James Hollifield, the director of the John Goodwin Tower Center at Southern Methodist University, writes that the United States “is trapped in a ‘liberal’ paradox:” it needs immigration to keep its economy strong, but it must also deal with the socio-political ramifications of that immigration. Through a review of past U.S. immigration policies, he explains that this paradox is not new and that we can balance openness with valid concerns for security and societal values. Past contributions to this project have posited that a society's demographics futures, unlike economic forecasts, are perhaps unique in their certainty. As a high immigration nation, however, the United States is enviably able to shape its demography through policy decisions and governance.

Finally, a changing economy and workforce will require an adaptive educational system, and Van Ton-Quinlivan, former executive vice chancellor of the California Community Colleges, explains how community colleges can contribute to preparing workers for 21st-century jobs. This is not the first time that our project's contributors have highlighted the role of nimble educational institutions. Drawing from her experience at the helm of California's overhaul of its community college-based workforce training programs, she outlines how other higher-education institutions can support continuing training and education for a modern economy.

We wish to thank our colleague John Taylor for sharing his insight into U.S. economic productivity and emerging technologies at both the roundtable and the public panel, and we extend our thanks to Gopi Shah Goda, from the Stanford Institute for Economic Policy Research, who moderated the discussions. Finally let us again thank our Hoover Institution colleagues who have supported this project, particularly Rachel Moltz.



George P. Shultz
Thomas W. and Susan B. Ford Distinguished Fellow



James Timbie
Annenberg Distinguished Visiting Fellow

How Will Machine Learning Transform the Labor Market?

By **Erik Brynjolfsson**, Massachusetts Institute of Technology, **Daniel Rock**,
Massachusetts Institute of Technology & **Prasanna Tambe**, University of Pennsylvania

Introduction

The twenty-first century will be the century of intelligent machines. Artificial intelligence (AI) has begun to transform the economy as it enables machines to do more and more of the cognitive tasks that were once done only by humans. In the coming decade, many existing tasks will be replaced by machines, while new ones will emerge. Almost every job will be affected in some way and most will need to be redesigned. Businesses will rise and fall depending on how well they understand, foster, and harness the changing skills that are needed to be productive. Economies will thrive if they can create and update the institutions needed to create these skills.

In particular, the branch of AI known as machine learning (ML) has advanced significantly in just the past decade, largely reflecting improvements in the area of deep learning, a technique that trains large neural networks on large datasets (Brynjolfsson and Mitchell 2017). Three different types of advances, each of about two orders-of-magnitude, have combined to make this possible: 1) an increase in the quantity and quality of digital data, 2) improvements in computational power, reflecting not only the march of Moore's Law, but also new specialized architectures like GPUs and TPUs, and 3) improved algorithms (McAfee and Brynjolfsson 2017). As a result, the performance of ML algorithms has improved significantly. In a highly cited example, the image recognition algorithms on the ImageNet Dataset improved from barely 70% in 2010 to over 97% today, and now surpass human level performance on the same data. Voice recognition and natural language processing, machine translation, recommendation systems, gaming and many other tasks have also seen striking improvements (Shoham et al. 2018). Because capabilities like vision, speech and decision-making are so fundamental for most occupations, these improvements to technology suggest that substantial changes in the nature of work can be expected.

Despite these impressive advances, however, ML is far from being capable of doing the full range of human cognitive tasks. This raises some obvious questions. What

tasks can ML do well, and what tasks are best done by humans? What are the implications for jobs, industries, and different geographies? How can we quantify the changing value of human skills for businesses? In this paper, we seek to address these questions by drawing on several streams of research that have been underway for several years. First, we report on work based on interviews with a set of leading experts in machine learning to develop a set of criteria, or a rubric, for distinguishing which tasks are most suitable for machine learning (Brynjolfsson and Mitchell 2017; Brynjolfsson, Mitchell, and Rock 2019). In turn, we applied this rubric to score 18,112 tasks in 950 occupations spanning most of the U.S. economy to create a guide to how different occupations, industries and regions would likely be affected as the use of ML becomes more pervasive (Brynjolfsson, Mitchell, and Rock 2019). Second, to illustrate how different types of human capital, including skills and education, affect firm value, we draw on ongoing work using data from LinkedIn, Compustat, and other sources to newly quantify these relationships.

Our first set of findings are that while existing ML technologies are not able to automate all the tasks that comprise any of the occupations we study, they are sufficiently advanced to do at least some tasks in almost every job. This suggests substantial redesign of work and significant reskilling will be needed to harness the potential of ML. Our findings suggest that people in lower wage jobs will be disproportionately affected as will those in retailing and transportation industries. People in smaller cities are also more likely to be affected than larger ones.

Investments into the redesign of work can yield significant value for firms. In a second set of findings, drawing on data from LinkedIn, we find that skills and education have value not only to the employees who acquire them, but also to the owners of the companies where those employees work. In fact, the value of IT-related investments has grown dramatically in recent years and, based on a sample of publicly traded firms, as of the end of 2016 amounts to about 39% of the value of installed property, plant, and equipment (about \$8-9 billion in

ITIC per firm in the sample). This implies that firms have a large incentive to invest in creating and updating the skill sets needed to take advantage of emerging IT—most recently, ML—as these technologies become increasingly pervasive. Although many skills will be affected, the skills needed to implement ML are a notable special case: their value has grown markedly as new technologies like Tensorflow have boosted their economic impact (Rock 2019a). The magnitude and scope of the reskilling and business process redesign needed to put ML breakthroughs into practice means that it will require years if not decades before the full effects are felt, just as with earlier technological breakthroughs. (Brynjolfsson, Rock, and Syverson 2018)

The Labor Market's Exposure to Machine Learning Technology

We apply a task-level approach to understanding the effects of ML. This is the most natural unit of analysis for specific capabilities. Detailed information about task-level exposure to ML can then be aggregated to improve our understanding of its effects on many aspects of the economy, including occupations, firms, industries, and regions. In particular, occupations can be considered useful bundles of tasks assigned to similar types of workers. The task-level approach relates worker labor inputs to new types of technological capital within a production function (Autor, Levy, and Murnane 2003; Acemoglu and Autor 2011). Inherent in the managerial decision to replace human tasks with capital services is a trade-off between wages paid to workers and capital rental costs for the machines that could do the same tasks. Increasing machine capabilities or decreasing capital costs for a given task increases incentives to substitute capital for labor in this class of models.

While they are typically stable in the short-run, the set of tasks within each occupation changes over time, as does the nature of many of the tasks themselves. New tasks are created and the value of old tasks changes, altering what the most productive mixture of tasks for a given occupation might be. A related class of models (Acemoglu and Restrepo 2018) elucidates the trade-off between investment in automation technologies and investment in creating new tasks. In these models, increased automation increases the returns to innovative activity in creating new tasks for human labor. In nearly all cases though, the impact of new technology on labor demand is contingent on more than simple human labor task replacement potential.

Making predictions about the impact of ML on labor demand is challenging because any given occupation most often performs a wide variety of tasks. Inevitably, some of those tasks are more suitable for machine learning than others. Focusing on what ML can do with respect to the tasks currently done by workers, however,

can yield insight into which tasks are most exposed to technology. Brynjolfsson and Mitchell (2017), Brynjolfsson, Mitchell, and Rock (2018, 2019), and Brynjolfsson et al. (2019) build, refine, and extend a rubric that seeks to identify the tasks in the economy which have the greatest Suitability for Machine Learning (SML). The rubric consists of 23 evaluative questions with potential answers ranging from 1 (very low SML) to 5 (very high SML). The rubric is used to create a score for 2,059 detailed work activities from the U.S. government's O*NET database. In turn, these scores were aggregated into 950 occupations consisting of 18,112 tasks which share detailed work activities across occupations. After being validated by experts in machine learning and assessed by a team at MIT on a representative set of tasks, rubric evaluation was scaled up to the full set of tasks by respondents on CrowdFlower, as described in Brynjolfsson, Mitchell, and Rock (2018). Subsequent iterations in Brynjolfsson, Mitchell, and Rock (2019) and Brynjolfsson et al. (2019) use data from Amazon Mechanical Turk respondents with some further refinements and improved quality control. The questions are designed such that a "1" (Strongly Disagree) corresponds to low SML and a "5" (Strongly Agree) corresponds to high SML, and neutral exposure corresponds to a score of 3 (Neither Agree nor Disagree). In most of the analyses, the values for each of these individual scores are essentially averaged to get an overall task-weighted occupation-level SML score.¹

Exposure to ML does not necessarily mean that the human labor will be replaced or even reduced in that occupation. As discussed in Brynjolfsson and Mitchell (2017), in addition to substitution, ML can also be used in at least five other ways: to complement labor, to increase demand for it by lowering costs, to change demand by changing overall income, to change information flows and thus information asymmetries, or to drive a reorganization of work. While there has been much emphasis on the first of these possibilities (automation and thus substitution) research suggest that the biggest effect in the coming years will be in driving a redesign of work, as only some tasks in most occupations are suitable for machine learning, while others will continue to require human labor.

Occupations vary considerably in their exposure to machine learning as measured by SML score. Figure 1 below shows the distribution of SML across jobs, tasks, and activities. Very few, if any, occupations are completely exposed to ML. The maximum SML value of any task is a 4.0, with a minimum of 2.13 across all tasks. Strikingly, nearly all occupations have at least one task with a relatively high SML score. Figure 2 shows the count of occupations (vertical axis) against their proportions of task SML above the 90th and 50th percentile of SML (horizontal axis). No occupation has tasks *entirely* in the 90th (or higher) SML percentile, but most occupations have at least some

tasks above this threshold, and almost all jobs have some number of tasks above the 50% threshold.

If ML could do all tasks in a particular occupation, there would be little need or opportunity to re-organize the tasks in that job. It would be fully automated. Likewise, if there were nothing ML could do in that occupation, there would be no reason to re-organize the occupation to unlock the gains from ML technology. The fact that most occupations fall between these extremes underscores the likelihood that machine learning will drive re-organization and re-engineering of how tasks are bundled and assigned into occupations. Indeed, Brynjolfsson et al. (2019) highlight that re-organization of work, not automation or substitution, is the labor demand force with the greatest economic potential for ML (see Figures 1 and 2).^{2,3}

The occupational impact of ML will be shared across many different types of labor, but it will also be uneven. Some occupations, and therefore industries and regions as well, are more exposed to ML than others. Figure 3 shows that lower wage occupations have relatively higher SML scores, though all wages levels have some occupations at either end of the SML spectrum.⁴ Figures 4A and 4B shows the standardized SML scores and Image Data scores (respectively) aggregated by occupation type to the region level. Large cities tend to have lower relative overall SML scores. In contrast, much of the potential for using ML for image analysis value is more concentrated in large cities. Finally, Figure 5 shows that employment-weighted SML by industry. Accommodation and Food Services, Transportation and Warehousing, and Retail Trade are relatively more exposed to the re-organization impact of ML than Education and Health Care (see Figures 3-5).^{5,6,7}

Higher SML tasks are often clerical tasks like balancing accounts or medical transcription or the type of routine work that might be done in a factory (e.g. inspecting items for defects). Most occupations have some component task that does something clerical. But taking advantage of this new technology will require adjustments to how these tasks are performed together. That means changing occupations, but also business processes.

Firms and organizations will have to build new kinds of intangible capital to complement the new types of technological capital created by machine learning advances. At the same time, know-how and tacit knowledge built for the old economic environment will lose value (Greenwood and Yorukoglu 1997). In this way, the coming changes from ML technologies are similar to many earlier varieties of information technology. Part of what is happening now is an extension of the recent past, which has been characterized by waves of investment in networked computing, databases, and other information technologies.

IT, Intangible Capital, and Value

The business process reengineering needed to unlock value from emerging technologies, such as ML technologies, can comprise a growing category of a firm's assets (R. E. Hall 2001; Brynjolfsson and Hitt 2000; Bresnahan, Brynjolfsson, and Hitt 2002). This "IT-related intangible capital" ("ITIC") is the result of investments that firms make into becoming information processing organizations, including investments in business process reengineering that facilitate rapid information acquisition, employee learning, and decision-making. These types of assets are likely to play an increasingly important role in explaining economic outcomes such as growth and firm performance. In some ways, these investments are similar to those that firms make in physical capital such as trucks or manufacturing equipment. Just as units of physical capital enable the conversion of raw materials to goods, units of ITIC enable firms to convert information and ideas into products and services.

The biggest difference is that the measurement of these IT-based intangible assets has proven elusive. Investment into this form of capital is largely invisible (to researchers) and depreciation rates are idiosyncratic and variable, so conventional methods for measuring capital stocks cannot be easily applied (B. H. Hall 1993). For other types of assets, one might use market transactions, such as a lease or resale price, to derive prices, but there are no observable markets for trading stand-alone ITIC—firms cannot sell their management practices or learning cultures.

As an alternative approach to measuring these assets, we can use methods from the literature on intangible assets to derive changes in the prices and quantities of ITIC in U.S. firms. This literature argues that the quantities of a firm's intangible assets can be inferred from the value of its securities (R. E. Hall 2001). The intuition behind this approach is that under reasonable conditions, the value of a firm's securities is equivalent to the value of its capital assets, which is in turn equal to the price of installed capital times its quantity, or equivalently, the ratio of market value to installed capital is equal to its price (equation 1).

$$\frac{V_t}{K_t} = q_t$$

Because firms' investments in these assets are governed by an adjustment cost condition (equation 2), from the marginal adjustment cost function⁸ evaluated at the investment rate at time t :

$$q_t = 1 + c' \left(\frac{i_t}{K_{t-1}} \right)$$

Thus, we have two equations with two unknowns which can be solved to recover quantities of tangible or intangible capital. In other words, even when investment into capital cannot directly be observed, its quantities can be computed using changes in the value of a firm's assets. In the case of IT, the value of a firm's ITIC is not recorded, but these ITIC values can be inferred using proxy measures of investment into a correlated input, such as IT infrastructure (Brynjolfsson, Hitt, and Yang 2002). We use the method described above to derive quantities of ITIC.

A challenge with this approach is that a lengthy series of firm-level IT investment data are needed, and such data have historically been difficult to obtain at the firm-level. To address this problem, we generated an IT series of one of the most important inputs into the construction of ITIC—IT labor—to enable the application of the methods described above.

The data were obtained from LinkedIn, a leading online professional network web site upon which individuals post their employment histories, including information for each job they have held on employer, job title, and years spent at the firm. Employer data generally include name, size, and industry. We use the employment histories of the workers on LinkedIn who identify themselves as IT workers in order to measure the distribution of IT employment in large public firms over a period of approximately three decades.⁹

The length of this labor series is important for two reasons. First, it is consistently available through the period of the dot-com bust, after the crash, and through the last decade including capturing the rise in spending around big data, data science, and machine learning, so we can test how much of the value of IT-related intangibles in the late 1990's was attributable to investor mispricing or when spending on this asset resumed. Second, in the absence of direct measures of spending on IT-related infrastructure, the wages paid to IT labor are among the principal inputs into the construction of ITIC. In other words, firms can build new information structures around old IT capital, so IT wages may be a more relevant input to ITIC investment than IT capital spending.

The key results from the application of these IT labor data to the model described earlier are shown in Figure 6 below. First, by 2016, the stock of IT-related intangible capital in our panel of firms¹⁰ had grown to about 25% of the value of physical capital stock. Despite swings in the value of ITIC around the dot-com boom and bust, firms continued to accumulate IT-related intangible capital well after the bust and through the 2008 recession (see Figure 6).

We also find evidence of significant heterogeneity in the distribution of these IT intangible assets within the U.S.

economy. Figure 7 plots quantities of ITIC by quintile in terms of ITIC stock value. This figure indicates that growth in this form of capital, far from being evenly distributed, has been concentrated in a relatively small set of firms, which is consistent with other patterns of digitization and firm dominance that have been observed in the recent economic data. Higher quantities of ITIC for these firms suggests that they will enjoy significant production-based advantages in the future (see Figure 7).

The Market Value of Skills Used for Implementing Machine Learning

What does our approach to assessing ITIC say about skills that can be used to implement machine learning? The types of intangible assets that firms build change over time to match the IT environment. Changes in technology alter the value of the installed capital base as well, including the value of human capital. Using the LinkedIn skills data, Rock (2019a) finds that Google's open source launch of TensorFlow lifted the market value of AI-using publicly traded firms by 4 to 7%. TensorFlow is a software package that makes it easier for software engineers and data scientists to implement deep learning models. Similar to the way that coding in Python or C++ is easier than coding in assembly language, TensorFlow saves a lot of the effort required to build and train neural networks. Because of that, companies that had intangible capital related to AI increased in value when TensorFlow was launched. The complementary workers these firms needed to generate value from their intangible assets became more abundant. While impactful to the companies involved, this is one relatively small example of the overall rising tide in IT-related intangible assets. Realizing the returns to AI skill investments often demands extensive investment in other IT skills like cloud computing, data engineering, and specialized management. All of these skill varieties require complementary investments in IT-related intangible assets as well (Tambe 2014).

Implications of the growth of ITIC

Our findings suggest that investment in information structures and related skills produces relatively long-lived, durable assets. For policy makers, these findings suggest that the large waves of investment in IT-related intangibles are associated with the development of significant productive capacity and, all else being equal, that this should boost prospects for long-run growth. The fact that ITIC assets behave similarly to other capital assets in recent years is itself interesting. This may be because translating organizational innovations into productive capital requires significant investment in reengineering and skills.

There are, however, important differences between ITIC and physical capital. Unlike most types of physical capital, ITIC has diminished value outside the context of the firm.

This has important implications for firm valuation and acquisition. Development researchers have traditionally looked at capital accumulation as an engine for growth. The lack of secondary markets for IT-related intangible assets ties these questions together in an important way to firm health. When firms are dismantled, ITIC is likely to disappear. Therefore, it is worth continuing to further our understanding how the rising importance of ITIC, in conjunction with notable differences between ITIC and physical assets, impacts economic behavior.

Conclusion

Since the industrial revolution, general purpose technologies like the steam engine and electricity have driven economic growth and higher living standards. We believe that the most important general purpose technology of our era is AI, especially ML. A key feature of general purpose technologies is that they enable complementary innovations and investments. Among the most important complements are new skills and new business processes. Unlike plant and equipment, skills and process innovations are intangible. This makes it difficult to measure and makes it subject to underinvestment and malinvestment. By using the SML framework, we can identify some of the tasks that are most likely to be affected by machine learning, as well as some of their associated skills.

The large increase in IT-related intangible capital in general, coupled with the surge in the value of skills needed for implementing ML in particular, suggest an important role for businesses reskilling the workforce. Our analysis shows that shareholders stand to benefit from when the employees of their firms have more of the right sets of skills. However, while the advances in technology have often been breathtaking, the reskilling of the workforce and the reinvention of business processes has lagged. This is reflected, for instance, in declining business dynamism according to work by Decker et al. (2016). The key bottleneck for unlocking value often is not technology but people. Therefore, for adapting to ML enabled work, the grand challenge of the 21st century will be speeding the adoption of new skills and organizational practices that support these technologies.

Acknowledgements

This work was done in conjunction with the MIT Work of the Future Task Force. We received access to LinkedIn data as participants in the LinkedIn Economic Graph Challenge. We thank the MIT Initiative on the Digital Economy for generous funding.

Bibliography

Acemoglu, Daron, and David Autor. 2011. Skills, Tasks and Technologies: Implications for Employment and Earnings. *Handbook of Labor Economics*. Vol. 4.

Acemoglu, Daron, and Pascual Restrepo. 2018. "The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment." *American Economic Review*.

Autor, David, Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4):1279–1333.

Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *The Quarterly Journal of Economics* 117 (1). MIT Press:339–76.

Brynjolfsson, Erik, Morgan R. Frank, Tom Mitchell, Iyad Rahwan, and Daniel Rock. 2019. "Machine Learning, Labor Demand, and the Reorganization of Work."

Brynjolfsson, Erik, Lorin M. Hitt, and Shinkyu Yang. 2002. "Intangible Assets: Computers and Organizational Capital." *Brookings Papers on Economic Activity* 2002 (1):137–98.

Brynjolfsson, Erik, and Lorin M Hitt. 2000. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *Journal of Economic Perspectives* 14 (4):23–48.

Brynjolfsson, Erik, and Tom Mitchell. 2017. "What Can Machine Learning Do? Workforce Implications." *Science* 358 (6370):1530–34.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock. 2018a. "Machine Learning and Occupational Change." *Unpublished Working Paper*. MIT.

———. 2018b. "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" *AEA Papers and Proceedings*, 43–47.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2018. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." In *Economics of Artificial Intelligence*. University of Chicago Press.

Decker, Ryan A, John Haltiwanger, Ron S Jarmin, and Javier Miranda. 2016. "Declining Business Dynamism: What We Know and the Way Forward." *American Economic Review* 106 (5):203–7.

Greenwood, Jeremy, and Mehmet Yorukoglu. 1997. "1974." In *Carnegie-Rochester Conference Series on Public Policy*, 46:49–95.

Hall, Bronwyn H. 1993. "The Stock Market's Valuation of R&D Investment during the 1980's." *The American Economic Review* 83 (2):259–64.

Hall, Robert E. 2001. "The Stock Market and Capital Accumulation." *The American Economic Review* 91 (5):1185–1202.

McAfee, Andrew, and Erik Brynjolfsson. 2017. *Machine, Platform, Crowd : Harnessing Our Digital Future*. Harnessing Our Digital Future.

Rock, Daniel. 2019a. "Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence."

———. 2019b. "Essays on Information Technology, Intangible Capital, and the Economics of Artificial Intelligence." Massachusetts Institute of Technology.

Shoham, Yoav, Raymond Perrault, Eric Brynjolfsson, Jack Clark, James Manyika, Juan Carlos Niebles, Terah Lyons, John Etchemendy, and Z Bauer. 2018. "The AI Index 2018 Annual Report." Stanford.

Tambe, Prasanna. 2014. "Big Data Investment, Skills, and Firm Value." *Management Science* 60 (6):1452–69.

¹ Four of the questions concern the presence or use of four different types of data: speech, image, text, and structured. One of the inputs to the overall SML score is the maximum of these four scores, while the data scores themselves are omitted. This is because an occupation need not have all types of data available in order to be high SML.

² Brynjolfsson, Mitchell, and Rock (2018a); Rock (2019)

³ Brynjolfsson, Mitchell, and Rock (2018a); Rock (2019)

⁴ Using BLS 2016 wage levels.

⁵ Brynjolfsson, Mitchell, and Rock (2018a); Rock (2019)

⁶ Brynjolfsson, Mitchell, and Rock (2018a); Rock (2019)

⁷ Brynjolfsson, Mitchell, and Rock (2018a); Rock (2019)

⁸ $c' \left(\frac{I_t}{K_{t-1}} \right)$ is the marginal adjustment cost for the firm at the rate of investment (investment at time t divided by capital stock in the previous period).

⁹ Not all workers participate on LinkedIn, so sampling error in the measures will tend to bias OLS estimates downward by an amount that is related to the share of the total variance of the measure that is attributable to the variance of the sampling error. However, the sample size suggests that any measurement error should be less than in the CI IT capital data, which prior work has claimed may be as high as 30-40% of the total measure variance (Brynjolfsson and Hitt 2003). LinkedIn includes the professional profiles of a very high fraction of U.S. IT workers. Therefore, measurement error related to differences in web site participation is likely to be fairly small.

¹⁰ To achieve a balanced panel, we limit the sample of firms in our analysis to those that appear in every year. Therefore, our sample skews towards larger, well established firms.

Erik Brynjolfsson is the director of MIT's Initiative on the Digital Economy and professor of management science at the MIT Sloan School. Daniel Rock is a doctoral candidate at MIT Sloan, and Prasanna Tambe is an associate professor of operations, information, and decisions at the University of Pennsylvania's Wharton School.

Supporting Data

Figure 1. Distribution of Counts of Suitability for Machine Learning (SML) Score for Occupations, Tasks, and Detailed Work Activities


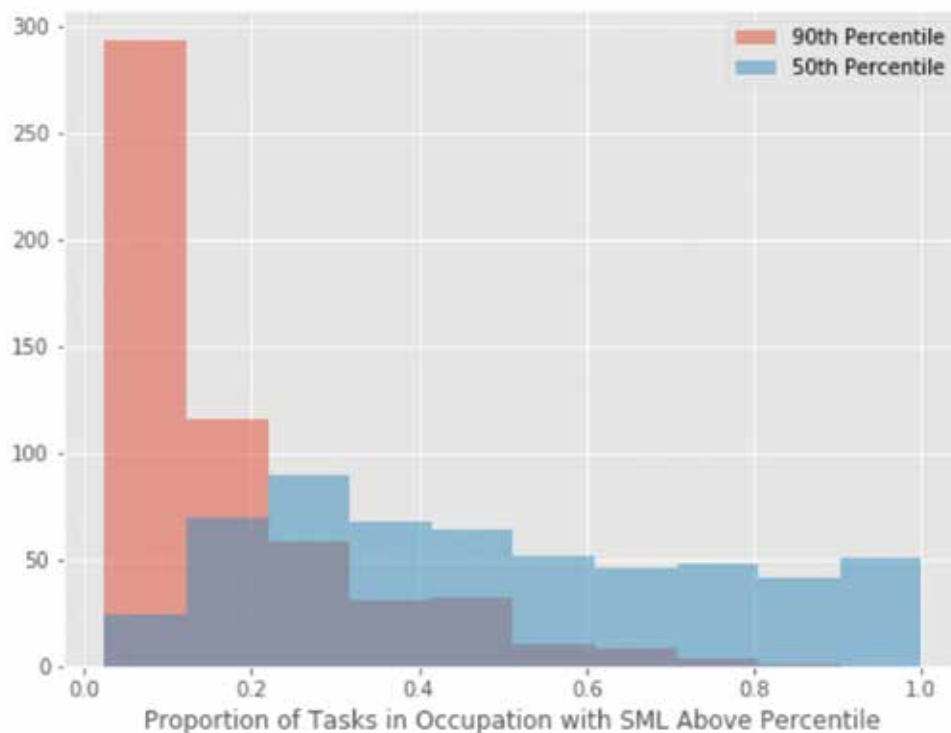


Figure 2. Histogram of Occupations by Proportion of Tasks with SML Larger than 90th and 50th Percentile Thresholds



Source for Figures 1 and 2: Brynjolfsson, Mitchell, and Rock (2018a); Rock (2019)

Figure 3. SML Score vs. 2016 Median Wage Percentile;
 Regression Coefficient: -0.0034 (t-stat = 18.5)

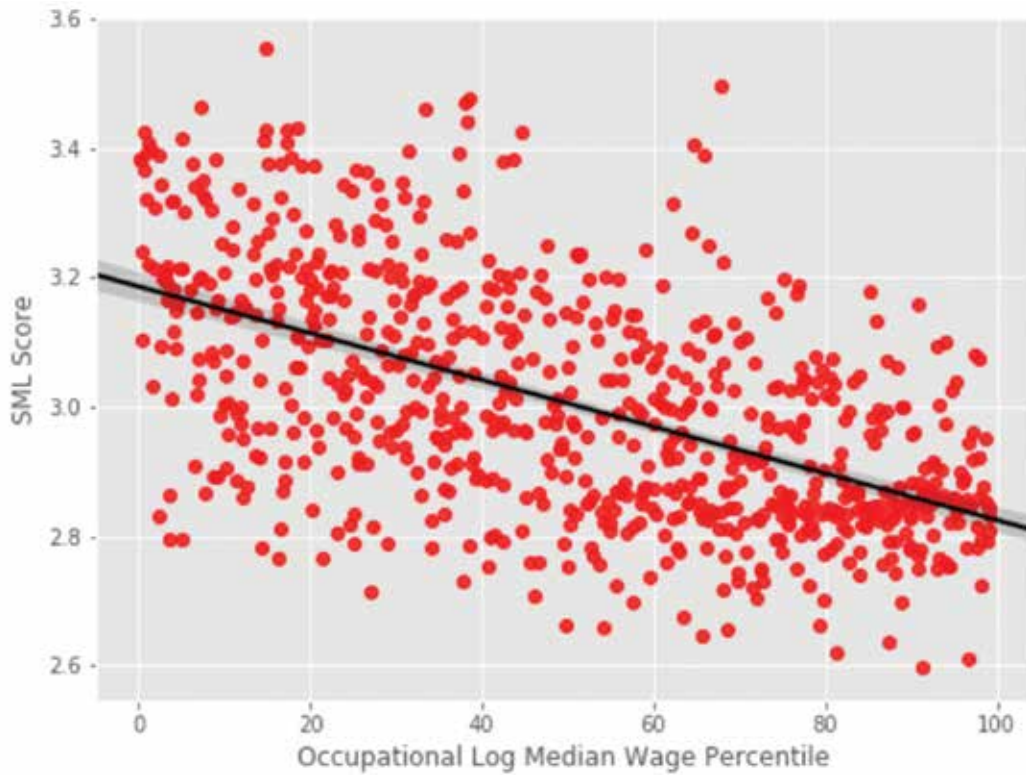
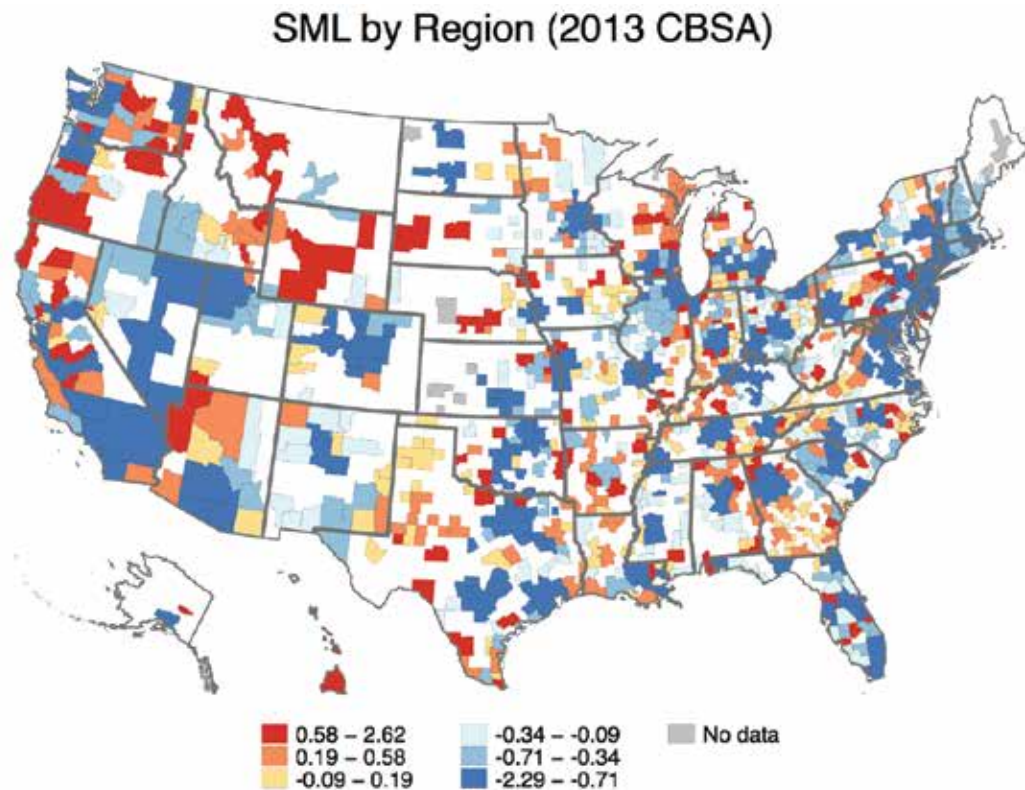


Figure 4A. Standardized SML Score by Region



Source for Figure 3: Brynjolfsson, Mitchell, and Rock (2018a); Rock (2019)

Figure 4B. Standardized Image Data Score by Region

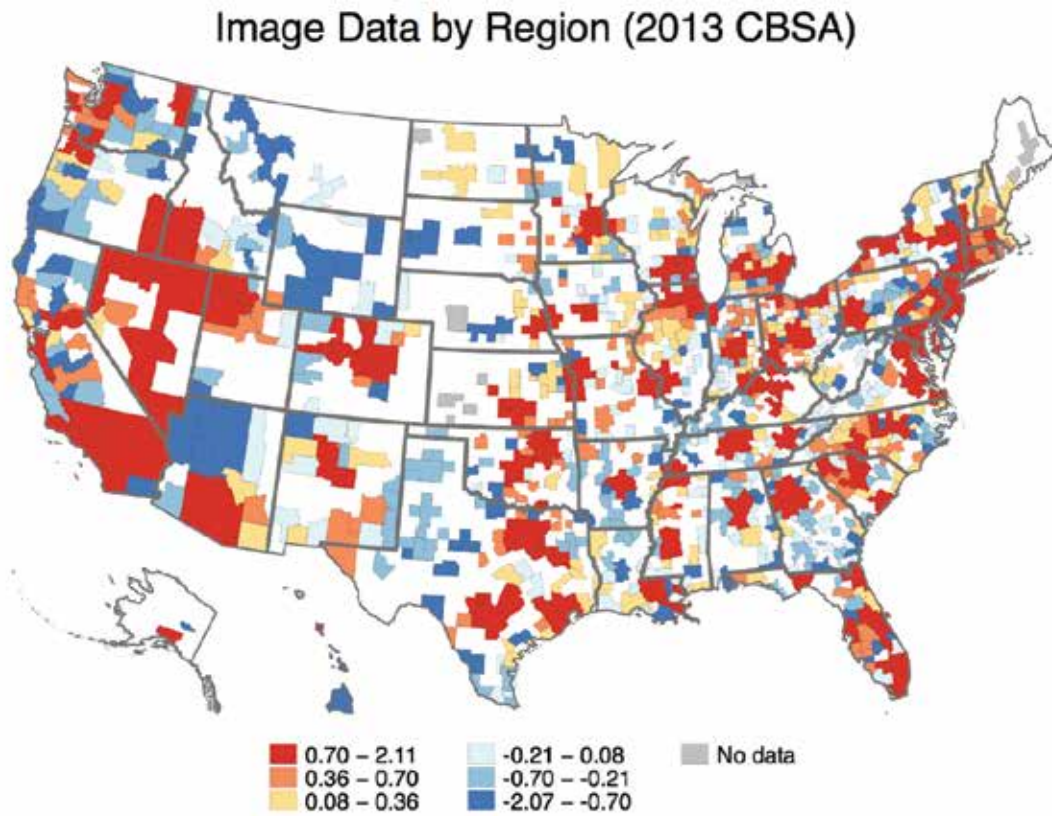
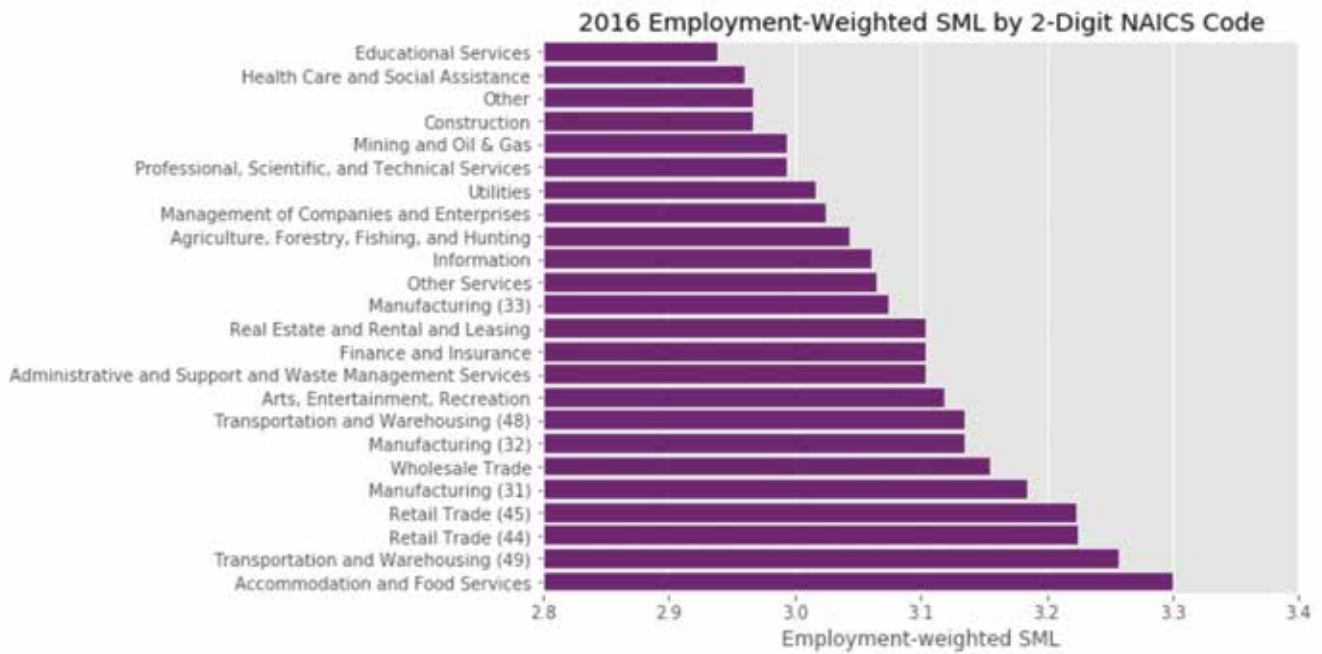


Figure 5. Employment-Weighted Average SML by 2-Digit NAICS Industry



Source for Figures 3-5: Brynjolfsson, Mitchell, and Rock (2018a); Rock (2019)

Figure 6. Change in Quantities of ITIC and PPE for Publicly Traded Firms from 1987 to 2016





Figure 7. Quantities of ITIC by Firm Quartile



The Commercialization of Decision-Making: Towards a Regulatory Framework to Address Machine Bias over the Internet

By Dipayan Ghosh, Harvard Kennedy School

Abstract

The consumer internet has exacerbated the discrimination problem. The business model that sits behind the front end of the internet industry is one that focuses on the unchecked collection of personal information, the continual creation and refinement of behavioral profiles on the individual user, and the development of algorithms that curate content. These actions all perpetuate the new pareto optimal reality of the commercial logic underlying the modern digitalized media ecosystem: that every act executed by a firm, whether a transfer of data or an injection of content, is by its nature necessarily done in the commercial interests of the firm because technological progress has enabled such granular profiteering. This novelty in the media markets has created a tension in the face of the public motive for nondiscriminatory policies; where adequate transparency, public accountability, or regulatory engagement against industry practices are lacking, it is directly in the firm's interest to discriminate should discriminatory economic policies suit its profit-maximizing motive. This paper discusses this technological development and offers policy responses to counteract these breaches against the subjects of internet-based discrimination.

Introduction: The Centrality of the Consumer Internet

The importance of the consumer internet in the context of the modern media ecosystem is unquestionable. Economic opportunities in housing, employment, and other objects of the consumer marketplace; national political concerns and the systemized dissemination of political communications; and social interactions that mirror or conversely define our sociocultural norms: these are all clear and evident results of the growth and present breadth of influence of the consumer internet.

The consumer internet is comprised of the firms that operate over the internet and interface directly with consumers—Facebook, Apple, Google, Netflix, Spotify, and Amazon among them. Consistent across the sector is a set of practices—constituting the business model that sits at the heart of the internet—driven by (1) the development of tremendously engaging platforms

that surface the ranked content that the firms predict consumers most wish to see and will therefore engage with; (2) the uninhibited collection of the consumer's personal information all to the end of generating behavioral profiles on the consumer that record the consumer's likes, dislikes, preferences, interests, routines and behaviors; and (3) the refinement of highly sophisticated but equally opaque algorithms that curate content to fuel the first practice and target ads by taking advantage of the second one. This economic engine, consistent across the consumer internet, is depicted in Figure 1.

To be sure, there are two caveats to clear before moving forward. First, while this business model is in clear use within the walled gardens of such firms as Facebook and Amazon, each firm adopts it in its own way, using its own proprietary processes and propensity for personal data collection along with its singular value proposition for the consumer market, to take advantage of the profits the general business model can yield. And second, this business model is utilized to varying degrees by the subject firm; that is to say that there may be other core practices and contributions to company revenue that are also critical to the subject firm and operated in parallel to the consumer internet offering. To mention a few examples, Amazon and Google are market leaders in the provision of cloud computing services; Apple's core revenue is generated from the sale of consumer device technologies; and Netflix maintains an order-by-mail video rental service that has relatively little to do with the aforementioned business model leaving aside the agglomerations of personal interests derivable from physical rentals. But to reiterate, it is the set of practices leading to the business model illustrated above that constitutes what I mean by the "consumer internet"—and which I hope to scrutinize further and critique in this essay. This is particularly because it is this business model that has instigated and perpetuated the negative externalities that we care about protecting the public from today and in the way forward—precisely because the business model has promoted an insidious economic logic that aligns the interests of nefarious actors with those of the internet platform firms.

To ask *why* these firms have uniformly adopted this business model is pertinent. As others have discussed, the internet industry is one that operates in a free commercial zone—it is, in other words, a radically free market that favors and rewards open capitalism.¹ In such jurisdictions as the United States, we continue to lack a federal standard on privacy and most other public interest concerns that would otherwise concern the firms in this sector. This fundamental lack of consumer and citizen rights in the United States has enabled the internet firms to have a free pass to take advantage of the free market zone. And take advantage they have, just as suggested by the institutional directive within Facebook to “move fast and break things.”

This lack of a regulatory regime has in turn meant that these firms have developed in a manner practically independent and uncaring of the public interest save when it serves their commercial interests. As with any business, the public interest need not be considered from their perspective; only the shareholders need be served. It is thus effectively unnatural to ask a chief executive in the industry to bend the knee to consumers; in the absence of meaningful economic regulations that target the capitalistic overreaches of the business model, nothing can save the public from the overreaches of the industry. And while public sentiment might swell to such a degree at times that it might appear the effective situation for the firms has changed for good because of ongoing public outcries, the lack of actual regulatory movement by the government equates to the free zone of commerce remaining intact. The public's memory is short; equivalently, the industry often moves directly back into the zone of commercial operation it did prior, perhaps under new disguises to protect itself from regulators. One could suggest that this is precisely what has happened in regard to the Cambridge Analytica disclosures of March 2018; while there was a cacophony of public outrage immediately following the whistleblower's revelations and the corresponding reports of sharing of 87 million Facebook users' information with the digital strategy firm engaged by Donald Trump's campaign for the presidency—outrage that led Mark Zuckerberg to testify before Congress mere weeks after the revelations—there is relatively little U.S.-led discussion now about what economic regulations should be passed to truly hold Facebook and like firms to account.² The stunted progress of the Honest Ads Act introduced by U.S. senators is ‘Exhibit A.’³

It is due to this unrestrained progress of the business model—particularly the constant quest by firms like Facebook and Google to maximize through whatever means necessary and possible the amount of time users spend on the platforms—that the leading internet platforms have overtaken the media ecosystem in the United States.

The Science of Machine Bias

Robust discussion has developed in recent years, particularly since the boom of the big data economy, concerning the potential of machine learning algorithms to systematically perpetuate discriminatory results in various fields from medical science to educational opportunities.⁴ Of primary concern is the development of machine learning models that engage in automated decision-making. While the methods underlying the application of machine learning are mostly taken from the traditional statistical literature and largely do not constitute mathematical novelty in general terms, cultural circumstances and advances in computing have popularized the term and expanded interest in the industry.

Supervised machine learning models are typically designed through a combination of human input and automated statistical analysis of a dataset. A dataset such as the demographic inferences Google draws over a class of users in a given region typically carries some implicit pattern; different classes of users might execute searches at various times of day, from various locations, and with varying frequencies—indicators called “features” since they are independent attributes associated with an instance, in this case an individual user. Machine learning models attempt to draw such relationships out of the data to develop inferences about the true nature of the population. A human—or in the case of unsupervised models, a machine—might code each user as participant in a particular class based on the user's individual features, inferred through analysis of the user's on-platform behavior, off-platform activity, and demographic data, thereby generating a set of “training data” that can be used to help the model learn how to classify future data points. A dataset including a population of such users might have some observable relationships consistent between certain classes of the population. These relationships are drawn into the machine learning “model” in the form of a set of decision rules—a series of inferences about the population developed from observation of the dataset that can then be implemented as a “classifier” of future objects subject to the model's classification regime. This implementation can then be executed on an automated basis such that when new observations come into view they can readily be analyzed and classified by the model.⁵ The continual refinement of learning models through feedback from real world routines and behaviors is illustrated in Figure 2.

Taking YouTube as an example, we would regard the platform's video recommendation system by which the user is suggested a video to watch next as the model or classifier, which operates over a set of decision rules established by the machine learning model developed and refined by the company on an ongoing basis. The company's commercial objective is to engage the user,

thereby enabling it to collect increasing amounts of information about the user's habits and preferences, and to generate ad space that it can sell to the highest bidder interested in persuading a set of users. Particularly at the outset of algorithmic design, a team of humans might be employed to classify a global set of users into various categories for each feature. A feature concerning type of use of the platform, for instance, might include classes such as "channel operators", "power users", "frequent users" and "occasional users." Additional features might include demographic details, information pertaining to the user's historical use of the platform including which videos the user watches and which channels he or she has subscribed to, information pertaining to the use of other Google services, and position in the relationship graph network among others. YouTube then might train a model that analyzes how the existing "observed" data points concerning the company's users were classified. This analysis of observed data points is used to develop and train, on an ongoing basis, a set of decision rules that constitute the classifier model that can determine based on statistical analysis which class new data points—new users of YouTube, for example—should enter. Finally, this algorithmic inference then determines what videos the platform will recommend to the user. Feedback loops incorporating accuracy of predictions (whether reported by the user or inferred by the platform due to a user's disengagement or other negative behavior) can be used to refine the model over time. This in turn leads many users down a path of watching a long series of highly engaging videos described by some as going down the YouTube "rabbit hole."

Many have described machine learning models—and more generally, algorithmic processes—as fair, or at the least, fairer than a human would be in making the same decision. This idea has been wholly rejected by most. While theoretically algorithms could be designed in a manner that is contextually "fair," one question that naturally arises is what fairness (even in context) should actually mean; different parties might have different definitions in practice, and even with consensus on the meaning of fairness, machine learning models have been shown to discriminate. Another concern is that it has proven to still be difficult to design machine learning algorithms in a manner that foresees all potential forms of fairness and preempts them through reorientation of the algorithm. In the case of YouTube, for example, reports suggest that the recommendation algorithm has had a longstanding tendency to suggest users watch conspiracy-laden videos including the "Momo" hoax that targeted children⁶ and the "flat-earther" myth.⁷

At issue is the propensity for most machine learning models to discriminate; in fact, this is precisely what they are meant to do: discern the characteristics of an incoming data point and infer, based on its features, which class

it belongs to. Presumably, such models are used to give potentially different treatment to data points that occur in different classes. In the case of YouTube, established sports fans might consistently be recommended to watch videos related to sports; those interested in foreign policy might be subject to recommendations to watch political videos.

Title VII of the Civil Rights Act in the United States offers protection from unfair decisions made on the basis of any protected categories, including race, gender, pregnancy, religion, creed, veteran-status, genetic testing status, ancestry, and national origin. (Importantly, political discrimination is not included here.) Various laws institutionalize further protections, among them the Age Discrimination and Employment Act, which states that employers cannot terminate an employee simply because of age; there must be some substantiation that the employee no longer can work effectively. Similarly, the Americans with Disabilities Act prohibits employer discrimination against those individuals who can work effectively despite their disabilities. Various state laws go further than the federal laws and institute further protections from discrimination, particularly through added protected classes and other expansions including, for instance, new lower age thresholds to trigger the age discrimination law.

Developing civil rights jurisprudence carries two principal mechanisms for protection from discriminatory outcomes: disparate treatment and disparate impact.⁸ In a typical disparate treatment case, a potential employer might suggest that the candidate should not be hired because he or she is the member of a protected category. This sort of determination would amount to an intentional violating decision to discriminate against the candidate because of his or her protected class status.

But in the realm of machine bias, disparate impact cases are typically of greater concern because of the manner in which learning algorithms engage in automated classifications over which decisions—which could be vitally important to the data subject in question—are automatically applied and implemented against many data subjects together according to a set of rules contained in the model. Disparate impact cases typically refer to instances in which a particular decision has greater resulting impact on a protected group than on the rest of the population. Harmful disparate impact can trigger an investigation against the liable party. And a decision such as a hiring policy might be "facially neutral"—where the decision rule does not appear to be discriminatory on its face—but if when carried out in practice it results in a harmful disparate impact against a protected group then civil rights protections may be triggered.

It may be the case that a learning model used to classify users in a consumer internet application—for instance, in

the context of identifying the consumer group at which to target an ad campaign that includes a set of political messages—might attempt to maximize clickthrough rates or some other engagement or revenue metric applied by the platform firm. The learning model might identify characteristics in regard to a number of signals (or in this case, features) about the messaging and the advertiser's intended target audience, for instance that political ads feature men and masculine themes, as well as issues that may appeal to certain socioeconomic classes, and that the geographic region the advertiser wishes to target is in the Midwest. In such cases, it is likely that the algorithm will determine that the target audience that will yield greatest engagement for the advertiser and the platform is some group that is primarily male, wealthy, and Midwestern—which, it could be said, is a necessarily harmful discriminatory targeting practice given that certain protected classes are not included in the target audience. All that said, such targeting is likely not illegal for several reasons. First and foremost, there might be no civil rights laws that covers the content of the ad campaign in question since American laws primarily cover various economic opportunities but not social or political ones. Second and perhaps more critically, it might be the case that, even if the ad content is covered by civil rights laws and pursues a discriminatory execution of dissemination that prevents certain protected classes from seeing it, the classifier was technically "fair." In such cases if a suit is pursued then the platform firm that enabled the targeting may have to respond to the question of why the algorithm screened out an inordinate proportion of, say, women. Should the firm be able to offer a justifiable business reason then it could be adjudged that it did not engage in unfair discriminatory practices leading to harmful disparate impact.⁹

Broadly, the utilization of learning models can produce discriminatory outcomes through two main means: the nature of the training or input data, and the design of the learning algorithms themselves. Underlying each of these primary themes is a more human concern: that the data miner him- or herself could be (intentionally or unintentionally) biased and carry that bias into the programming of the model and analysis of the data.

Discriminatory Concerns Related to Training Data

There is a longstanding refrain in the field of computer science: "garbage in, garbage out." Machine learning models are "trained" through the analysis of the aforementioned training data, which in supervised learning schemes might be classified by humans. Data points—such as a typical Google user—has a set of attributes about his or her use of the company's platforms that can be used to classify the user into certain audience clusters. Inferences about new users to the company's systems are then made by the learning model. But as was discussed in a recent White House report, poor design

of training data can promote discriminatory outcomes.¹⁰ There are two primary mechanisms by which flaws in training data can perpetuate discriminatory decision-making.

The first is in the process by which the data is organized. Historical datasets on which training data is based typically come with certain mutually exclusive class fields as discussed above in the YouTube example—but the selection of class fields and attribution of data subjects to them occurs at the hands of humans in supervised learning premises. The individuals who organize these class fields—the data engineers responsible for development of learning models—attempt to define a paradigm through the identification of class fields that they believe most fairly and effectively reflects the situation of the real world. For instance, it might be the case that to make determinations about the creditworthiness of a loan applicant, credit agencies decide that it is most critical to understand his or her net worth, demographic information, profession, education level and relate categories—but that it is less important or particularly difficult to include information related to the individual's personal life goals, trustworthiness, and commitment to paying back the loan. This can germinate a form of bias in the designation of class fields, as such determinations to include and exclude certain categories could diminish the chances of a positive decision for certain demographic groups while elevating the chances of others. The creditworthiness example can be translated readily to the consumer internet context: firms continually refine ad-targeting algorithms so as to advance the commercial interests of the advertisers by offering them maximized bang for buck with the data that they have at hand. Whether the advertiser is a credit or housing or employment agency or another client, the tendency for all the parties at hand will be to promote profits over protecting the consumer's interest given the lack of any sort of legitimate nonpartisan scrutiny over firms in the digital advertising and consumer internet sectors.

The second major family of discrimination concerns that might arise from poor design of training datasets is attributable to the data that populates the datasets itself. Two main problems can be responsible for this: incorrect data and selection bias. In the first case, data might be outdated or otherwise contain inaccuracies about the population that perpetuates bias since the incorrect data is used to train the classifier model. For instance, if loan payback periods are incorrectly reported to be longer for some individuals than others, then those individuals might be adversely affected by the decisions executed by the resulting model trained on the inaccurate dataset. The second case, selection bias, is often subtler and involves the collection of data that is not representative of the population which, if used to train the resulting learning model, projects the inferences learned from the

biased training set on current decisions, likely resulting in biased decisions. A simple example of biased input data occurred in the case of the StreetBump application developed in Boston; the mobile application was designed to enable residents to report the occurrence of potholes to the app developer and the idea was seen as so successful in enabling crowdsourced reports that the municipality engaged the developers to know when and where to dispatch repair teams. After some time of use, however, it was found that repair teams were disproportionately dispatched to wealthier and younger neighborhoods—parts of the city that presumably had more people who owned smartphones and greater local propensities to participate in the crowdsourcing functionalities offered through the application. The city was, in other words, receiving a biased selection of the data; a truly representative set of data would report relative frequencies of potholes across neighborhoods in the city in proportion to their true occurrence. The repair service dispatching decision process thus could only produce biased results without some counteractive measure to replace a more representative sampling of data or tweak the algorithm such that it could correct for the direct harms that came to the neighborhoods that were less well-off.¹¹

Related to the concerns around bias emerging from the training data is the capacity for learning models to suggest discriminatory decisions based on such biased datasets. Training data might only contain information at a level of granularity that disadvantages certain groups. Such issues around the granularity of the datasets in question lead to such potentially discriminatory practices as redlining, in which certain inferences are drawn about individual neighborhoods—inferences that are extended to advise decisions made about any residents living in those neighborhoods. If the data suggests that on average a certain zip code earns relatively little the inference could be that it will therefore yield low click-through rates on ads and eventual purchases of interesting market opportunities—and thus anyone living in that neighborhood could be subject to a discriminatory outcome that may constitute a harmful disparate impact upheld by the courts should the harm occur in regard to, for instance, a housing opportunity.

Algorithmic Design

Machine learning algorithms carry the bias contained in data inputs and reflect those biases as the model learns based on the makeup of the training data. But critically, there are additional concerns that can result from the mechanics of traditional statistical analysis as well.

Foremost is the common fallacy in statistical analysis that correlation necessarily implies causation. We know this not to be true; it might be the case, for example, that certain racial groups have higher education levels than others,

but this does not suggest that certain races are more intelligent or hard-working than others. Though this issue has been surfaced with machine learning models, there are mechanisms to curtail its prevalence proactively, in much the same way that certain explanatory variables are excluded from regression models because they are redundant or misleading.

Perhaps more deeply concerning, a poorly designed machine learning model—or one that is ill-equipped to fully handle the problems of discrimination, especially in fields that are not subject to strict regulations like personal finance or housing—may drift over time in such a way that perpetuates biased outcomes for marginalized people. This problem is distinct from the initial training of a model; indeed, trained models implemented in the consumer internet industry are refined on an ongoing basis so that they reflect the user's desires to the greatest degree possible. But what happens when an algorithm exceeds its intended purview and presumes things about us as individuals or as a population that just are not true—or even worse, encourages engagement of our less virtuous tendencies? There is a widely known statistical concept that describes a related tendency: “confirmation bias,” whereby the model—or its designer—finds what might be expected given cultural norms, instead of the reality. The broad propensity for machine learning to “drift” in such directions presents a veritable thicket of concerns regarding bias. For instance, a model might learn from original training data that has been carefully engineered and monitored by the data miner to limit occurrences of unfair discrimination—but at their heart, learning models are designed to cut corners, to efficiently make decisions and determinations about a population in a way that approximately understands the true nature of the real world and reflects that in its algorithmic design, and as such, they are designed to discriminate. This natural tendency for them to attempt to find ways to discriminate in whatever legal manner possible organically forces them to tend toward overstepping the boundaries that have been set for them through secondary backdoors, and this enforces within the model an economic logic that drives them to acquire new behaviors through novel discoveries about the real world. But what happens when those so-called “discoveries” that advise the decision-making algorithm are outsized or otherwise biased? This is the type of model drift—through ongoing observation of the real world—that can engender discriminatory behavior. It is this characteristic of machine learning that can cause models to systematically feed voter suppression content to underrepresented minorities or send nationalistic groups down hateful pathways on social media. A generic conclusion depicting how this might work is illustrated in Figure 3; while the learning model might treat representative cases across a sample population by developing a reasonably accurate decision system for the majority, it might not reflect the particular situation of the minority.

An additional concern that can subvert antidiscrimination efforts is the organic generation of so-called proxies as the model is trained from the input data. It may be that a machine learning model is designed to exclude the use of any protected class data in the course of statistical analysis so as to explicitly protect against discriminatory outcomes against those protected classes. Models might learn, however, that there are alternative “proxies” that are equivalently descriptive of the protected class categorization as the protected class data itself. For instance, an algorithm prevented from accessing race information pertaining to the population might determine that some combination of other class fields—such as location of residence and name—might be used in tandem to generate through the back door an understanding of the individuals’ racial group category. Further, such inferences might be completely non-transparent to the model’s engineers, since they typically occur silently premised on the data already provided as input to the model, and proxies are not proactively reported to the designers as they are generated by the learning algorithm in the course of maintaining and updating the classifier model.

And as a final note, there is a robust active conversation in regard to what should be considered “fair” in the first place. Should fair mean whatever is lawful—and correspondingly that everything outside the reach of the law is on-limits and therefore fair? That is essentially how the industry today operates—and it is the underlying free market economic design of the United States that in fact enables and encourages such capitalistic “innovations” as discriminatory decision-making executed by artificial intelligence so long as it does not constitute harmful disparate impact in the areas of industry protected by federal civil rights law. In this way, the vast majority of the consumer internet’s industrial activity falls directly outside the purview of federal laws in the United States—unless of course the business activity concerns American civil rights laws as has been suggested by the American Civil Liberties Union about a narrow sliver of Facebook’s advertising platform.¹²

These harmful effects are supercharged when it becomes the direct commercial interest of the party developing the learning model to develop a classifier that maximizes revenue. In such an environment, potential discriminatory outcomes are a mere afterthought.

The Radical Commercialization of Decision-Making

One could question whether or not the fact that the internet firms have overtaken and now define the western media ecosystem is in fact a negative thing; perhaps it is for the best in that it breaks the centralization of the creation of content. A truly social platform elevates not

necessarily the content generated by actors in the mainstream media like mainstream newspapers but rather those issues and elements raised, reported and reposted by the common user, and particularly a mix of those posts that are (1) predicted to be interesting to the user in question and (2) which have received wide circulation. (Atop these factors are more including the explicitly expressed preferences of the user, who might for instance choose to see the News Feed in chronological order, obviating some of the concerns recently associated with social media platforms.) Thus the traditional central power of large media companies—that epitomized by say the Hearst Corporation among other examples of the twentieth century—is somewhat diminished by the nature of the internet and the internet platforms themselves as social media receives more attention from the younger generations of the population than traditional forms of print media that also offer access to the news. And in fact, most of the appeal of social media originates from its capacity to connect us with issues and ideas that matter in our individual lives—issues that would not appear on traditional media formats at all—more so than the more abstracted concerns of the mainstream media.

Where power has waned amongst the producers of news media, however, the power of the internet platforms has quietly emerged—albeit in very different form. While power for traditional media firms lies largely in defining and producing content for broad dissemination and consumption, internet firms in large part do not participate in content production. Google’s value proposition is instead focused on offering the efficient and effective classification and searchability of content (including news); for Facebook, it is offering seamless connection and engagement across the user’s friend graph; and for Twitter, it is the attribution of ideas and engagement against them by the broader user population.

But it is not only provision of these services in broad terms that distinguishes and strategically separates Google and Facebook from their competition—if that were the case then there would be far more effective competition against these firms. A key part of their ongoing commercial strength in fact lies in their first-mover advantage¹³ in seizing the reins of the consumer internet business model premised on the creation of advertising exchanges at a time when we also saw the coinciding rise of capacity in two technologies: data storage and computing. Just as Google and Facebook settled on their advertising-based business models these two technologies surpassed a key threshold that triggered the rise of the “big data” economy.

It is the combination of these technologies—the novelty of the targeted advertising regime created and commercially promoted across the media ecosystem alongside the coinciding rise of big data capacity—that, along with their nominally unique consumer services,

set them on their historic trajectory. What has emerged, though, is a commercial regime underlying the entire consumer internet that is algorithmically trained for the maximization of monetary opportunity subject to few constraints.

It is throughout the three pillars of the aforementioned business model that describes the consumer internet's practices that advanced machine learning systems are implemented for gains in profit—and equivalently, it is throughout each of these core practices that there is tremendous capacity for discriminatory results pushed onto the individual consumer. On a continuous basis, algorithms are trained to understand the consumer's preferences, beliefs and interests all of which are shuffled into the individual behavioral profile; keep the user engaged on the platform by understanding and ranking all content existing in the realm of posts that could be populated in the user's News Feed; and push digital advertisements at the user with which he or she will be likely to engage. In a sense, then learning algorithms are continuously and ubiquitously used by the firms leading the internet industry to infer as best as possible what the individual's true nature is and what arrangement of content should be pushed at the individual to maximize profits for the service operator.

I describe this as the “commercialization of decisions”—and it is radical because of its continual engagement and refinement, and its total ubiquity across the sector. All decisions made by learning algorithms in the context of the consumer internet are now necessarily commercialized in light of the combined strengths of supercharged big data technologies and platform power. That is to say that each decision made by a consumer internet learning algorithm—be it over determination of what content to push at the user or inference of the user's character, or some other narrower practice—is incentivized by the pursuit for profits; there is currency tied to every decision-making process that occurs in the industry no matter how impactful or important it is. This is a critical distinction from prior times: the commercialization of decision-making has inseminated novel opportunities to disseminate any sort of speech—whether organic, commercial, or otherwise nefarious in nature—and inject it throughout the modern media ecosystem.

We have thusly moved on from the formative “public good” conceptualization at the inception of the world wide web; we are in the age now of the “commercial good”—explicitly, of the firms leading the industry. The media ecosystem of the twentieth century, in contrast, did not involve the commercialization of fine-grained information dissemination. This was perhaps true even in the early stages of the internet through the turn of the millennium. But now algorithmic developments including the deployment of sophisticated learning models by the most cash-rich firms in the world—alongside their

data-gathering practices and advantageous pseudo-monopolistic positions in a market with a paucity of true or would-be competitors—have collectively introduced a vicious situation by which commercial operators have the opportunity to initiate, advertise, and host a market for commercialized information dissemination in such a way that it is those willing to pay-to-play in this commercial regime who exclusively have the capacity to push information at the individual.

That is not to say that this power of decision-making was to an extent true of past instantiations of the American media ecosystem as well. The prior world dominated by broadcasting, radio, and print materials too had the capacity to produce and perpetuate bias. But there were some key differences. Their reach was not as granular or personalized because of the nature of the technology in question; a consumer internet laden with learning algorithms evolving and operating over corporate servers and producing results within milliseconds on the Search results page generates different impacts entirely. Furthermore, these more traditional past instances of the media ecosystem were heavily regulated either directly by the government or indirectly through combination of measures instituting industry-wide transparency and public accountability.¹⁴ Examples include federal election regulations for the broadcast and radio formats as well as journalistic standards across the news media. Thus, overall their capacity to engage in unfair practices leading to potential consumer harms was constantly policed. While they did nevertheless have tremendous power—these formats collectively constituted the media ecosystem—they experienced continual pressure and possessed limited capacity to perpetuate damaging impacts on the public.

Individual capacity to determine what we will see and be subject to has been holistically undermined and diminished by the consumer internet firms. Whereas the individual's consumption of information in decades past was one of open space or human thought it has now been invaded by a silent form of commercial speech in that the content displayed before us at the call of the firm responsible for populating the results page. Each time we open the laptop or checks the phone and utilizes the services central to information consumption today we are subjected to an array of information preselected and ordered for us at the determination of a mercenary machine that works for the profit of Facebook or Google, with nothing else trained into its decision modeling besides profit maximization. Scholars contend that human minds were not meant to deal with this kind of ease: instead we are biologically trained to see a wide unlabeled array of content and contend with its merits and demerits to the end of deciding for ourselves whether we shall support and take up the opinion-driven arguments or objective information contained therein.¹⁵

This is Thoreau's civil disobedience; but the Twitter feed has subverted the very concept of civil disobedience and subjugated the human interest to such an extent that it is not only our democratic processes and progress but our moral humanity itself that is currently under direct and immediate threat perpetrated by the consumer internet's business model. It is that which is in the crosshairs of the modern commercialized information dissemination system in America.

To examine this conundrum from a different angle, it is the 'third layer' of the infrastructure of the media system that has now been radically industrialized. The other two—the first being the physical network infrastructure and the second the content—already were in decades past. The third is the content dissemination network—but it could be said that the third layer of the infrastructure never should have been a free market in the way it is now in the first place. Leaving aside whether and how much the first two layers should have been opened to the industry at all and inspecting only the third, we can note that the industrialization of the dissemination layer clearly subverts the consumer's interest if left to the free market, given the observable negative externalities including the perpetuation of the disinformation problem and the wide spread of hate speech over these platforms.

Nissenbaum argues the approach to consumer privacy protection undertaken by the Federal Trade Commission and Department of Commerce is dangerous, noting that the U.S. government's "interest has been limited...by a focus on protecting privacy online as, predominantly, a matter of protecting consumers online and protecting commercial information: that is, protecting personal information in commercial online transactions. Neither agency has explicitly acknowledged the vast landscape of activity lying outside the commercial domain."¹⁶ Nissenbaum's reference is to the manner in which U.S. governmental agencies focus not on privacy concerns at large as and when they occur across society including governmental agencies and regulated entities like hospitals and banks, but rather only on those occasions when the data transfer affects "consumers"—those individuals party to some monetary transaction in the marketplace. Based on the discussion above we can extend Nissenbaum's point to the lack of effective oversight over the commercialization of decision-making—precisely because the narrow and independently minor decisions made using the classifier models developed by learning algorithms do not necessarily have dollars attached to them. But they are nonetheless designed in such a way as to yield the greatest possible profit margin for the service operator—and even perpetuate provably discriminatory decisions against individuals and classes of individuals so long as doing so remains non-transparent to the public and is aligned with the profit motives of the platform firm.

To that end, the collection of personal information is ubiquitous and its transfer amongst firms involved in the digital media ecosystem multidirectional. Indeed, the *modus operandi* of leading internet firms is to at once be at the center of and reach its tentacles throughout the commercial information sharing network stretching across the digital ecosystem. Firms like Google accordingly utilize a multitude of technologies and technological protocols to collect personal data, including over its own platforms, as well as through web cookies and physical equipment technologies deployed throughout the world. Critically this information is maintained by the firm and others like it within the company's walled gardens—its proprietary systems so that Google can maintain hegemony over the knowledge of the customer's individual profile for content-targeting purposes. Further, the firm "leases" the information out in anonymized formats—enabling advertisers to target certain classes of the population at will. Sometimes, the advertiser might inject its own data into Google's advertising platform, encouraging Google to help it reach audience segments to a remarkable degree of precision. This bidirectional relationship is critical to the functionality of the consumer internet—and operates as the grease at the joints of an industry enabling the aforementioned radical commercialization.

Bias in the Consumer Internet

The commercialization of decision-making in the consumer internet plays out in various ways potentially detrimental to marginalized groups including protected classes of the population. When the markets elevate currency over values the resulting economic logic tends toward enabling the pursuit of highest profit margin at the expense of any other concern, particularly if it is an unpoliced one extant over a largely unregulated market. Machine learning is the tool that enables the collation and exploitation of information, thus reducing transaction costs even further—with the profits generated thereof typically being drawn up by the industrial entities responsible for implementing the learning models in integrated manner.

Indeed the internet is effective as a means for communication—to the extent it is now humanity's social medium of choice—because it reduces costs of transaction in the exchange of information relative to the communication media of the past which typically did not enable personalization of rendered services nor collection of information on the consumer in the first place; the internet thus enables a two-sided exchange in a manner we had no capacity for in years past.¹⁷

But it is precisely this reduction of transaction costs that has enabled discriminatory outcomes that disfavor marginalized communities, particularly in the United States where the internet is in such wide use, the internet industry has such tremendous political power, and our

demographic heterogeneity and national political economic tradition and trajectory are such that the capacity for internet-enabled discrimination has been supercharged.

In this part we discuss a non-exhaustive set of common practices in and features of the internet industry that illustrate its capacity for discrimination though they have nevertheless reduced transaction costs for individual consumers.

Targeted Advertising Platforms

The creation of the commercial regime underpinning the consumer internet economy—targeted advertising—has enabled both intentional and unintentional discriminatory outcomes. Typically, ad targeting regimes take advantage of the commercial interests of two types of parties: the advertisers that wish to communicate their products and services to consumers and persuade purchasing decisions as possible; and the platforms and publishers that have access to consumer attention and therefore own ad space.¹⁸

Usually, platforms also possess and analyze large, refined stores of information on consumers. The raw data might include data collected about the consumer's "on-platform" activity including what products, social posts and search results displayed on the platform in question the consumer interacts with; the consumer's "off-platform" web activity pertaining to activity on third-party websites, including mouse clicks, browsing pathways, and content consumed; location information shared with the platform via the consumer's smartphone should the consumer have opted into location sharing with the platform service (or through other means in certain cases¹⁹); location and behavioral data collected through other device technologies such as beacons and routers that interact with the consumer's devices in the physical world; data purchased from or voluntarily shared by third parties such as data brokers and advertisers; and many others.

Advertising platforms—including those implemented and hosted by Facebook, Google, and Twitter—take advantage of such data collection regimes to infer behavioral advertising profiles on each user participating on the company's internet-based services. Those behavioral profiles are maintained by the platform firms and largely remain non-transparent to third parties. But should advertisers such as apparel designers and retail banks wish to target certain audience segments—young people of a certain income in Manhattan and San Francisco, for example—the platform firm typically analyzes its data stores, and determines which grouping of consumers in the target geography would be most likely to purchase the advertiser's wares. It is this determination of who should go into the targeted audience segment that

clearly has the capacity to engender harmful disparate impact. A recent suit put forth by the U.S. Department of Housing and Urban Development illustrates this tension clearly: the used its authority under the Fair Housing Act to allege that Facebook enables harmful disparate impact in making available housing opportunities because of the way that advertisers can target certain groups according to their membership in various consumer classes—including protected classes such as race and gender.²⁰

Perhaps most dangerously, civil rights laws in the United States only cover certain key areas that are absolutely critical to maintaining a modicum of economic fairness—including in housing and employment. Unfortunately, such protections against a commercial operator enabling disparate impact in the majority of other areas does not necessarily trigger a civil rights violation despite the clear discriminatory outcomes that can arise from only certain marginalized communities being subject to shady scams or, conversely, more mainstream communities exclusively being pushed very favorable ads such as investment opportunities.

Meaningful Social Interaction

Consumer internet firms deal in a novel form of currency: the collective combination of the user population's personal information and attention. By raking as much of this as possible and amassing it to generate collated ad space that can be sold off to the highest bidder via intelligent auctions for the purpose of enabling targeted commercial speech, the internet companies maximize their value proposition to businesses that wish to advertise back at the consumer. It is a vicious cycle that takes advantage of other efficiencies as well—in particular, the need to continually engage users such that they spend as much time on the platform as possible and furthermore engage with it to the greatest possible extent.

In 2018, Facebook chief executive Mark Zuckerberg proclaimed that his company would institute new changes to the algorithm driving the social media network's core News Feed service that ranks the universe of content available to a given user in the home screen; he noted that the company would now focus on promoting "meaningful social interactions."²¹ That is not to say that this was not always in the company's designs: he discussed how recent events had illustrated more clearly that there was too much passive interaction with content, particularly posts shared by "businesses, brands, and media."

What does meaningful social interaction really entail? Conveniently for Facebook, it is a metric that if effectively maximized can contribute to the two resources it principally cares about: the consumer's attention and personal information. Effective meaningful social interaction would keep users on the platform because

if done right it would connect users to more personal social content that they actually want to see—and if they engage more with such content then Facebook will know it and thereby know the user better such that ads can be disseminated more efficiently at them and ad space can be increased.

This is where the power of commercial machine learning—and resulting machine bias—come in. There is no scientific way to determine what types of content matter for an individual user; it is nigh impossible for a machine to infer precisely what the individual consumer truly cares about. Only broad inferences can be drawn—but it might be more difficult to infer what academic subjects and scholars resonate for an individual, or which particular players on a team he or she likes, or which shade of blue he or she likes the most. This is the fallacy of data, and by extension, learning models; it is used to estimate the real feature but cannot ever offer a precise representation of the real world, and yet it is readily used to make determinations about what the individual actually cares about in the real world. Thus, the leading consumer internet companies' quest toward enabling meaningful interaction—whether in the context of a search engine or e-commerce platform or social media network—is flawed at best.

The industry's use of highly sophisticated artificial intelligence systems including neural networks for real time analysis of user behaviors—in conjunction with social science research conducted within the industry itself—powers the refinement of the models used to rank such features as the News Feed. But regrettably, such systems have the propensity to supercharge the deployment of assessments about the individual in ways that implicate the individual's interest. If a user does not interact with some mundane piece of content because it does not personally resonate at a social or intellectual level, the platform must reorient its assessments about that individual user. It is this dynamic that has led Facebook down the path of grouping individuals by political allegiances and which has caused YouTube to be unable to screen certain inappropriate videos.

When looking through the lens of discriminatory practices, the platforms are designed to necessarily make assumptions about the nature of the individual based on the individual's demographic profile—including protected classes such as race and gender but also more precise ones including interests in certain forms of ethnic culture, music, and other instances of intellectual content. This is an online commercial landscape in which disparate impacts can run riot—where only certain marginalized classes are shown (or not shown) certain forms of content. And even if the content does not trigger civil rights protection in the United States, there are other manners in which it might damage the economic prospects of the individual. If Facebook decides that

an individual is likely more interested in basketball than microeconomics, for example, it might be the case that that individual is never subjected to content that would encourage better practices around personal finance, better awareness of the political state of the nation, and better awareness of broader economic opportunities that might be available should the user know where to look for them.

Whether such ranking models are fair or not all depends in the end on the design of the algorithm that maximizes so-called meaningful social interaction, defined and algorithmically trained to service the commercial objectives of the platform operator.

The Initial Pursuit of High-Value Customer Audiences

There is a tradition in Silicon Valley, particularly in the consumer internet industry, whereby fledgling firms tend to serve those niches that are already well-off first; should they be able to prove the efficacy of the business by serving those high-value customer segments then they might receive investment funds to tackle broader growth as well. Indeed, companies leading the sector have variously been party to such practices: Facebook first invited only Harvard students to participate on the network²²; Airbnb initially served only those cities where real time hotel prices were high²³; and Gmail's beta version was distributed first to a few hundred opinion leaders and those friends they wished to invite to use the service as well.²⁴

Needless to say, such communities—namely, elite universities, would-be hotel patrons in rich cities, and public intellectuals—are not representative of any community beyond the elite and tend to deprioritize or exclude marginalized communities that are most often subject to harmful discrimination. Nonetheless, it is through the observation of these initial groups' interactions with the platforms that computer engineers attempt to design the form their platforms will take at steady state. This culture of serving the privileged first and rolling out consumer products to the rest of society should the product gain in popularity is seemingly part and parcel of the investment culture that bleeds through the venture capital industry.

But it is a culture that considers the desires of lower socioeconomic classes last. And when overlaying the development and refinement of learning models over this conundrum, the potential for machine bias leading to disparate impact becomes resoundingly clear. The argument could be made that the platform firms protect against this potential harm in various ways—for instance by protecting against in-built propensities for learning models to perpetuate biased outcomes—but the fact remains that the design of the platforms necessarily must favor the elite and wealthy first. In a capitalistic regime favoring free markets no other approach would be viable

for venture capitalists and founders; if they do not take advantage of the economic opportunity of serving the well-off first then the competition will eventually do so and overtake them. In fact it could be said that if at any point a platform such as Facebook were to lose high-value users to the competition then the company would have to either acquire those competitors or reorient the ways in which the fundamentally platform works so as to increase the probability that high-value users might come to the platform. Indeed, this is exactly the strategic circumstance Facebook finds itself in now as it considers how to reclaim the high customer lifetime values at hand with respect to the young users who opt for non-Facebook internet-based services.

Public Policy Interventions to Counter the Spread of Machine Bias

Experts contend that model designers can protect against bias through development of technologies that check the representative nature of the training data and fairness of the outcomes. But while one absolutely can engineer such technological solutions to counter the overreaches of learning models, what forces companies to be fair when it is in their commercial interests to discriminate, even unfairly so, as long as the discrimination is not illegal? I would thus suggest a slightly different remedy: implementation of such technologies by the industry backed up by accountability forced on the industry through smart and earnest governmental regulation.

Machine learning technologies have come to the fore because of their tremendous efficiency. No longer do we require humans to monitor traffic systems to infer areas of congestion and manage the network; Uber, Waze and Lyft can accomplish the task much more effectively on an algorithmic basis. No longer do we need news editors to determine what information should or should not go front and center before our individual attention; Facebook, Apple and Twitter can infer who we are, what we want to see, and route the relevant content to us. No longer do we need to ask the contracting expert what flooring suits our apartment the best; Amazon will find out for us and assure it arrives post haste. And no longer do we need to rely on the guidance counselor to help decide where to apply and what college to ultimately attend; Google can address all our concerns.

As machine learning algorithms and artificial intelligence system become more ingrained in our daily lives and influence our behaviors throughout the day, so too does humanity necessarily become increasingly dialogical with the machine underpinning the consumer internet. Society's observable actions and behaviors are actively feeding the decisions executed by the machine that sits quietly behind the internet, and the corresponding commercially-driven decisions in turn influence our

actions and outlooks in the real world. Beyond the obvious questions this conundrum presents in regard to individual autonomy, psychological dependence, mental health, and the broader concern of empowering a civilization-wide overdependence on machine technologies and implicit bias against sentient real-world expertise, however, is the apparent reality: machines are discriminatory by design. Indeed, the more discriminatory they can be—the more incisive their predictions about individual behaviors and the collective outlooks of population classes—the more they add to the industry's pocketbook. This is ad targeting and content curation 101: if a machine can understand your mind, it is doing the job Facebook designed it to do. But in the course of so doing, the machine is bound to make frequent mistakes; there is no real-time learning system that can effectively model the human psychology without making mistakes along the way—and it is that in noise that pervades the system where harmful bias lurks.

There is no reasonable solution, then, but to utilize the full agency of the public interest to intervene and clarify for commercial entities what is right and what is wrong. Without making such rules of the road explicit, it is in the industry's interest to breach the public interest so long as it is legal to do so and unintelligible to the public; if Instagram leaves such opportunities on the table, Snapchat will pounce—and vice versa. They both thus have to take such opportunities up unless the consumer market reproaches them through expression of collective sentiment in the marketplace or unless the government intervenes. And consumer outrage expressed through purchasing behaviors will take too long or have minimal long-term impact in a space that offers little transparency. We need look no further than the voluntary reforms instituted by Facebook since the Cambridge revelations; while it has ceased certain activities, firms not under the public eye have taken them up, taking advantage commercial zones of operation cast aside by Facebook on the back of vociferous public advocacy.

We can conversely inspect the industry's actions in the face of governmental inquiries. If the industry earnestly wished to protect against these harms, why would it not wish to submit to governmental review and sectoral oversight? It is a problem of the interests of private commerce versus the interests of human rights. The culture engendered by the Facebook cultural insignia "move fast and break things" necessarily implicates machine bias and other challenges wrought by the radical capitalism seen in this industry. The industry's tendency is accordingly not to take on challenges presented by algorithmic design earnestly until it becomes popular to do so—by which time it is too late; the consumer internet industry's systems may have by that time contributed inordinately to systemic bias, prominent as it is in the American media and information universe.

A novel approach for governmental intervention should include the following, offered in increasing order of political difficulty given the inevitable policy pushbacks each measure would face.

- Federally funded research into techniques to protect against algorithmic bias. Computer scientists have developed novel techniques to protect against machine bias in recent years.²⁵ But as a general matter these approaches are variously applicable only to certain types of models or are otherwise not always feasible because of cultural norms that dictate companies will fail to pause to question whether their models are fair before deployment, or because of other practical hang-ups in the sector. More robust research is needed to develop more industry-grade mechanisms to help protect against machine bias. Further research is also needed to develop greater understanding regarding the impact of computing machines on society, and what public policy measures should be taken to counter industrial overreaches and contain harms. A good start has come from the Defense Advanced Research Projects Agency's Explainable Artificial Intelligence program²⁶; the government should channel further resources to such pursuits.
- Federally endorsed multi-stakeholder standards development toward a guiding framework for ethical artificial intelligence. Mathematicians and computer scientists have variously come together with ethicists and philosophers from across the industry, civil society, and academia to produce a slew of ethical codes for artificial intelligence and machine learning in recent years. The fact these discussions exist is a positive development. But we must assure history does not regard them as fluff. One commonly cited framework²⁷ highlights five key principles: "responsibility" such that those with grievances in regard to an algorithmic outcome have redress with a designated party; "explainability" so that the algorithms and data used to develop them can easily be explained to the public or those subject to their decision-making; "accuracy" so that the model's errors can be identified and proactively addressed; "auditability" so that third parties including public interest agents can investigate the algorithms and assure their integrity; and "fairness" so that the models do not perpetuate biased outcomes. This represents a start to developing a comprehensive set of principles on the governance and execution of fair machine learning models. The government should work with these and more stakeholders to coordinate a multi-stakeholder conversation concerning the development of an ethical framework for artificial intelligence. These conversations should be focused on the particular issue of the technological nature of the algorithms and data inputs themselves—leaving other important but less relevant contemporary conversations regarding the technology industry and the governance of artificial intelligence to the side. The government can use the National Institute of Standard and Technology's Cybersecurity Framework developed under the auspices of the Obama administration as a blueprint for how such multi-stakeholder guidance can come together.²⁸ Of particular importance throughout the process will be the assurance that public interest advocates are represented.²⁹ Such conversations can focus on the industrial use of artificial intelligence over the internet to maintain group focus while also addressing internet algorithms' outside influence over the information ecosystem.
- Industrial auditing and oversight of high-impact commercial internet algorithms backed by governmental enforcement to assure fairness. Consumer internet firms extensively implement machine learning models to drive growth, engagement, behavior profiling, and revenue collection and management—among many other activities. These have a tremendous impact on public interests from fairness to democratic process and should be subject to general governmental oversight in some capacity. A model like that settled by the Federal Trade Commission with Facebook³⁰ and Google³¹ through consent orders nearly ten years ago may be appropriate, whereby in response to industry overreaches the agency settled new conditions with each company, including the ongoing auditing of their practices with regard to maintaining consumer privacy. This condition from the consent orders effectively enforced a sea change on the companies; it forced them to install what are now known throughout the industry as privacy program teams—staff that are charged by the company to work with product managers and engineers to understand every single proposed product innovation, including the most minor of features, and help the subject firm coordinate a cross-functional decision as to whether the proposed changes would harmfully implicate the user's privacy or not. The personnel in the privacy program teams interact with external professional auditing consultants who

verify the integrity of the privacy practices of the subject firm, and develop periodic reports shared with the federal regulators that help affirm that the subject firm's privacy program is effectively protecting users from privacy overreaches. It could be argued that, in light of the reality that PricewaterhouseCoopers failed to find Facebook's missteps that were eventually revealed by Christopher Wylie, these types of setups are bound to fail.³² They can suffer, for example, from the traditional auditor's paradox, by which the auditing firm becomes close and collegial with the subject firm and fails in its role as an independent review agency working in earnest for the public interest. Culturally there can be a lack of incentive to report concerns accurately, largely because sharp criticisms will be seen by the subject firm. This is where the government can come in by holding all parties accountable. As the U.S. government pursues actions against the internet industry on the basis of further breaches of privacy, security, public trust, and algorithmic integrity, it should consider mechanisms to additionally force the companies to work with independent external auditors to assure internet-based artificial intelligence systems are not implicating public interests.

- *Radical data and algorithmic transparency for the public.* Centrally responsible for the exceedances of the algorithms underlying the consumer internet architecture is the lack of transparency into how they are developed, how they operate, and what they accomplish. Consumer transparency into this regime—through consumer understanding of what data corporate actors hold on them, how behavioral profiles are developed through inference, how machine learning models are used to develop such features as the News Feed and YouTube recommendation algorithms, and what the practical outcomes of these algorithms are—is critical to limiting the harmful discriminatory effects of the internet platforms. Indeed, many have attempted to develop tools to layer such transparency over the sector—including the political ad transparency projects led by ProPublica, Mozilla, and Who Targets Me, which were all stopped by Facebook in early 2019.³³ That Facebook was so determined to block the aforementioned services by tweaking its code is illustrative of the tension at the heart of true transparency measures: transparency breaks the impetus of the business model of internet companies like Facebook. These companies want to protect information pertaining to how

their curation and targeting algorithms work for two primary reasons. First, helping the public peer into the targeting metrics pertaining to Facebook ad campaigns can shine a much-needed light on how this company's algorithms perpetuate bias including by feeding the filter bubble problem, stoking hateful conduct online and offline, and enforcing damaging disparate treatment and impact in areas including politics and media exposure. Second, exposing the design behind algorithms enables the company's competitors to understand important strategic elements of the commercial makeup of Facebook and adjust their strategies in real time to challenge the company's strength in the market. In other words, it is all a form of commercial protectionism. I suggest a novel regime—a radical form of transparency as I have discussed with colleagues in related work—that can truly hold the industry accountable for the negative effects pushed by its models onto the public. Such transparency would enable users—or at the least, governmental or nonprofit organizations working in the public interest—to see what inputs go into the development of algorithms developed in the internet industry, and what outcomes those algorithms produce.

In addition to these proposals, reforms concerning privacy and competition policy are much-needed and should be pursued as well. I discuss with colleagues what form such reforms could take in related work.³⁴

Conclusions: An Ethical Approach in the Way Forward

The tide of public sentiment is closing on Silicon Valley internet firms. Over the past year, the Cambridge Analytica revelations, frequent disclosures about privacy and security breaches, and historic regulatory fines have demonized the sector and turned our attention toward Facebook, Google, and Amazon.

What distinguishes the consumer internet sector is that it is not subject to a rigorous regulatory regime like the telecommunications, healthcare, or financial industries are; the operation of online digital services over a physical infrastructure is still largely a novel practice as far as the laws are concerned, and the U.S. Congress has not yet acted. In this mostly regulation-less environment, these firms have had the opportunity to grow profits toward the combination of business practices that most effectively yields highest margins—in just the way that Karl Marx suggests capitalists would. These companies have, in the view of many scholars, subjugated the national public interest. The disinformation problem; the spread of hate speech; the persistence of extremist content; and the present concern of algorithmically-charged outcomes that perpetuate harmful bias: these negative externalities

are the symptoms of the commercial regime that sits behind the consumer internet, a silent machine that is designed algorithmically only to seek the highest possible profit without consideration of the public interest.³⁵

Centrally concerning is the currency these firms deal in and the opaque mechanisms by which they rake it. Some industry executives suggest the services they offer are “free”—a misleading conjecture. True, consumers do not pay monetary fees for their services, but the most effective consumer internet firms develop as two-sided platforms that amalgamate a complex combination of user attention and data on the end-consumer side of the market, and translate it through an automated digital advertising exchange into monetary reward in the advertising market. Further, these firms have inordinate market power in the end-consumer market; Facebook for example has near-monopolies in traditional social media and internet-based text messaging, Amazon has a near-monopoly in e-commerce, and Google has near-monopolies in online video, email, and search. Thus, these firms are able to hoover currency in the form of attention and personal data on one side of the market and charge monopoly rents for it on the other side of the market.

This hegemony over the market has been shown to tread over the public interest. The industry’s disincentive in protecting the public from such negative externalities as the disinformation problem is a mere symptom of its unwillingness to bend this highly profitable business model along with its use of social power to protect the business model from regulation through influence over policymakers. What the public now needs is a novel regulatory regime that can effectively rebalance the distribution of power between the industry, government, and citizen—a digital social contract. Jean-Jacques Rousseau suggested the danger of radical property rights—such as those that the capitalistic Silicon Valley now has over the individual’s attention and personal data—when he noted we must “beware of listening to [the first man to claim property rights]. You are lost if you forget that the fruits of the earth belong to us all, and the earth itself to no one.”³⁶

I do not contend that we should abolish the industry’s ownership of intellectual property in the consumer internet industry altogether—nor that we compromise the targeted advertising business model entirely. I would rather suggest development of a regulatory response that effectively responds to the capitalistic overreaches of the business model that sits behind the consumer internet. And this must include measures that can effectively hold the industry’s artificial intelligence platforms accountable, including through transparency that would enable public visibility into the darker effects of learning models implemented by the industry that systematically make decisions that are not in the interests of the individual.

Free market capitalism is the principal hallmark of the American approach to national economic design, but the government has never hesitated to strike down the market when its practices have implicated the nation’s commitment to democracy. This is the very situation we now find ourselves in with respect to the internet.

¹ Robin Mansell, New visions, old practices: Policy and regulation in the Internet era, *Journal of Media & Cultural Studies*, Volume 25, Issue 1, 2011.

² Cecilia Kang and Sheera Frenkel, Facebook Says Cambridge Analytica Harvested Data of Up to 87 Million Users, *The New York Times*, April 4, 2018.

³ S.1989, Honest Ads Act, 115th Congress, 2017-2018.

⁴ See e.g. Solon Barocas and Andrew Selbst, Big Data’s Disparate Impact, *California Law Review*, Vol. 104, No. 3, June 2016.

⁵ Yoshua Bengio and Yann LeCun, Scaling Learning Algorithms towards AI, in *Large-Scale Kernel Machines*, MIT Press, 2007.

⁶ Amanda Sakuma, The bogus “Momo challenge” internet hoax, explained, *Vox*, Mar 3, 2019.

⁷ Ian Sample, Study blames YouTube for rise in number of Flat Earthers, *The Guardian*, Feb 17, 2019.

⁸ Steven Willborn, The Disparate Impact Model of Discrimination: Theory and Limits, *American University Law Review*, Vol. 34, No. 3, Spring 1985.

⁹ Susan Grover, The Business Necessity Defense in Disparate Impact Discrimination Cases, *Georgia Law Review*, Vol. 30, No. 2, Winter 1996.

¹⁰ Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights, *Executive Office of the President*, May 2016.

¹¹ When Digital Dust Is Gathered, Constellation May Be Muddled, *National Public Radio*, April 13, 2013.

¹² Facebook agrees to sweeping reforms to curb discriminatory ad targeting practices, *American Civil Liberties Union*, Mar 19, 2019.

¹³ Thomas Eisenmann, Internet companies’ growth strategies: determinants of investment intensity and long-term performance, *Strategic Management Journal*, Oct 30, 2006.

¹⁴ See e.g. Advertising and disclaimers, Federal Election Commission.

¹⁵ Cass Sunstein, #Republic, Divided Democracy in the Age of Social Media, *Princeton University Press*, 2017.

¹⁶ Helen Nissenbaum, A Contextual Approach to Privacy Online, *Daedalus*, Fall 2011.

¹⁷ B. Mahadevan, Business Models for Internet-Based E-Commerce: An Anatomy, *California Management Review*, July 1, 2000.

¹⁸ Robbin Lee Zeff and Bradley Aronson, *Book Advertising on the Internet*, John Wiley & Sons, 1999.

¹⁹ Joseph Cox, I Gave a Bounty Hunter \$300. Then He Located Our Phone, *Vice*, Jan 8 2019.

²⁰ The Secretary, Department of Housing and Urban Development, on behalf of Complainant Assistant Secretary for Fair Housing and Equal Opportunity v. Facebook, Charge of Discrimination, U.S. Department of Housing and Urban Development.

²¹ Adam Mosseri, Bringing People Closer Together, *Facebook Newsroom*, Jan 11, 2018.

²² Alexis Madrigal, Before It Conquered the World, Facebook Conquered Harvard, *The Atlantic*, Feb 4, 2019.

²³ Jessica Salter, Airbnb: The story behind the \$1.3bn room-letting website, *The Telegraph*, Sep 7, 2012.

²⁴ Chris Welch, Gmail is 10 years old today, *The Verge*, Apr 1, 2014.

²⁵ See e.g. Nicholas Diakopoulos, Accountability in Algorithmic Decision Making, *Communications of the ACM*, Feb 2016.

²⁶ David Gunning, Explainable Artificial Intelligence (XAI), Defense Advanced Research Projects Agency.

²⁷ Nicholas Diakopoulos et. al., Principles for Accountable Algorithms and a Social Impact Statement for Algorithms, *Fairness, Accountability, Transparency in Machine Learning*.

²⁸ Cybersecurity Framework, National Institute of Standards and Technology, U.S. Department of Commerce.

²⁹ Yochai Benkler, Don't let industry write the rules for AI, *Nature* 569, 161, 2019.

³⁰ Agreement Containing Consent Order, File No. 092 3184, *In the Matter of Facebook, Inc.*, Federal Trade Commission.

³¹ Agreement Containing Consent Order, File No. 122 3237, *In the Matter of Google, Inc.*, Federal Trade Commission.

³² PwC had cleared Facebook's privacy practices in leak period, *Reuters*, Apr 20, 2018.

³³ Jeremy Merrill and Ariana Tobin, Facebook Moves to Block Ad Transparency Tools — Including Ours, *ProPublica*, Jan 28, 2019.

³⁴ See e.g. Dipayan Ghosh and Ben Scott, Digital Deceit II: A Policy Agenda to Fight Disinformation on the Internet, *New America & Shorenstein Center on Media, Politics and Public Policy at the Harvard Kennedy School*, Sep 24, 2018; and Dipayan Ghosh, A New Digital Social Contract to Encourage Internet Competition, *Antitrust Chronicle, Competition Policy International*, Apr 2019.

³⁵ Joseph Nye, Protecting Democracy in an Era of Cyber Information War, *The Hoover Institution*, Nov 13, 2018.

³⁶ J. Rousseau, On the Origin of the Inequality of Mankind: The Second Part

Dipayan Ghosh is a Shorenstein Fellow and co-director of the Platform Accountability Project at the Harvard Kennedy School. He has served as an economic policy advisor in the White House during the Obama administration and a privacy and public policy advisor at Facebook.

Supporting Data

Figure 1. Consumer internet platforms engage in exchanges of personal information, content including personal posts and news, and dialogical feedback.




Figure 2. Commercial learning models designed to infer behavioral profiles and curate content are continually refined by consumer internet firms through feedback from real world routines and behaviors.





Figure 3. While learning models might efficiently design decision regimes for large populations, poorly designed systems may fail to detect that minority populations defined along protected class lines have a different nature, which can systemically perpetuate harmful discrimination.



The Migration Challenge

By James F. Hollifield, Southern Methodist University



Introduction

International migration has been steadily increasing in every region of the globe since the end of the Second World War. In 2017, approximately 258 million people reside outside of their country of birth (3.4 percent of the world's population) and over the past half century, individual mobility has increased at a steady pace. Tens of millions of people cross borders on a daily basis, which adds up to roughly two billion annually. International mobility is part of a broader trend of globalization, which includes trade in goods and services, investments and capital flows, greater ease of travel, and a veritable explosion of information. While trade and capital flows are the twin pillars of globalization, migration is the third leg of the stool on which the global economy rests.¹

Clearly, migration is a defining feature of the global era in which we live. It is in many ways connected to trade and investment, yet it is profoundly different. *People are not shirts*, which is another way of saying that labour is not a pure commodity. Unlike goods and capital, individuals can become actors on the international stage (they have agency) whether through peaceful transnational communities or violent terrorist/criminal networks. In the rare instances when migrants commit terrorist acts, migration and mobility can be a threat to the security of states. However, the benefits of migration far outweigh the costs. Immigrants bring labour, human capital, and new ideas and cultures (diversity) to their host societies; and in liberal democracies, they come with a basic package of (human and civil) rights that enables them to settle and become productive members of society, if not citizens of their adoptive countries. Conversely, they may return to their countries of origin where they can have a dramatic impact on economic and political development.²

Lest we forget, not all migration is voluntary—in any given year millions of people move to escape political violence, hunger, and deprivation, becoming refugees, asylum seekers, or internally displaced persons. In 2017 the number of “persons of concern” to the United Nations High Commissioner for Refugees (UNHCR) was 65.6 million, including 22.5 million refugees, 2.8 million asylum seekers, and 40.3 million internally displaced people. Wars in the Middle East (especially Syria and Iraq), East and West Africa, and instability in South Asia and Central (Northern Triangle) and Latin America (Venezuela) continue to feed a growing population of forced migrants. Among

the most recent and fastest exoduses of people from their place of origin were the movement of Rohingyas from the Rakhine State in western Myanmar into neighbouring Bangladesh, and Venezuelans fleeing political instability and economic deprivation (almost four million to date). Europe (as in the European Union) and Germany, in particular, struggled to cope with waves of forced migration—almost 1 million asylum seekers arrived in Germany alone in 2015. In 2018-19, tens of thousands of Central Americans fled the Northern Triangle countries, most heading north to seek asylum in the United States. Because it is so complex and multi-faceted, migration of all types poses a challenge for individual states, for regions like the European Union (EU) and for the international community as a whole.³

Four factors drive migration policy—security, cultural and ideational concerns, economic interests, and rights. National security—the institutions of sovereignty and citizenship—economics (markets) and rights are all part of a multi-dimensional game in migration policymaking. In ‘normal’ times, the debate about immigration revolves around two poles: markets (numbers) and (status) rights, or how many immigrants to admit, with what skills, and what status? Should migrants be temporary (guest) workers, allowed to settle, bring their families, and get on a ‘path to citizenship?’ Is there a trade-off between rights and numbers (markets) as Martin Ruhs and others suggest?⁴ All good questions—but cultural concerns (where should the immigrants come from, which regions of the globe, with which ethnic characteristics, and issues of integration) often trump markets and rights, and the trade-offs are more intense in some periods and in some countries than in others.

With the September 11, 2001, terrorist attacks in the United States and again with the November 13, 2015, attacks in Paris, France, immigration and refugee policymaking has been dominated by a national security dynamic (with a deep cultural subtext, fear of Islam) and the concern that liberal migration policies pose a threat to the nation and to civil society. In times of war and political crises, the dynamic of markets and rights gives way to a culture-security dynamic and finding equilibrium (compromise) in the policy game is much harder—this is the policy dilemma facing leaders in the United States and across the globe.

The four-sided game (see Figure 1) is difficult at the national, state and local levels, and it is rendered more complex by virtue of the fact that migration control has important foreign policy implications. The movement of populations affects international relations and security in myriad ways. Hence, political leaders are always engaged in a two- or even three-level game,⁵ seeking to build local and domestic coalitions to maximize support for policy but with an eye on the foreign policy consequences.

The Historical Context: *E Pluribus Unum*

In 2004 one of the most respected political scientists of his generation, Samuel P. Huntington of the “clash of civilizations” fame, published what would be his final major work, a book entitled *Who Are We? The Challenges to America’s National Identity*. In this book, Huntington argued that American national identity, and by extension, U.S. national interests are threatened by a growing wave of Hispanic immigration. He asserted that Mexican immigrants in particular are engaged in “*la reconquista*” or a re-taking of territory lost during the Mexican-American War—not through military conquest but through a peaceful “invasion,” the result of which has been to undermine Anglo-Protestant (Puritan) values of hard work, loyalty to the “founding principles” of the U.S. Constitution, and rule of law. Huntington decried the wave of undocumented immigration of impoverished and poorly educated Mexicans and Central Americans, the rise of dual nationality, bilingualism, and what he sees as the loss of a clear national identity and purpose—all the result of too much immigration. He begins the book by outlining three waves of immigration in U.S. history, first in the mid-19th century with the Irish and Germans and continuing through the late 20th and early 21st centuries with Hispanics and Asians. He omits the first wave of immigration from the British Isles in the 17th and 18th centuries, because he views this wave as a period of settlement and founding during which the new American nation was forged with a fundamentally Anglo-Puritan outlook.

Huntington’s argument underscores the enduring controversy over immigration as a force shaping and reshaping American society. His critics accuse him of being a latter-day nativist and “know nothing,”⁶ echoing the controversy in earlier periods of American history when immigration was seen as a threat to basic “American” values. In the 18th century, for example, Benjamin Franklin was very concerned about German immigration in Pennsylvania, because he thought that uneducated German peasants who were coming from a semi-feudal society had little understanding of what it was like to live in a Republic based on rule of law and individual liberty. Later in his political career Franklin changed his views on German immigration, as German-Americans became an increasingly important part of the

electorate in Pennsylvania; and less than two centuries later a descendant of those German immigrants, Dwight Eisenhower, was elected President of the United States. It is important to keep in mind that immigration from the colonial period through the Civil War and Reconstruction (roughly the first hundred years of U.S. history) was controlled by the individual states, to the extent that it was regulated at all. Immigration was largely driven by the demand for labor to fuel the fires of industrialization—as in later periods private employers were instrumental in recruiting immigrants—by westward expansion, and by a seemingly unlimited supply of labor displaced by the industrial revolution and the concomitant rural exodus in Western Europe.

In *The American Kaleidoscope: Race, Ethnicity, and the Civic Culture* (1990) Lawrence H. Fuchs argued that three ideas have dominated the American approach to immigration and citizenship. They are the Massachusetts and Virginia models, dating from the early colonial period, and the Pennsylvania model, which took shape in the early years of the Republic. He admits that these are ‘ideal types,’ but he contends that traces of each model can still be found in contemporary debates. The **Massachusetts model** most closely conforms to Samuel Huntington’s ideal of Anglo-Puritanism (what might be called a White Anglo-Saxon Protestant or WASP view of American national identity). In this view immigrants are to be welcomed if they are willing to assimilate, learn English, and adopt the dominant religion and culture. In colonial Massachusetts that meant conformity to ascetic Puritan ideals, and Samuel Huntington, who lived on Beacon Hill and spent most of his career at Harvard, clearly wanted to make respect for Anglo-Puritan values the basis for selecting and naturalizing immigrants.

The **Virginia model** revolves around the demand for labor. In the early colonial period (17th and 18th centuries) the Virginia and Carolina planters needed stoop labor to pick tobacco and cotton. They acquired this labor initially through coercion—the impressment of Native Americans and the enslavement of Africans brought to the New World in bondage. Since both groups were considered to be sub-human, no thought was given to their naturalization and assimilation. Indeed, many Europeans were brought to work on the plantations and in shops and factories as indentured servants with limited rights. We hear echoes of the Virginia model in contemporary debates about guest worker programs whereby foreigners are brought to the United States as bonded labor on a temporary basis with no right to settle or naturalize.

Finally, the **Pennsylvania model**, which Fuchs sees prevailing in the Nationality Act of 1790 establishing a uniform rule of naturalization, calls for equal treatment of newcomers, welcoming them to settle, live, and worship as they see fit so long as they respect the law and the basic values of the Republic. It should be noted that the

same act limited naturalization to 'free white persons of good moral character,' thus enshrining race (and to some extent class as it excluded indentured servants) into U.S. immigration law. Still the first President of the Republic, George Washington, reflected an expansive ideal of citizenship when he said, "the bosom of America is open to receive not only the Opulent and respectable Stranger, but the oppressed and persecuted of all Nations and Religions; whom we shall welcome to a participation of all our rights and privileges."

The Pennsylvania model was reinforced after the Civil War with the ratification in 1868 of the Fourteenth Amendment to the Constitution, which extended citizenship to "all persons born or naturalized in the United States."⁷ The Amendment was intended primarily to overturn the *Dred Scott* decision of the Supreme Court (1857) and to grant citizenship to former slaves, but in the process, it codified birthright citizenship with far-reaching implications for immigration policy. Barely two decades after the end of the Civil War, the Statue of Liberty—a gift from one fledgling Republic, France, to another, the United States—was erected in New York harbor (1886); and it would become the most visible symbol of an open and tolerant America, welcoming immigrants from the four corners of the globe. Inside the pedestal of the statue is inscribed the most famous immigration sonnet in American history, *The New Colossus*, by Emma Lazarus, which reads in part

"Give me your tired, your poor,

Your huddled masses yearning to breathe free,

The wretched refuse of your teeming shore.

Send these, the homeless, tempest-tost to me,

I lift my lamp beside the golden door!"

All three "models" have been present historically in debates over immigration and citizenship, which have followed the unofficial national motto, *e pluribus unum* (out of many, one). At times, Americans, like Samuel Huntington, have been more concerned about the "*unum*" and the need to maintain a clear national identity and purpose; at other times, Americans have hewed to the Pennsylvania model, showing a greater willingness to accept immigrants and celebrate diversity, the "*pluribus*."

As mentioned above, we can identify four waves of immigration in U.S. history (See Figure 2). The first from the British Isles before 1820 was made up largely of the English and Scots who came for a variety of religious, political (many of the early English settlers were dissenters), and economic (the promise of land and a new start) reasons. The second wave beginning around 1840 and running through the Great Depression of the 1870s was more economic in nature (the Irish were fleeing starvation

and deprivation during the potato famine) while other northern and west European groups like the Germans and Scandinavians were mostly farmers and artisans, attracted by land in the vast expanse of the Great Plains. Because many of the newcomers were Roman Catholic, the second wave provoked an anti-Catholic backlash, which found its greatest expression in the "know nothing" movement of the mid-19th century. The third wave started in 1880 and continued to 1914, when the Great War brought an end to the transatlantic migrations. This wave proved even more controversial than previous waves, because it was ethnically diverse. Male Chinese laborers were brought into the west to build the transcontinental railroad and to work in the mines; southern and eastern Europeans flooded into eastern cities, and into the mid- and southwest, increasing the Catholic and Jewish populations in these regions. It was during the third wave that the federal government began to assert control over immigration, starting with the Chinese Exclusion Act of 1882, which, like Prop 187 over a century later, was the direct result of a nativist backlash in California against a rising tide of immigration and a seemingly complacent federal government (see Figure 2).

By the early 1900s, political pressure was again building to slow the rate of immigration. Congress set up the Dillingham Commission in 1907 to study "the problem" and to recommend new ways of selecting immigrants. The pendulum was swinging back in favor of those concerned about national identity (the *unum*), and the Commission report issued in 1911 concluded that the United States was threatened by the increasing number of immigrants from "non-traditional" source countries. The Commission called for literacy tests and—relying on the pseudo-science of eugenics widely accepted at the time—argued in favor of a racially-based immigration policy. The Commission concluded that immigrants from southern and eastern Europe had more "inborn socially inadequate qualities than northwestern Europeans."⁸

After World War I inflows of immigrants from Europe recovered briefly, but in 1921 Congress enacted the first quantitative restrictions on immigration and in 1924 passed the National Origins Quota Act, which restricted immigration to northern and western Europeans, essentially locking out all other nationalities. Inflows fell rapidly and the onset of the Great Depression in 1929 brought a halt to immigration. The foreign population was quite large in the interwar period, but immigration (inflows) would not start again until after World War II. The 1924 National Origins Quota Act, which established a racial hierarchy as the basis for immigrant selection, remained in effect until its repeal in 1965. During the turbulent decade of the 1930s through the Second World War avenues for legal immigration were restricted, and the United States had no official refugee policy. Refugee admissions were decided on an *ad hoc* basis, and many European Jewish

refugees fleeing Nazi persecution were turned away from American shores.

Notwithstanding the wave of nativism, racism and restrictionism in the 1920s, the American political landscape was transformed by the third wave of immigration. Attention shifted from stopping immigration to assimilating immigrants. This was the heyday of Tammany Hall⁹ and big-city political machines in places like New York, Boston, and Chicago where first the Irish, then the Italians, and eventually southern and eastern European Jewish immigrants would come to play a larger role in urban politics. The Democratic Party was the major beneficiary of the support of the newcomers and Franklin Roosevelt would forge a New Deal coalition between working class, largely Catholic and Jewish immigrants in the north, and poor whites in the Protestant south. Even though the muscle of the big city machines was not enough to overcome nativist politics in the interwar period, Americans found a new metaphor to describe the assimilation of immigrants: the *melting pot* was popularized in a play by Israel Zangwill, which premiered in 1908. The notion of immigrants from many different cultures melting into a new society became synonymous with immigration and the “American dream.” The protagonist in Zangwill’s play proclaims “Germans and Frenchmen, Irishmen and Englishmen, Jews and Russians—into the crucible with you all! God is making the American!”

All was not love and light in immigration politics following the third wave. In a dispute in 1930 with a Congressman from New York, Fiorella LaGuardia, President Herbert Hoover wrote in a letter to his fellow Republican, “the Italians are predominantly murderers and bootleggers [and you and your Italian supporters] should go back to where you belong [because] like a lot of other foreign spawn, you do not appreciate this country which supports you and tolerates you.”¹⁰ In the presidential election of 1928, Al Smith, the Irish Catholic Governor of New York and Democratic candidate, would lose to Republican Herbert Hoover, but by winning the Democratic Party nomination, Smith had broken an important cultural barrier, overcoming anti-Catholic and anti-immigrant sentiments. Thirty-two years later another Irish Catholic politician, the Democrat from Massachusetts, John Fitzgerald Kennedy, would overcome the final hurdles to the full participation of Catholics in American political life.

The Politics of Immigration Control

The first cracks in the National Origins Quota policy occurred during and immediately after World War II with the repeal of the Chinese Exclusion Act in 1943, the launch of the *Bracero* program in 1942, and the arrival after the war of large numbers of refugees and war brides from Europe and Asia. These groups did not fit within any of the existing quotas. China was an ally in the war against Japan, and Congress decided that the longstanding ban

on immigration and naturalization of Chinese nationals was bad for the war effort. Chinese immigrants living in the United States were allowed to naturalize, but strict quotas on Chinese immigration remained in effect. The United States in World War II was leading the fight against fascism and the racist ideology underpinning it. The contradictions of American immigration and refugee policy—not to mention segregation and Jim Crow—were increasingly anomalous and at odds with American foreign policy.

The War also brought new demands for foreign labor. The *Bracero* program was launched to fill gaps in the American labor market resulting from the draft. This guest worker program would have major long-term consequences for U.S. immigration policy. Following the ‘Virginia model,’ the program allowed for the recruitment of tens of thousands of temporary or guest workers from Mexico in the 1940s, first in agriculture and subsequently in the railroad and transportation sectors. It marked the beginning of large scale immigration from Mexico, which continued until the ‘great recession’ and financial crisis of 2007-08 when the flows reversed and net migration from Mexico turned negative, with more returns than arrivals. Attempts were made to reverse the flows with “Operation Wetback” in 1954 in which hundreds of thousands of Mexican workers and their families, including many who were U.S. citizens, were voluntarily repatriated or summarily deported to Mexico. The *Bracero* program remained in effect until its repeal in 1964 and the passage of the Immigration and Nationality Act of 1965.

Known as the Hart-Celler Act, the 1965 INA was a landmark piece of legislation, which repealed the National Origins Quota system, thus eliminating race and ethnicity—from the law on the books if not from the law in action—as the principal criteria for selecting immigrants. The pendulum of immigration politics was swinging back in favor of greater diversity (*pluribus*) and tolerance. The 1960s would see the triumph of the Pennsylvania model and the rise of what I have called ‘rights-based’ politics.¹¹ The relationship between individuals, groups, and the state was redefined through a process of political struggle (the civil rights movement) that would sweep away Jim Crow¹² and racial discrimination and, in the process, expand the rights of immigrant and ethnic minorities. A new type of *rights-based politics* was emerging at every level of the polity, from partisan and interest group politics, to the legislature and executive, and especially in the federal judiciary, which became increasingly active in protecting minority and civil rights. A similar trend in rights-based politics can be seen across the western world. Beginning in the 1960s, the courts would play an important role in immigration policymaking, restraining state and local authorities in their treatment of immigrants, helping to consolidate the rights of immigrants and minorities, and reasserting the plenary power doctrine.

The move away from the Massachusetts and Virginia models in favor of the Pennsylvania model of immigration and citizenship after World War II can be attributed to two political developments: the cold war and the civil rights movement.¹³ Public opinion remained hostile to immigrants and refugees in the 1950s. Congress passed the McCarran-Walter Act in 1952, which made it a felony to "harbor, transport, and conceal illegal immigrants;" but under the *Texas Proviso*, those employing 'illegal' immigrants were exempt from the law. Employers, particularly the growers in the southwest, had enough political clout to keep cheap Mexican labor flowing into the U.S. market. McCarran-Walter also loosened racial restrictions on immigration ever so slightly, but without repealing the National Origins Quota system. Reflecting the fear of communist subversion during the early years of the cold war, McCarran-Walter contained provisions for screening immigrants to catch communists and subversives, a move which was in keeping with McCarthyism and the new red scare. President Harry Truman vetoed the bill, calling it "un-American," but Congress overrode his veto. Congressional efforts to placate xenophobic and McCarthyite groups made it difficult for the President to ease restrictions on refugees coming from communist countries. Immigration and refugee policy were important foreign policy tools and the President needed a freer hand to accommodate cold war refugees in particular. Ultimately the civil rights movement, which had as its primary objective to overturn Jim Crow and achieve equal rights and full citizenship for African-Americans, swept away the last vestiges of the racist and discriminatory National Origins Quota system, leading to the most radical reform of immigration policy in American history. The INA of 1965 was passed on the heels of the 1964 Civil Rights Act and the Voting Rights Act of 1965. Immigrants were among the most important beneficiaries of the civil rights movement, as laws designed to end racial discrimination against blacks helped open up new political and legal spaces (rights) for immigrant minorities, setting the stage for the fourth (and largest in absolute numbers) wave of immigration in American history.

The fourth wave began slowly in the 1970s, in part because of the severe economic slowdown that was the result of two oil shocks and a steep recession in 1981-82. But as the economy recovered in the 1980s immigration accelerated rapidly, and by the first decades of the 21st century the foreign-born population (stocks) would climb to an all-time high of 44.5 million in 2017 (see Figure 2). The civil rights movement and the INA of 1965 had laid the political and legal groundwork for a more expansive immigration policy, but it was the soaring American economy in the 1980s and '90s that propelled immigration to new heights. The free market policies of the Reagan and Clinton administrations made the United States increasingly immigrant-friendly. Demand-pull forces in

the American labor market were strong and there was a relatively unlimited supply of workers in Mexico, Central America, and Asia ready to fill this demand.

Strange bedfellow coalitions of civil rights liberals (northern Democrats, many of them—like Senator Edward Kennedy of Massachusetts—descendants of the second and third wave immigrants) and business-oriented, Wall Street Republicans helped to pass some of the most expansive immigration laws in U.S. history. The Refugee Act of 1980 incorporated the 1951 UN Refugee Convention into U.S. law. During most of the cold war period, U.S. policy favored refugees fleeing persecution in communist countries, whereas the Geneva Convention defined a refugee as anyone with a "well-founded fear of persecution." Signatories of the Convention were bound by the principle of *non-refoulement*, whereby anyone who met the Geneva standard for asylum could not be returned to the country from which they were fleeing. The 1980 Refugee Act brought the United States in line with international law, giving new impetus to a more rights-based approach to immigration and refugee policy. With the winding down of the Cold War in the late 1980s and '90s, only Cuba retained its special status as a communist country from which refugees would be accepted with almost no questions asked. However, the Mariel boat lift at the end of the Carter presidency in 1980—in which Fidel Castro opened the Cuban port of Mariel to emigration allowing 125,000 Cubans to flee, including criminals and the mentally ill who were released from prisons and hospitals—forced the United States to rethink blanket asylum policies for Cubans.

In 1979 Congress set up the Select Commission on Immigration and Refugee Policy (SCIRP), chaired by Father Theodore Hesburgh and under the direction of Lawrence Fuchs—the first such commission since the Dillingham Commission. As the SCIRP went about its work in the early 1980s, holding hearings, gathering data, and conducting research, immigration soared—not only legal immigration, already opened up as a result of the 1965 INA, which made kinship and family ties the primary criterion for admission, but also 'illegal' immigration. The 1965 INA repealed the National Origins Quota system, creating avenues for immigration from non-traditional sources, particularly Latin America (Mexico), Asia, and eventually Africa and the Middle East. The INA also imposed numerical limits on the number of visas, including the first such limits on immigration from the Western Hemisphere (120,000 annually). These limitations would lead to a big imbalance between the demand for and supply of visas. Rather than waiting in long queues that could last years, many immigrants chose to come illegally, either slipping across land and sea borders or coming on a tourist visa and overstaying.

The majority of undocumented immigrants were (and are) visa 'overstayers,' i.e. individuals who entered the

country on a tourist visa and simply remained in the U.S., melding into society, and joining a growing black market for labor (see Figure 3). When the SCIRP made its recommendations to Congress early in the Reagan administration, illegal or undocumented migration was the biggest policy issue; and the foreign-born population, as a percentage of the total population, was rapidly approaching a historic high. By 2008 foreigners surpassed 14 percent of the total population—a level not seen since the beginning of the 20th century (see Figure 3).

Clearly, immigration was reshaping American society, and immigrants were coming to play an increasingly important role in the economy. Policy debates in the 1990's and 2000's would evolve along four lines: (1) Economic—what are the costs and benefits associated with high levels of immigration, especially 'illegal' or undocumented immigration? (2) Social—how are the newcomers and their children (the second generation) assimilating?¹⁴ Are they learning English and are they succeeding in the labor market? (3) Political—will the newcomers be good citizens? Will they participate in politics, and if so, how? Will they be Democrats or Republicans, liberals or conservatives? Will they constitute a "swing vote?" And (4) security—with the terrorist attacks of September 11, 2001 (9-11), immigration and refugee policy was in the spotlight. Border enforcement and screening of persons wishing to enter the United States took on a new urgency. How did the terrorists enter the country? Was the 9-11 attack the result of lax border enforcement and an overly liberal immigration and refugee policy? Debates about terrorism and national security intensified in the 2010s with attacks by jihadi groups in major European cities—Madrid in 2004, London in 2005, Paris in 2015, and Brussels in 2016—and with the election of Donald J. Trump as President of the United States. Yet, despite the security concerns, it proved difficult to roll back the rights of migrants. Sealing the border, building walls, summarily deporting large numbers of migrants (as happened during Operation Wetback in 1954), stopping family reunification, turning back refugees and asylum seekers, rolling back civil rights (due process and equal protection), and cutting immigrant access to social services, is not so easy. It turns out that in liberal democracies like the United States, rights have a long half-life and they are deeply institutionalized.

Congress attempted to regain control of immigration, especially **undocumented immigration**, in 1986 with the Immigration Reform and Control Act (IRCA). IRCA, known as the Simpson-Mazzoli Act. IRCA was the result of a compromise between "restrictionists," those who wanted to stop undocumented immigration, including Republicans led by Senator Alan Simpson of Wyoming and some southern Democrats, and "admissionists" those who wanted to legalize the large population of undocumented migrants by granting them amnesty, including northern liberal Democrats, led by Senator

Edward Kennedy of Massachusetts. In the end a rights-markets coalition formed in the Senate and the House, and a compromise was struck, allowing for the amnesty of undocumented migrants in exchange for sanctions (fines and imprisonment for repeat offenders) to be imposed on employers who knowingly hire undocumented migrants. The amnesty succeeded in bringing over 2.7 million people out of the shadows. To qualify for amnesty, the undocumented had to get certification that they were employed and that they had come to the United States prior to January 1, 1982. Critics of the amnesty argued that it created a **moral hazard**. More people would be willing to take the risk of immigrating illegally on the assumption that they would be amnestied at a later point in time. Opponents of more liberal policies argued that until the borders are secure there should be no expansion of rights for immigrants, legal, illegal or otherwise—this is the 'enforcement only' approach to immigration reform.

Employer sanctions, on the other hand, represented the first attempt by the federal government to pursue an internal control strategy, using labor laws to control immigration. IRCA created the I-9 form, which requires all persons seeking employment to present documentary evidence that they are legal residents. But out of concern that the new law could lead to discrimination against foreign-looking or foreign-sounding job applicants, provisions were inserted in IRCA to ensure that the rights of ethnic minorities would be protected—more evidence of the power of rights-based politics. Under IRCA, employers were not liable for hiring anyone who presented documents that "looked official," and they were not required to verify the authenticity of documents. This loophole made employer sanctions very weak, and it led to the creation of a new black market for false papers, especially social security cards and drivers licenses. Concerns for privacy and civil liberties prevented Congress from creating a national identification card, which is common in many other democracies. The American Civil Liberties Unions (ACLU) is strongly opposed to a national ID.

Agriculture posed a specific regulatory problem, because of the informality and seasonal nature of employment in this sector. In the run-up to the passage of IRCA, growers lobbied for a guest worker program (again visions of the Virginia model), but labor unions, especially the United Farm Workers of America (UFW), co-founded by the charismatic labor leader, César Chávez, opposed what they considered a system of bonded labor. The result was the creation of a Special Agricultural Worker (SAW) legalization program under which 750,000 mostly Mexican farm workers were amnestied. Finally, with respect to the impact of IRCA on overall levels of immigration, it is important to remember that each person covered by the amnesty was able to bring relatives (spouses, parents, brothers, and sisters) into the United States under the family reunification provisions of the 1965 INA (see Figure 4.)

The IRCA did little to slow the pace of undocumented immigration into the U.S. (see Figure 4). Over the course of the 1990s and into the first decade of the 21st century, undocumented would come to rival documented immigration, setting the stage for a backlash against immigrants and a series of anti-immigration laws. First came Prop 187 in California (1994), then the Illegal Immigration Reform and Immigrant Responsibility Act or IIRAIRA (1996), leading to the Sensenbrenner Bill (2005), and contemporary debates over what to do about the undocumented population, which reached a peak of around 12 million in 2007 (see Figure 5). It is important to keep in mind, however, that not all immigration is illegal or undocumented, and not all is unwanted (or unskilled). So-called illegal immigration dominates the headlines and there are powerful anti-immigration lobbies, like the Federation for American Immigration Reform (FAIR) that seek drastically to reduce immigration; but there are equally powerful pro-immigrant lobbies, some of them like the Mexican American Legal Defense Fund (MALDEF) and LULAC are devoted to defending the rights and interests of Latinos. Still others like the American Chamber of Commerce and various trade associations represent powerful business interests. Bill Gates, the founder and former head of Microsoft, organized a successful lobbying campaign by high-tech industries to stop Congress from restricting high-skilled immigration at the time of debate (1995-96) over the IIRAIRA (see Figure 5).

We might have expected Congress to restrict or slow immigration during economic hard times; but at the start of the relatively mild recession of 1990-91, Congress enacted another expansive immigration reform. The Immigration Act of 1990 was designed to reform **legal immigration**, setting an overall annual ceiling of 675,000. Because of family reunification and the fact that visas not used in one year can be carried over to the next, actual levels of legal immigration are much higher, averaging over 1 million/year throughout the 1990s and into the 2010s. Many undocumented migrants were able to “adjust their status” and become legal permanent residents (LPRs or green card holders, see Figure 3), thus adding to the annual totals.

In fact, the United States immigration system relies heavily on adjustments of status to deal with backlogs of individuals who find themselves in legal limbo; and this “adjustment of status” system creates a demand for immigration lawyers and other specialists who advise millions of immigrants and potential immigrants, as well as their employers. The American Bar Association (ABA), specifically the American Immigration Lawyers Association (AILA), is among the most important pro-immigration interest groups. Lawyers are essential for the smooth functioning of the system, because they help to adjudicate and manage hundreds of thousands of cases on an annual basis. This gives the American system

for managing immigration greater flexibility to deal with admissions on a case-by-case basis, even though quotas and quantitative caps on the numbers of visas available for specific nationalities and regions make the system cumbersome and inefficient. The highly individualized nature of this regulatory system is consistent with the broader trend in rights-based politics and policy, which began with the civil rights movement of the 1950s and ‘60s.

The 1990 Immigration Act created a new category of visas (the H1-B) for highly skilled immigrants, thus adding an important economic and human capital (as opposed to family and humanitarian) dimension to U.S. policy, and generating more work for lawyers who specialize in helping employers recruit individuals with skills that match their needs. The 1990 Act set an annual cap on H1-B’s of 65,000, but during the high-tech boom of the late 1990s Congress adjusted the cap in response to higher demand for skilled workers and pressure from business groups. The H1-B visa was designed for high-skilled immigrants, and the H2-A and H2-B visas were created for non-agricultural seasonal workers. However, the number of job-based green cards, whether for the unskilled (capped at 10,000/year) or the skilled (capped at 140,000/year), was too low to accommodate the overall demand for immigrant labor (see Figure 6). The H1-B became a magnet for foreign students, creating a pathway for them to stay and work in the United States (see Figure 6).

Throughout the boom years of the 1990s and into the 2000s, the result was the issuance of more temporary visas (over 600,000 in FY 2005) and rising levels of undocumented immigration (see Figures 4-5). It is difficult for Congress to create an employment-based visa system that mirrors the business cycle and perfectly matches the needs of the labor market. With the bursting of the high-tech bubble in 2001, the demand for H1-Bs declined and a binding cap of 65,000 was brought back in 2004, only to see demand rise again in 2004-07 (see Figure 7). The bursting of the housing bubble in 2008 and the ensuing financial crisis led to declining demand for unskilled immigrant workers, especially in construction; and unemployment reached 10 percent of the labor force as a whole in 2009. Immigration does not follow the business cycle (or employment levels), because of lags between the demand for and supply of visas, the difficulties of quickly adjusting policy, and the rise of rights-based politics (see Figure 8).

To combat undocumented immigration in the 1990s, the Immigration and Naturalization Service (INS) developed ever more sophisticated strategies for border enforcement (external control), increasing the number of border patrol agents and redeploying them at critical entry points along the U.S.-Mexico border. Operations Hold the Line in Texas (1993) and Gatekeeper in California (1994) were designed to seal the border in

urban areas like El Paso and San Diego, and to force clandestine crossings away from the cities into remote, desert areas—a policy of deterrence. These external enforcement policies succeeded in redirecting flows, but levels of undocumented immigration continued to rise, reaching a peak of 12.2 million in 2007 (Figure 5 above). Thousands of migrants would die in the deserts of the southwest, leading some to argue that the policies of the Clinton administration were nothing more than symbolic and cynical attempts to show the public that the government was regaining control of the border—an “out of sight, out of mind” approach to immigration control—but with deadly consequences for the migrants themselves.

The failure of external enforcement policies in the 1990s combined with the 9-11 terrorist attacks, led to a massive reorganization of the immigration bureaucracy. In 2003 the INS, formerly an agency of the Department of Justice, was reorganized into two agencies—one for enforcement, Immigration and Customs Enforcement (ICE) and one for services, Citizenship and Immigration Services (CIS)—and placed in the new Department of Homeland Security (DHS). A third agency, Customs and Border Protection (CBP), was created to coordinate border control. All three agencies were tasked with protecting the U.S. homeland from another terrorist attack, as the security function of immigration and refugee policy came to the fore (see Figure 1). The Real ID Act, passed in 2005, established new standards for driver’s licenses and non-compulsory state ID cards, to make it more difficult to counterfeit these documents and for individuals to obtain false papers. The law was intended to reinforce checks on individual identity, stopping short of creating a national ID card.

The new emphasis on security made travel and immigration to the United States more difficult, especially for anyone coming from a Muslim country; and the issuance of visas in U.S. consulates around the world came under much greater scrutiny, slowing an already cumbersome and inefficient process with elaborate background checks for visa applicants and refugees—it takes on average over two years to process individual asylum claims. Overworked Foreign Service Officers (the front line of immigration control) were fearful of admitting someone who might carry out another terrorist attack. The 9-11 hijackers entered the United States legally on tourist and student visas, but 7 of the 19 had false passports and 3 were on terrorist watch lists, leading the 9-11 Commission—set up to investigate the attacks—to conclude that better immigration and border enforcement might have prevented the terrorists from entering the country. Security considerations aside, the debate over immigration reform during the George W. Bush administration (2000-08), as in previous eras, revolved primarily around the economic effects of immigration, especially undocumented immigration. In May, 2006

the Bush administration proposed “comprehensive immigration reform,” to match “willing workers with willing employers,” by creating a new guest worker program (a return to the Virginia model) and an “earned legalization” program for the millions of undocumented migrants already working in the United States (visions of the Pennsylvania model). Opponents of comprehensive reform charged that it would be a repeat of the IRCA amnesty, creating another situation of moral hazard that would lead to yet higher levels of illegal immigration. The rallying cry of opponents was “fool me once, shame on you, fool me twice, shame on me!” The Sensenbrenner Bill of 2005-06 represented an alternative, “enforcement only” strategy, placing a premium on enforcement of existing laws, reinforced border control, the arrest and deportation of the millions of so-called ‘illegals,’ and the criminalization of undocumented migration, making it a felony.

The collapse of the Bush reform effort in 2006 led many state and local governments to take up the cause of immigration control, further dividing communities and the electorate and intensifying the three-level game (see Figure 1). It proved impossible to resurrect the rights-markets coalitions in Congress that enacted earlier reforms during the cold war period. The Republican Party in particular was divided between a culturally conservative—nativist and xenophobic—wing, which refused to compromise, and a more moderate, business-oriented wing (the latter represented by the late Senator John McCain and the former by President Donald Trump), which wanted to give the Grand Old Party (GOP) a more immigrant-friendly face. The fear among many Republicans, like George W. Bush’s political “Svengali,” Karl Rove, was that demographic changes resulting from high levels of immigration were changing the electorate, and that Hispanics in particular constituted a swing vote in many key states and districts. Some leaders of the GOP did not want to return to the nativism of the 1920s when the Republicans ceded third-wave immigrants to the Democratic Party for the better part of two generations. They became ‘Reagan Democrats’ in the 1980s. In the run-up to the 2008 presidential election and flush from their successes in the 2006 mid-term elections, Democratic leaders in Congress decided against compromise with moderate Republicans, like Senator John McCain, and the Bush White House, preferring instead to leave the immigration issue open, like a festering wound, and to use it against Republicans. In the event, the Democrats fought two successful presidential elections under Barack Obama in 2008 and 2012, only to suffer a massive reversal in the 2016 election when Donald Trump succeeded in using immigration as a wedge issue to divide the electorate.

The Economics of Immigration and Integration

Arguments about the economic impact of immigration and the integration of immigrants abound. Two things are clear, however: the United States is more ethnically diverse than ever before in its history—Latin Americans and Asians have replaced Europeans as the dominant immigrant groups and Africans are not far behind—and immigrants play an increasingly important role in economic growth, providing not only labor, but human capital, and higher rates of entrepreneurial activity. Rather than concentrating in traditional immigrant cities—like New York, Boston, and Miami on the east coast, Chicago and Saint Louis in the mid-west, or San Francisco and Los Angeles in the west and Houston in the southwest—immigrants are settling in new gateway cities—like Dallas-Fort Worth, Atlanta, Phoenix, Washington, DC, Charlotte, Nashville, and Las Vegas to name a few—and in states and cities far from the main ports of entry. States with the fastest growing foreign populations are in the south (North and South Carolina, Georgia, Tennessee, and Arkansas), the west (Nevada, Utah, and Washington), and non-traditional destinations in the east (New Hampshire and Pennsylvania) and mid-west (Iowa, Wisconsin); and there is a close correlation between increases in the foreign population and employment (see Figure 9). Immigration accounted for roughly 30 percent of U.S. population growth from 1980 to 2018. Without immigration, the U.S. population would have stagnated and begun to decline decades ago. The four biggest immigration states in 2018 were California (almost 11 million foreign born), Texas (almost 5 million), New York (4.5 million), and Florida (over 4.3 million each). The leading countries of origin in the fourth wave were Mexico (31 percent of the foreign-born), followed by the Philippines, India, China, and Vietnam (see Figure 9).

The fact that so many of the newcomers are of Latin American and Asian origin has increased the visibility of immigrants across the country, giving greater impetus to debates about immigrant integration. Mexican and Central American immigrants in particular are predominantly unskilled, many are undocumented, and they often speak little English; hence the concerns expressed by Samuel Huntington and others for American national identity (the *unum*). Many of these newcomers—much like their counterparts a century earlier—live in ethnic enclaves in large cities. While the first generation may experience significant improvements in their welfare (compared to their situation in the country of origin), their children, the second generation, may experience significant downward mobility—what the sociologist, Alejandro Portes, termed “segmented assimilation,” which is non-linear and does not lead to “mainstream” outcomes. The theory purports to explain why some second-generation immigrants engage in deviant or criminal behavior, joining gangs or terrorist organizations,

for example. It is important to note that this is not a new phenomenon—visions of the anarchists at the turn of the 20th century or the film of the 1960s, *West Side Story*! Apart from debates over assimilation, the cost of educating immigrant children and providing health care for indigent families have become central features of debates about the fiscal impacts of immigration (see Figure 10).

But given the rapidly declining percentage of the native-born work force with less than a high school education (for men this number fell from over 50 percent in 1960 to about 8 percent in 2018—a great success of mass-based public education), unskilled immigrants fill a niche at the low and the high end of the labor market (see Figures 10 and 11).

High levels of unskilled and undocumented immigration have provoked intense debates among economists over the long-term costs and benefits of immigration, with some, like George Borjas, arguing that the service-based, high-tech U.S. economy no longer needs so many unskilled and poorly educated immigrants; while others, like Gordon Hanson, David Card and Giovanni Peri, point out that key sectors of the economy (agriculture, construction, consumer services such as restaurants and hotels, and health care to name a few) would suffer without access to foreign labor. Critics counter that without immigrants, wages for native-born workers would rise significantly and the jobs would still get done. This assumes that immigrant workers are substitutes for the native-born, and vice-versa; rather than complements, as some would argue. Alan Greenspan, former Chair of the Federal Reserve Board, weighed into the debate, arguing that immigration keeps inflation down (by keeping wages and prices in check), and that high skilled immigrants in particular are a boon for the U.S. economy.¹⁵

Not all immigrants are unskilled and undocumented, especially those from Asia, who often come with much needed skills (Filipino nurses and Indian doctors, for example) and high levels of human capital (the foreign-born share of U.S. scientists and engineers is large and increasing, see Figure 13). These highly skilled immigrants, many of whom come as students, enter at the top of the labor market and in a short period, their earnings rapidly surpass those of natives. Immigrants and foreign students have high levels of entrepreneurial activity, with a willingness to work long hours at low pay, accumulating capital, reinvesting it, and in the process creating new jobs for immigrants and natives alike—a rising tide that lifts all boats, contributing to productivity and economic growth. Certain immigrant groups, like the Koreans and Iranians, for example, have exceptionally high levels of self-employment, 28 and 20 percent respectively, which is much higher than among the native-born (13 percent). Clearly, fourth-wave immigrants are highly diverse, in ethno-cultural terms and in terms of social class, education, and economic achievement.

Not surprisingly, rates of naturalization vary significantly from one group to another. In the 1980s, there was concern that the new immigrants were not naturalizing; but with the political backlash against immigration that began with Prop 187 in the 1990s and continued with the election of Donald Trump in 2016, naturalization rates have increased dramatically in recent decades. Refugees, like the Vietnamese and the Iranians, have very high rates of naturalization, a trend that is reinforced by the fact that they tend to be highly educated and in higher income brackets. Mexicans and Central Americans at the other extreme tend to be poorly educated, they have a higher propensity to return home, and since 1998 Mexicans are allowed to have dual nationality. All these factors have led to lower rates of naturalization among Hispanic immigrants, but the rates have been going up steadily since the mid-1990s, reaching into the 40 percent range for Mexican LPRs (green card holders) by 2018. Of course, these numbers do not take into account the large number of the undocumented, who are not eligible for naturalization.

Hispanics constitute one of the fastest growing demographic groups in U.S. society (58 million strong and 18 percent of the total U.S. population in 2016), but they have low rates of naturalization because many are undocumented. Clearly the biggest immigration issue in American politics going forward (the elephant in the room) is the fate of some 12 million undocumented immigrants who contribute their labor to the U.S. economy but pose a challenge to **state sovereignty and security** (control of borders/territory is a central attribute of sovereignty and vital to security), **rule of law** (it is illegal for employers to hire individuals not authorized to work), and **civil society** (large numbers of individuals living in the shadows at the edges of society is detrimental to the social contract). What to do about this segment of the immigrant population and how to reform immigration policy (how many immigrants should be admitted, from where, and in what status?) are big, unresolved questions.

The former Democratic Speaker of the House of Representatives, Tip O'Neill—when pushed by some of his colleagues to bring the 1986 IRCA to the House floor for a vote—said, “gentlemen, immigration is political death.” O'Neill feared that if you open the issue for debate, everyone will be angry and no one will go away happy. Despite O'Neill's reservations, the 1986 bill did come to the floor, and it narrowly passed. In 2016, immigration is even more controversial and divisive than in 1986, but candidates cannot afford to tread lightly around the issue.

President Barack Obama faced a choice of kicking the can (in this case immigration reform) down the road or opening a divisive political debate. He chose to double down on border enforcement, deporting by far more people than any president in history, and try to bring relief

to some of the undocumented migrants by protecting them from deportation through executive actions like Deferred Action for Childhood Arrivals (DACA) and Deferred Action for Parents of Americans (DAPA). Obama may have alienated some of his new Hispanic constituents—during the 2008 campaign he promised them that he would propose comprehensive immigration reform during the first year of his presidency, but his advisers convinced him that this would be futile, given Republican intransigence. He vowed to pursue reform that would make immigration orderly and legal and to find a “pathway to citizenship” for undocumented migrants, but he was forced by a recalcitrant Congress to back away from these promises, and he never succeeded in passing any significant immigration legislation, much less comprehensive immigration reform. Being the champion of deportations and tough border control policies did not win him support among Republican lawmakers.

Back to the Future? Trump's Immigration and Refugee Policies

In light of the long history of nativism in the United States, the anti-immigrant policies of President Donald Trump are not so surprising. During the campaign, candidate Trump called for “a total and complete shutdown of Muslims entering the United States until our country's representatives can figure out what the hell is going on.” He asserted, “Mexicans are rapists, murders and drug traffickers,” that the United States should build a wall along the entire southern border to keep them out. In the first two years of his administration he has enacted by executive fiat some of the toughest restrictions on immigration and asylum-seeking in decades, separating children from their parents at the border, and reallocating funds from the Department of Defense budget to build more wall along the southern border, adding for good measure that we should not accept immigrants from ‘shithole’ countries in Africa. One of the biggest challenges for managing migration in the 21st century is climate change and the displacement of millions of people in vulnerable regions of the globe, especially Africa, but also Central America and South Asia. The Geneva Convention has no provision for environmental refugees, but this is rapidly becoming one of the most important sources of forced migration.

President Trump made good on other campaign promises, issuing executive orders banning refugees from seven Muslim-majority countries, placing a ceiling of 30,000 on refugee admissions for 2019. How can we assess the impact of these policies on American society and the economy, specifically their effect on U.S. foreign and national security policy? The Trump administration seems to ignore the fact that immigration and refugee policy—while providing a good opportunity to gain votes and to pursue symbolic politics that allow him to shore up support among some elements of his political base—has major implications for foreign and national security

policy, not to mention future economic prosperity and future growth.

Trump's immigration and refugee policy is couched in cultural or 'civilizational' terms, pitting Christians and Jews against Muslims (per Sam Huntington and the clash of civilizations) and Mexicans/Hispanics against whites and blacks. In so doing, the President has divided the electorate and created a perfect storm of opposition to his policies at the national, state and local levels. The battle over so-called sanctuary cities is raging, and 3.3 million Muslim-Americans (roughly 1 percent of the population) feel threatened by Trump's executive orders. At the domestic level the policy shift contributes to an environment of intolerance and intimidation in which hate crimes have been increasing, giving succor to domestic and foreign terrorists. At the international level the policies have alienated allies in the Muslim world (the governments of Iraq, Afghanistan and Pakistan to name a few have expressed opposition) and in the Western Hemisphere.

It is hard to see how these policies will allow President Trump to build domestic coalitions and find a stable equilibrium outcome to the game of immigration politics (again see Figure 1), or how the policies will align with U.S. foreign policy and national security interests. By opting to pursue policies defined almost exclusively in symbolic and ideological (civilizational) terms, Trump is setting himself and the Republican Party (the party of government) up for major domestic opposition. Interests will coalesce to maintain access to much-needed foreign labor (skilled and unskilled), to defend the rights of immigrants (legal and illegal) and refugees, and to protect minorities (Hispanics and Muslims) from prejudice and discrimination. Nativism is shaky ground upon which to build a political consensus, as former Governor Pete Wilson of California discovered. Wilson managed to win reelection as Governor by pursuing anti-immigrant policies (Proposition 187) but his political career went nowhere after that Pyrrhic victory.

Likewise, the intricate demands of foreign and security policy cannot be met simply through appeals to nationalism and a 'clash of civilizations.' At the end of the day, there is no substitute for the intelligence work needed to stop terrorists before they strike. While nationalism (America first) and symbolic politics scapegoating migrants and refugees may make some Trump supporters feel good, banning refugees will not make the country safer. Long-term foreign policy and security interests, such as the need for allies in the Muslim world, for practical solutions to the refugee crises in Africa, the Middle East, Europe, and Central (the northern triangle) and Latin (Venezuela) America, and for a stable, friendly, prosperous, and democratic regime

in Mexico must outweigh the short-term electoral high that comes from nativism and symbolic politics.

Conclusion: The Liberal Paradox¹⁶

Like other democracies, the United States is trapped in a '**liberal**' paradox—in order to maintain economic competitiveness, the United States must keep its economy open to trade, foreign investment, **and** immigration. But immigration, unlike trade in goods and services, or the movement of capital, involves greater political risks. The liberal paradox highlights some of the risks and contradictions inherent in U.S. immigration policy. As with any sovereign nation, it is essential for the United States to maintain control of its borders (a degree of political and legal closure); otherwise the government risks undermining the social contract and rule of law, cheapening citizenship, and deepening the political and social divide. The central challenge therefore is to maintain openness while at the same time protecting the rights of individuals—citizens as well as denizens.

In the 21st century managing migration is a central function of the modern state, and the state must make strategic choices about how many immigrants to accept, from where, and with what status (Hollifield 2004). From the end of World War II until the recession of 2008, immigration in the United States has been increasing and it has played a key role in U.S. economic growth, providing adequate sources of labor/manpower and much-needed human capital. In 2009, the foreign population stands at a historic high of 44.5 million (14 percent of the total population) and climbing. The rise in immigration is a function of market forces (demand-pull and supply-push) and kinship (family) networks, which reduce the transactions costs of immigration. Economic and sociological forces are the necessary conditions for immigration to occur, but the sufficient conditions are legal and political. States must be willing to accept immigration and to grant rights to outsiders. How then can a liberal democracy, like the United States, regulate immigration in the face of economic forces that push it toward greater openness, while security concerns and powerful political forces push it toward closure?

Historically U.S. immigration policy has been driven by three concerns, epitomized by the Massachusetts, Virginia, and Pennsylvania models. The first (Massachusetts) revolves around concerns for national identity, cultural and ideological cohesion (the *Unum*). To what extent is the United States a White Anglo-Saxon Protestant nation and how much diversity (*pluribus*) is acceptable? The second (Virginia model) is primarily concerned about the need for adequate supplies of labor and human capital in a dynamic and fast-growing economy. The third (Pennsylvania model) is open to diversity, tolerant of differences, but stresses respect for the values and

ideals of the Republic. We continue to see each of these ideas at work in debates over immigration reform and how they play out in the coming decades will be vital to political and social stability, as well as economic growth and prosperity.

What does a 'rational' immigration policy look like in the 21st century? Open borders are a non-starter politically and socially, even in the European Union where freedom of movement is one of the 'four freedoms.' Given demographic decline and stagnation in the United States and across the OECD world, there will be increasing competition, not only for high-skilled migrants, but for reliable sources of labor. Japan has been wrestling mightily with this dilemma. One solution is to go back to the Virginia model, relying on guest workers to fill gaps in the labor market. However, to coin a phrase, 'nothing is more permanent than a temporary worker.' The guests will come if invited, put down roots, intermarry, have families, and stay. The only way to make 'guest worker' programs work is to have circular migration, where the migrants can come and go with the seasons and as their family situation evolves. This is how Mexican migration to the United States was managed throughout U.S. history, prior to the hardening of the border in the 1990s, which froze migrants in place, creating a large undocumented population. A rational immigration policy requires more openness, more visas, and allowing markets to work, not closing or hardening borders. It also requires regional and international cooperation to manage flows and asylum seeking.

Without legal avenues (sufficient numbers of visas and green cards) for migration, black markets for labor will grow, thereby undermining the social contract and leading to a nationalist and nativist backlash. As important as the economics of immigration are, governments have a responsibility to treat immigrants with the same high standards we hold for citizens. Immigration is not just about economics — it is about identity, politics. Migrant rights must be taken into account. The social contract will suffer as long as millions of people live in the shadows. Illegal immigration is detrimental to society and to the immigrants.

In the short term, the United States must fix the status of the undocumented immigrants; legalize and regularize this population, because the federal government cannot deport millions of people. How quickly undocumented are put on a path to citizenship is an open question. In the long-term, the United States needs an immigration policy that works for our national interest, with sufficient visas for those people willing to risk migrating and for whom there is high demand. In the words of former President George W. Bush, we need to 'match willing workers with willing employers,' a policy that balances the economic

and social needs of the United States, while respecting the rights of migrants and giving them a path to full membership and citizenship. Only then will we gain the full benefits of immigration.

¹ James F. Hollifield, "Migration and International Relations," in *The Oxford Handbook of the Politics of International Migration*, eds. Marc R. Rosenblum and Daniel J. Tichenor (Oxford: Oxford University Press, 2012).

² James F. Hollifield, Pia Orrenius, and Thomas Osang, eds., *Trade, Migration and Development* (Dallas: Federal Reserve Bank of Dallas, 2006); Philip L. Martin, Susan F. Martin, and Patrick Weil, *Managing Migration: The Promise of Cooperation* (Lanham, MD: Lexington Books, 2006).

³ James F. Hollifield, Philip L. Martin, and Pia M. Orrenius, eds., *Controlling Immigration: A Global Perspective*, 2nd ed. (Stanford, Calif: Stanford University Press, 2014).

⁴ Martin Ruhs, *The Price of Rights: Regulating International Labor Migration* (Princeton, NJ: Princeton University Press, 2013).

⁵ Robert Putnam. "Diplomacy and Domestic Politics: The Logic of Two-level Games." *International Organization*. 1988/42: 427–460.

⁶ The so-called "Know Nothings" were part of a secret political movement in the 1840s and '50s who were opposed to the largely Catholic German and Irish immigration of the second wave. When asked about membership in the organization, members were instructed to say "I know nothing." The movement had some successes in state and local politics but eventually fell apart largely because of divisions between its members over the issue of slavery.

⁷ Section 1 of the Amendment states that "all persons born or naturalized in the United States, and subject to the jurisdiction thereof, are citizens of the United States and of the State wherein they reside."

⁸ Taken from *The Dillingham Commission Reports*, available online at <http://library.stanford.edu/depts/dlp/ebrary/dillingham/body.shtml>.

⁹ A political patronage organization, Tammany Hall gave political favors to local leaders of the Democratic Party in New York, especially Irish and Catholic immigrants. Tammany Hall held sway in New York politics from the time of the creation of the Republic at the end of the 18th century until the 1960s when the big city machines began to lose control to the national parties.

¹⁰ E. Digby Baltzell, *The Protestant Establishment: Aristocracy and Caste in America* (New York: Random House, 1964), p. 30.

¹¹ I first used this term, "rights-based politics" or "rights-based liberalism" in *Immigrants, Markets, and States* (1992), p. 169. The extension of rights to minorities and foreigners in the decades following World War II is one of the most salient aspects of political development in the advanced industrial democracies. The creation of new legal spaces for marginal groups (including foreigners) in societies as different as Germany, the United States,

and Japan is linked to a much broader change in international and democratic politics, which originated with the adoption of the Universal Declaration of Human Rights in 1948 and in the American civil rights struggles of the 1950s and 1960s. A new type of rights-based politics has taken shape at many levels of the democratic polity and in the international system itself: in legislative acts, partisan and interest group (especially ethnic) politics, and, most important of all, in judicial rulings—Hollifield, Martin and Orrenius, *Controlling Immigration* (2014), pp. 8-9, and Schuck (1998).

¹² Jim Crow refers to a system of laws in the states of the old Confederacy put in place from the period of Reconstruction and designed to enforce racial segregation. These laws would be overturned during the civil rights struggles of the 1950s and '60s.

¹³ James F. Hollifield and Carole J. Wilson. 2011. "Rights-Based Politics, Immigration, and the Business Cycle: 1890-2008," in Barry R. Chiswick, ed. *High-Skilled Immigration in a Global Labor Market*. (Washington, D.C.: The AEI Press, 2011).

¹⁴ By the 1990s it was considered politically (and scientifically) incorrect to use the terms assimilation. American sociologists, like Alejandro Portes (1990) prefer the term incorporation, which sees the process as a two-way street. See also Alba and Nee (2003).

¹⁵ "Greenspan Backs Increase in Foreign Skilled Workers," *Wall Street Journal*, April 30, 2009.

¹⁶ I advanced the notion of a liberal paradox in *Immigrants, Markets, and States* (1992) to explain the dilemmas of immigration control in the Western democracies.

James Hollifield is the Ora Nixon Arnold Chair in International Political Economy and the director of the Tower Center at Southern Methodist University. He is also a global fellow at the Woodrow Wilson International Center.

Supporting Data

Figure 1. The Dilemmas of Migration Governance

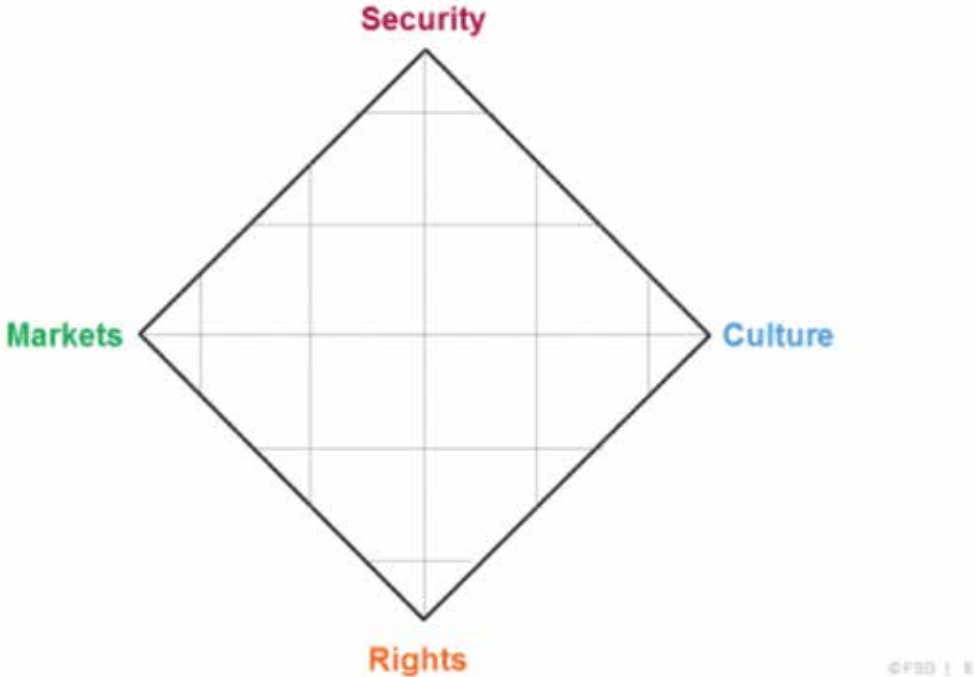


Figure 2. Foreign Population in the United States

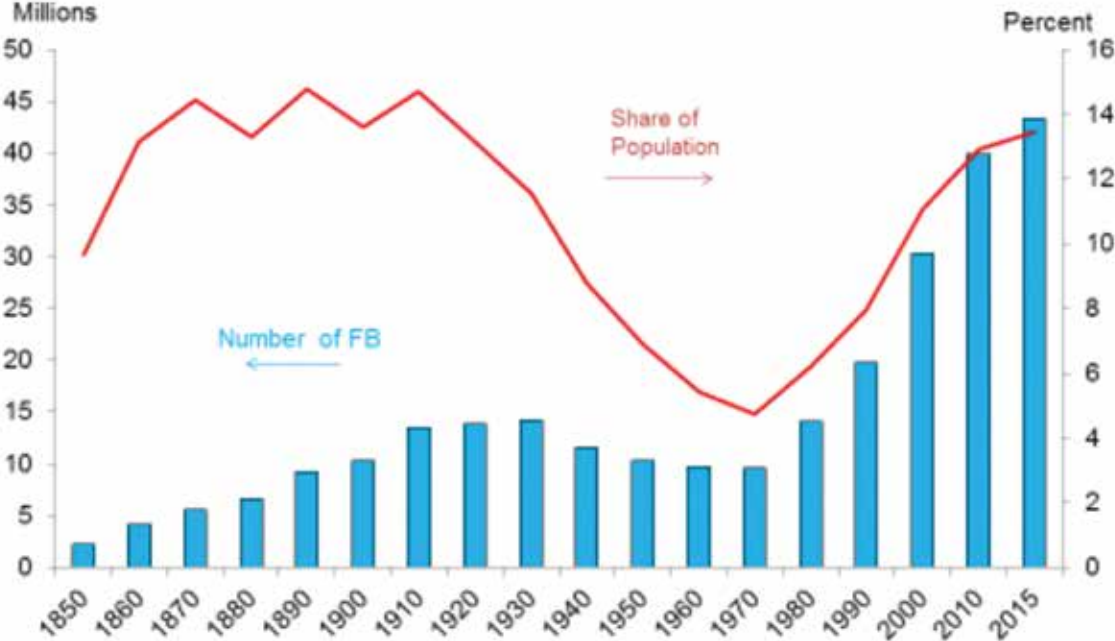


Figure 3. Status of Immigrants in the United States

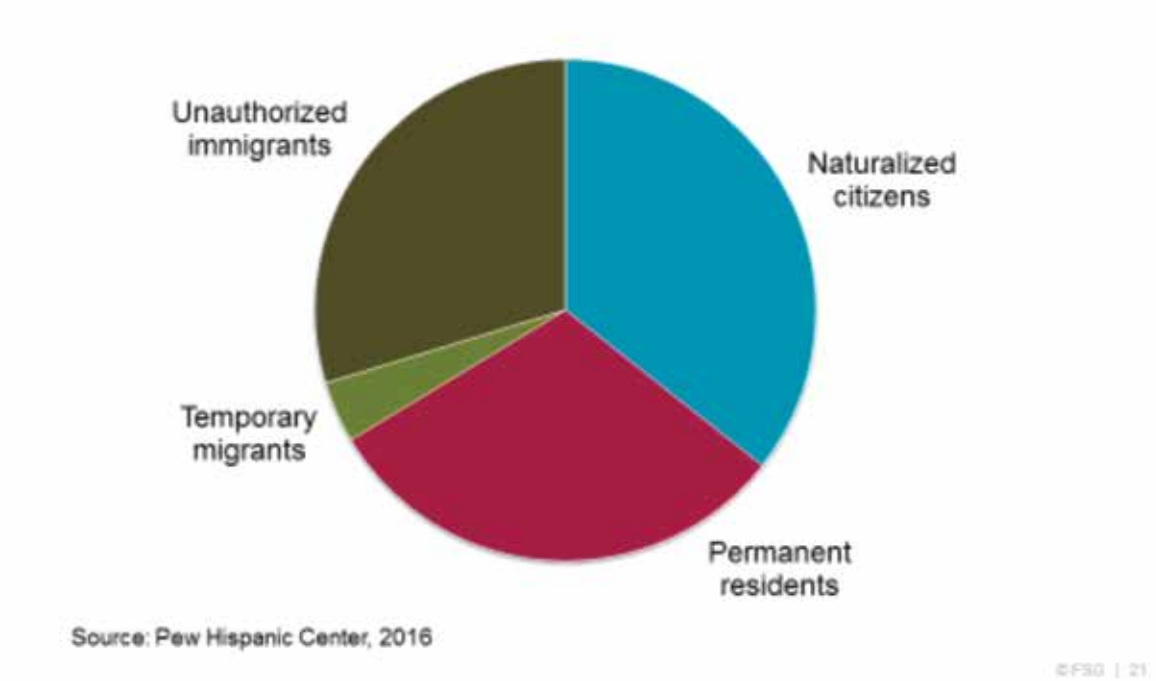


Figure 4. Immigration Flows by Status




Figure 5. Unauthorized Immigration




Figure 6. Shortage of Green Cards

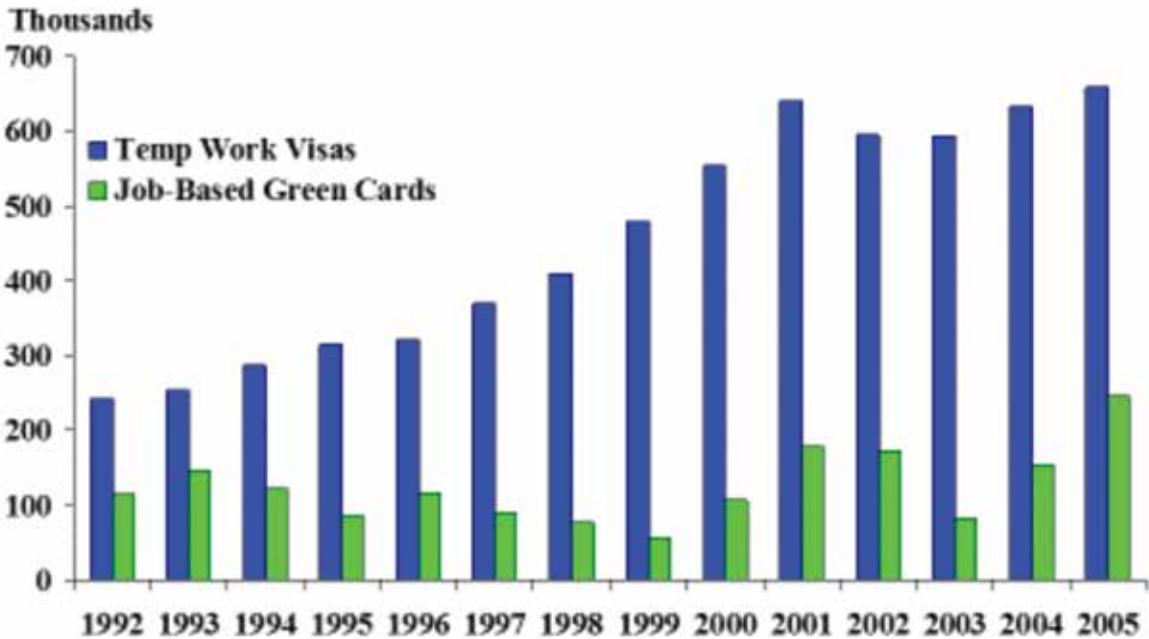



Figure 7. H-1B Visa Cap



Source: Department of Homeland Security, petitions approved for initial employment

Figure 8. Employment and Immigration Diverge

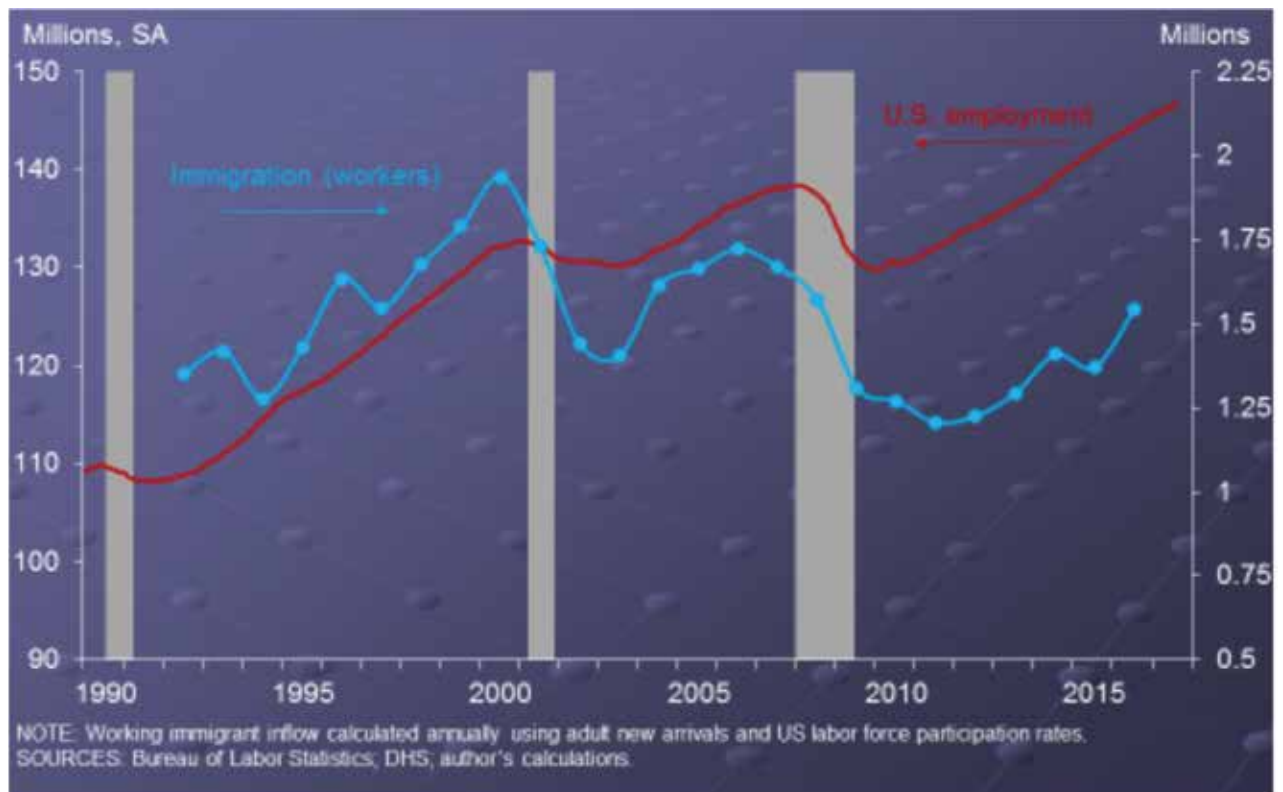


Figure 9. Immigrants Go Where Jobs Are Growing

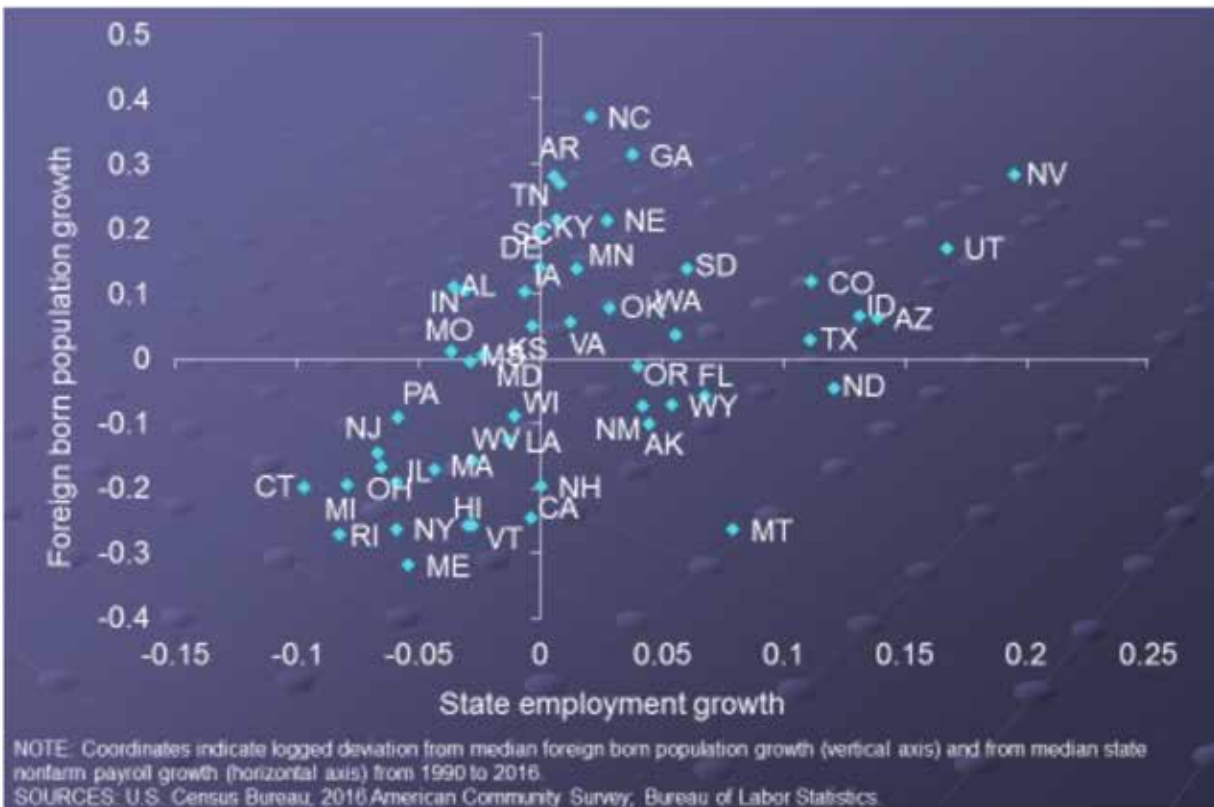


Figure 10. Fiscal Impact of Immigration, Age Is Key

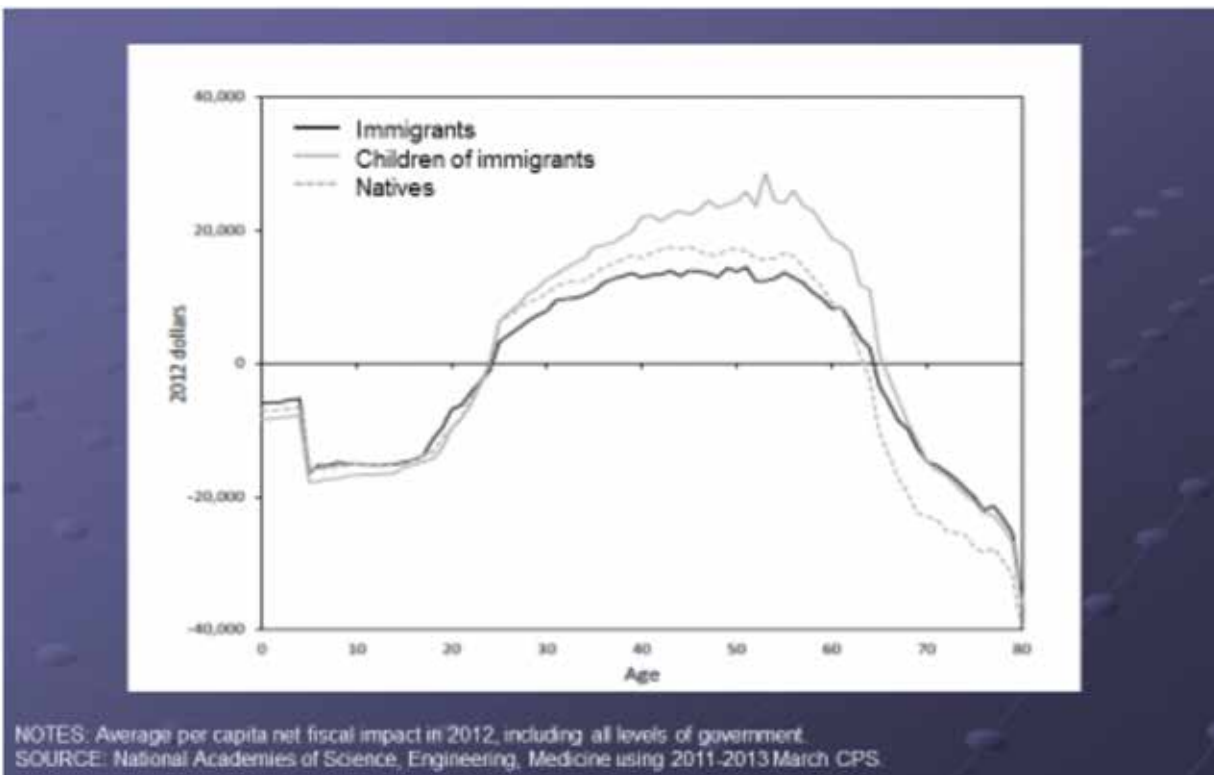


Figure 11. Immigration and Labor Market

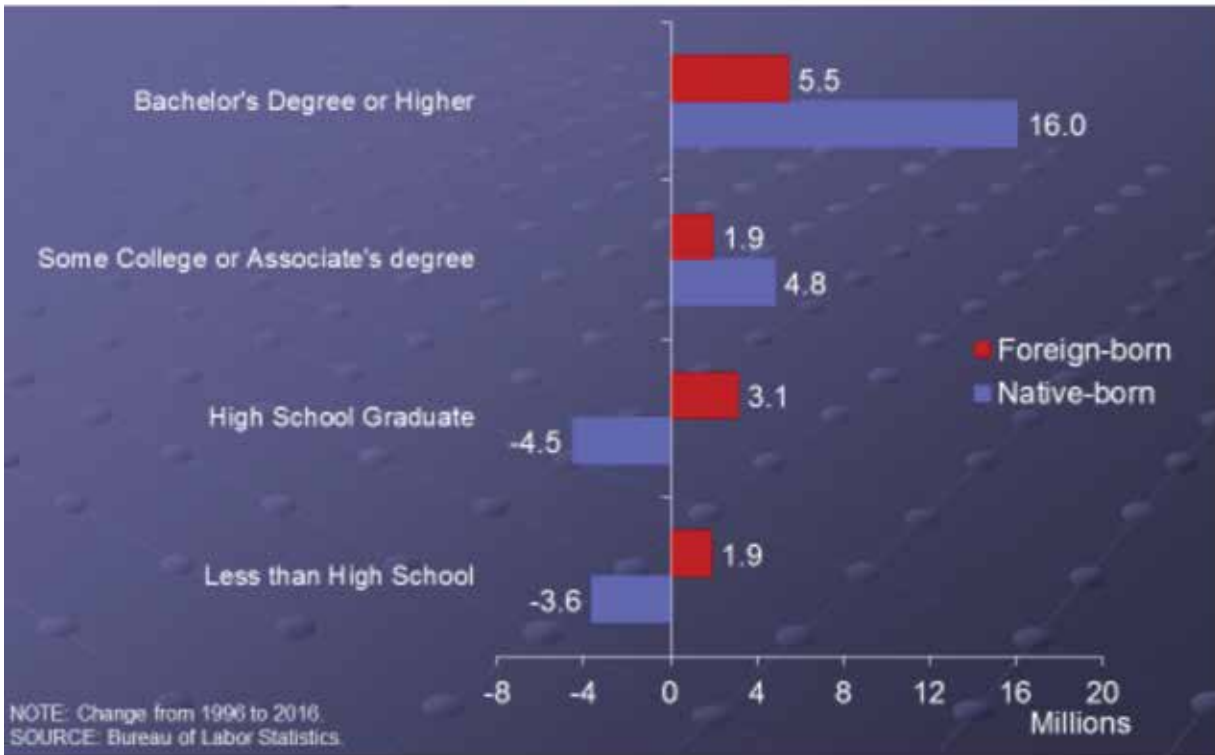


Figure 12. Foreign-Born Share of Labor Force

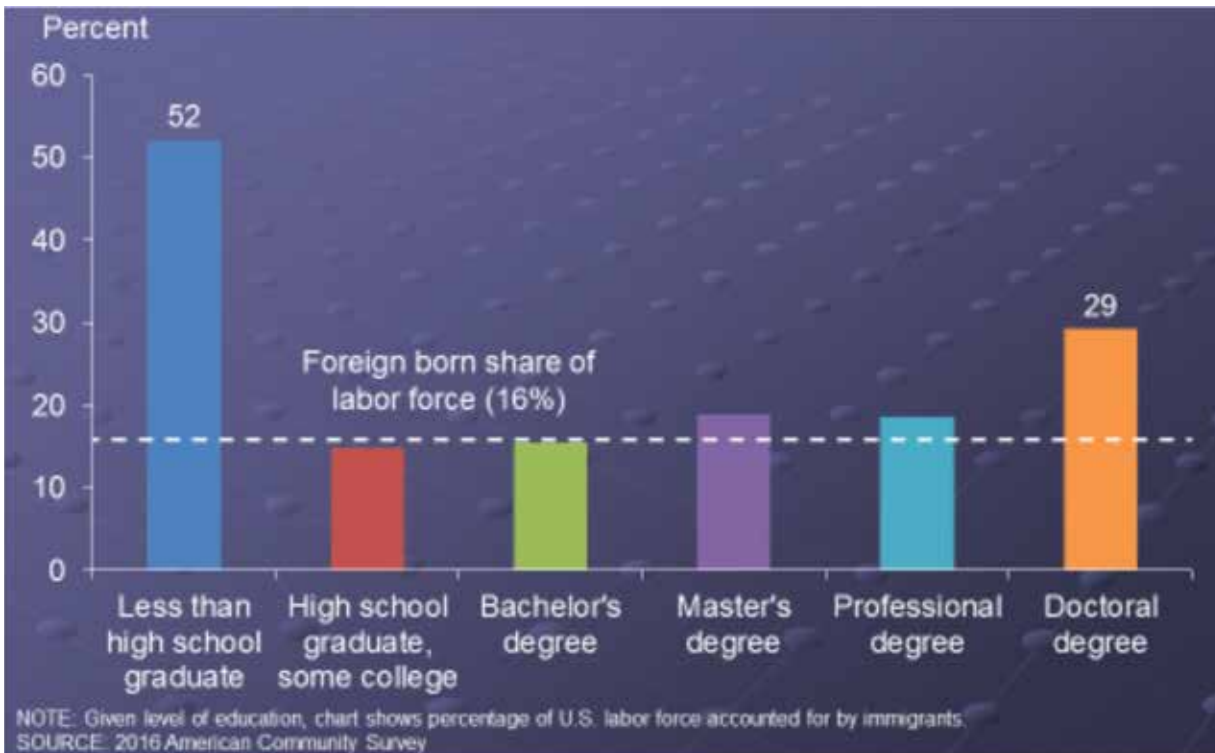
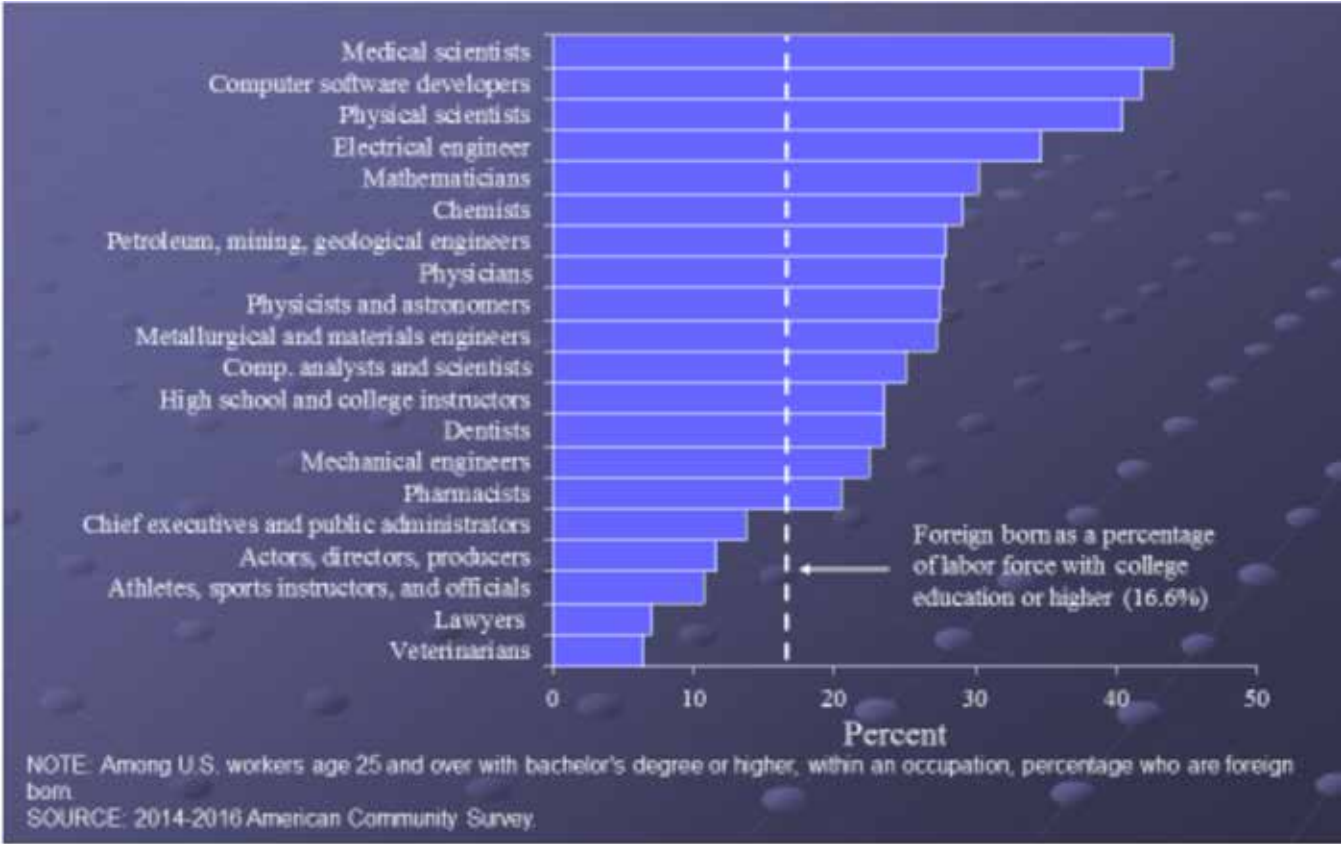


Figure 13. Foreign Born Concentrate in STEM Sectors



Pathways to Economic Opportunity in the 21st Century: A Case Study on How the California Community Colleges Modernized to Deliver on Its Workforce Mission

By Van Ton-Quinlivan, California Community Colleges emeritus



Introduction¹

Since 2012, the California Community Colleges (CCC) system has been driving transformation of its workforce mission to better address labor market needs. From 2012–2018, the California Community Colleges Chancellor's Office (CCCCO) deepened its commitment to modernizing the system's career and technical education (CTE) programs and infrastructure. The approach taken by the CCCCCO was informed by two public policy principles surfaced through the convenings of the California Economic Summit: 1) approach the State as a set of regional economies rather than a monolithic one, and 2) expand CTE capacity in order to provide skilled workers needed by regional economies.² The efforts ushered in an era of experimentation, innovation, and collaboration by California's community colleges and their ecosystem of partners to deliver more workforce programs with better student outcomes.

This renaissance pervaded the state and drove conversations about workforce development from an afterthought to a statewide policy priority. The purpose of this paper is to set the context for and explore this renaissance. The paper is divided into three parts. Part I provides an overview of the nation's shift from a manufacturing- to a knowledge-based economy; presents current demographic data about postsecondary students; reviews research on pathways to good jobs; and outlines program elements necessary for building CTE capacity at community colleges. Part II provides a case study of the CCCCCO's efforts to reimagine and modernize workforce programs across its colleges. Finally, Part III offers higher education six lessons to meet the nation's workforce needs—as informed by the experience of the California Community Colleges.

Part I: The Knowledge-Based Economy

Over the last 50 years, the United States has shifted from an industrial- to a knowledge-based economy. Technological innovations and globalization, among other factors, have profoundly affected the pathway to

economic opportunity. With each passing year, there are fewer good jobs capable of supporting a family for those with a high school education and more jobs that require at least some postsecondary education and training. As a result, workers must increase their skills and credential attainment, or “upskill,” to participate in the job market. “Whereas two out of three entry-level jobs in the industrial economy demanded a high school diploma or less, now two out of three jobs demand at least some education or training beyond high school.”³

As jobs have changed, so have the needs of employers. Employers seek “graduates with a broad array of knowledge and skill—not just specific content knowledge, but transferable skills like critical thinking, the ability to solve unscripted problems, and to communicate effectively.”⁴ Despite high need and good salaries, employers often cannot find enough educated and/or skilled workers. As of January 2019,⁵ the number of open jobs in the United States hovered at 7.6 million, up from 4 million in 2014.⁶ An examination of the nation's postsecondary attainment, confirms employers' struggles. Today, the United States ranks 13th in global postsecondary attainment behind Canada, Japan, South Korea, the Russian Federation, and the United Kingdom. In 2017, the average postsecondary degree attainment in the United States was 48 percent of 25-34-year-olds. In contrast, South Korea's average degree attainment for 25-34-year-olds was nearly 70 percent.⁷

In response to these changes, the Obama Administration⁸ and the philanthropic community began a campaign with a national goal to increase the number of Americans with a high-quality degree, certificate, or other credential from an historical rate of 39 percent to 60 percent of Americans by 2025.⁹ The intensified focus on postsecondary attainment appears to be working. Nationally, “the share of 25-34-year-olds with a tertiary (postsecondary) degree has increased by 7 percentage points since 2007, reaching 48 percent in 2017.”¹⁰ Between 2019 and 2025, the Lumina Foundation

estimates that approximately 24.2 million Americans will earn postsecondary credentials, given current rates of attainment. Though this number is substantial, it falls short of the 60 percent goal. To achieve that goal requires an additional 16.4 million earned degrees or credentials.¹¹

Today's Postsecondary Student

Postsecondary education in the United States has undergone dramatic changes. Women, who attended institutions of higher education in low numbers well after WWII, now comprise almost 60 percent of undergraduates.¹² No longer are college students predominantly recent high school graduates ages 18 to 21. Students today are a diverse group managing multiple commitments and challenges. Data collected by the Lumina Foundation illustrates this point. Of today's college students:

- 37 percent are 25 or older
- 46 percent are first generation
- 42 percent are people of color
- 64 percent work, and 40 percent of these work full time
- 49 percent are financially independent from their parents
- 24 percent have children or other dependents
- 31 percent come from families at or below the Federal Poverty Guideline.¹³

In addition to these demographic changes, there has been a shift in the primary reasons motivating students to pursue higher education. In the 1980s, students reported that they sought to learn more about themselves and their potential first and the ability to gain employment as a result of their degree second. Sensitive to the rapidly changing job market, now students' primary focus is on identifying a clear degree or credential path that will lead to a job.¹⁴ Students are also aware of how expensive postsecondary education and training can be, and the likelihood that they will have to borrow to complete their educational goals. Students have good reason to be concerned. "Tuition has increased 503 percent more than inflation over the last 35 years."¹⁵ For students, there is a real need for a return on their educational investment so they can support themselves, provide for their families, pay their loans, and save for their retirements. Students need good jobs.

Pathways to Good Jobs

In *Three Educational Pathways to Good Jobs*, three pathways to jobs capable of supporting a family are identified, including: high school, middle skills (often

earned through community colleges), or bachelor's degree. The report emphasizes the key role community colleges play in helping students prepare for good jobs.¹⁶ While much of the opportunity has been for students earning bachelor's degrees, a significant area of growth has been in the middle skills pathway for students completing associate's degrees, certificates, certifications, licenses, and some college but no degree. Middle skills jobs account for 24 percent of good jobs and "good jobs are growing the fastest among workers with associate's degrees (AAs)."¹⁷ Between 1991 and 2016, the middle skills pathway added 3.5 million jobs while the high school pathway lost 1.8 million.¹⁸ The middle skills pathway benefited from a surge of skilled-service jobs that include healthcare, education, business, finance, and leisure and hospitality.¹⁹ "The growth of good middle-skills jobs in skilled-services industries has proved a wellspring of opportunity."²⁰

In response to these factors, President Obama held the first summit of community college leaders in 2010. In his remarks, the President noted, "[i]n the coming years, jobs requiring at least an associate's degree are going to grow twice as fast as jobs that don't require college. We will not fill those jobs—or keep those jobs on our shores—without community colleges."²¹

Building CTE Capacity at Community Colleges

Building on the Summit, the Obama Administration sought to modernize the nation's public workforce system. In 2014, the Workforce Innovation & Opportunity Act (WIOA), previously the Workforce Investment Act, was signed into law.²² Simultaneously, the White House and the National Economic Council in partnership with the Departments of Commerce, Education and Labor developed an action plan to implement a job-driven and employer-led system of workforce development "connected to real jobs in local or regional labor markets, and resulting in skills and credentials that help individuals succeed in the labor market."²³

As the pathway to economic opportunity continues to shift from a manufacturing- to a knowledge-based economy and beyond,²⁴ higher education will be asked to respond to rapid changes in job creation with corresponding education and training, community colleges are an important partner in the alignment of degree attainment and workforce needs. Every year community colleges enroll approximately 13 million students in credit and non-credit courses.²⁵ Community colleges are open admission, close to home, and affordable. In addition to traditional general education pathways, community colleges offer a variety of career technical education (CTE) programs. CTE programs provide training for a variety of careers including but not limited to nurses, emergency medical technicians, welders, utility lineworkers, plant operators,

maintenance and repair technicians, and early childhood education providers.

Building CTE capacity and availability at community colleges institutions requires many program elements to be successful, including:

- *Understand labor market needs.* Responding nimbly to changing workforce demands requires policymakers, administrators, and faculty understand labor market needs and trends at all levels—local, regional, and state—and allocate resources in accordance with an understanding for those needs.
- *Engage employers.* Traditionally, developing new programs has been the domain of faculty. Designing relevant career technical education programs requires that faculty authentically engage employers in curricular design and implementation—to ensure that the skills students developed are current.
- *Respond to the life cycle of new jobs.* For many of the new jobs emerging in the workplace, the traditional life cycle of developing CTE curriculum for a new program is too long. In five years, for example, jobs can fundamentally change or disappear. As a result, the timeline to develop and implement a new program needs to be timely to match the lifecycle of workforce demands.
- *Modularize education and training programs.* Modularizing curriculum offers one way to close the gap between the skills students have and the skills employers need. Modules can augment existing CTE programs to allow for a more rapid response to changing workforce needs. In addition, shortened curricular pathways can be more palatable for working adults who may need only a skill booster rather than a degree to be attractive in the job market. Creating curricular components that can be mixed and matched allows customization for geographical differences in industry needs.
- *Pool employer demand.* For a new program to be viable, colleges need to enroll a minimum cohort of 15-25 students. Hence, colleges produce graduates in volume. In contrast, employers tend to hire a limited number of people trained for the same occupation at any one time, leaving most of the cohort without employment. To match supply of graduates with demand of hires, two options are available: 1) an employer must aggregate job postings, or 2) multiple employers must pool their hiring to time with the graduation.
- *Braid resources.* As discussed, students are a diverse group with wide-ranging needs. For many students,

financial aid for tuition does not meet their full cost of attending postsecondary education and training. For example, students may need assistance with child or other dependent care, buying books, or paying for transportation. Braiding together multiple funding streams available for tuition and other student support services is essential for enrolling students and retaining them in programs.

- *Balance technical and transferable general education skills.* While acquiring industry-relevant and occupation-specific skills, students also need to build transferable skills—like strong English, math, and requisite social skills—required by employers. Some students arrive at community colleges with these skillsets underdeveloped,²⁶ whether from inadequate K-12 preparation, immigration across national and/or state borders, frequent school changes as a result of disruption forces, or other reasons. With the guided pathway movement²⁷ at some community colleges, general education faculty have partnered with CTE faculty to create courses and curricular pathways that integrate general education into workforce pathways. For example, students in a health pathway can be taught reading comprehension using healthcare-themed texts or receive mathematics instruction through calculating dosages. From an employer's point of view, development and mastery of basic, transferable skills should be the responsibility of the K12 and higher education and are essential for employment.

To help illustrate this point, see Figure 1. The energy industry considers Tiers 1 and 2 to be foundational skills that ought to be groomed in students through their K12 and postsecondary education. Gaps in student skillsets, especially in these tiers, frustrate employers the most. Moving up the stack, ownership for developing the skills in Tiers 3 and 4 can be shared between industry and postsecondary education, with some of these skills best gained through work-based learning programs, apprenticeships, internships, and cooperative experiences. Tiers 5 through 8 are industry- and occupation- specific skills that employers consider their purview to develop (see Figure 1)²⁸.

- *Ensure predictable funding* - New CTE programs are often funded through grants and have three primary cost elements: 1) one-time expenditures for equipment and consumables; 2) on-going maintenance; and 3) on-going faculty personnel. While grants allow for new programs to begin, these funds often do not last long enough to ensure program stability, especially with the latter two cost categories. It is not surprising, then, that institutions

are reluctant to hire full-time faculty for programs using temporary funds. Additionally, often only full-time faculty can introduce new curriculum into the approval process. Without the predictable funding to hire full-time faculty, new CTE options cannot be introduced. Fresh, new programming requires funds that will address both one-time and ongoing cost.

- *Develop K-12 and community college connection.* For many students, there is a disconnect between career pathways identified in high school and those that are available at the community college. This is because K-12 and community colleges have not developed synchronized pathways corresponding to regional labor market needs. Developing pathways that are shared between both segments, allows for K-12 students to engage in early career exploration opportunities and identify pathways into meaningful employment.

Part II: California Community Colleges for Educating and Training Students and Workers: A Case Study

The CCC is a public system of 115 community colleges. Serving 2.1 million students, it stands the largest system of higher education in the nation.²⁹ Compared to the University of California and the California State University systems, both of which have selective admissions processes, the California's community colleges are designed by the California Higher Education Master Plan to be broad access, available to all students.

Since its inception, there have been numerous attempts to update the Master Plan. The most recent of these in 2018 "directed the Governor's Office of Planning and Research (OPR) to 'conduct a review of state policies developed under the Master Plan for Higher Education and future workforce needs.'"³⁰ The report made important findings about the ability of the current higher education system to address student needs, tackle changing demographics, and meet workforce demands. The report argues that:

Addressing the needs of California's current and future population and ensuring that Californians have viable and accessible pathways to meaningful degrees and credentials requires a more student-centered perspective than existed in the days of the Master Plan...Meeting student and workforce needs requires continued work to eliminate barriers to access and completion, to create coherent pathways to degrees and credentials and to provide support to students on those pathways.³¹

Concerns about workforce needs are not new. In September 2012, 67 percent of Californians felt that "jobs and the economy [we]re the most important issues facing the state."³² Despite the urgency many Californians felt, a review of 2001–2011 data shows an overall diminishment

of the CTE portfolio in the CCC (see Figure 2). At the time, there was an inherent financial disincentive in the way community colleges were asked to support workforce programs. "Despite higher than-average costs in many CTE³³ fields, CTE programs generally receive no more per-student state funding than liberal arts and science programs."³⁴ Consequently, the ten-year downward trend in the system's career technical education portfolio (as a percent of full-time equivalent students) was not the result of CCC system directives but rather choices made by individual colleges (see Figure 2)³⁵.

Commencing in 2012, the CCCCCO conceptualized and began operationalizing *Doing What MATTERS for Jobs and the Economy* (DWM)³⁶ as the framework for transforming its workforce mission. Over the span of seven years from 2012 through 2018, policymakers gained greater appreciation for the value of CTE and grew public investment—from \$100 million to over \$1 billion—through dedicated 'categorical' funds. Updated metrics and new data tools, along with numerous legislative, regulatory, and administrative changes, accompanied increases in state funding. To better understand CCC's execution of the DWM framework, implementation is broken into three sections: 2012-2013, 2014-2015, and 2016-2018.

2012–2013

With the advent of new leadership in the Workforce and Economic Development Division in the Chancellor's Office, the CCCCCO initiated a set of administrative changes designed to streamline and integrate existing resources designated for CTE programs. One of the most important changes was redesigning the use of the existing \$100 million in 'categorical' funds. With the goal of improving student workforce outcomes, these resources were restructured to advance the DWM framework along three themes, including: 1) industry sector strategy; 2) regional coordination; and 3) technical assistance.

Industry Sectors

Ten industry sectors were important to California's regional economies based on labor market need: health; information and communications technology; advanced manufacturing; energy, construction, and utility; advanced transportation and logistics; agriculture, water and environmental technologies; life sciences/biotech; retail, hospitality and tourism; small business and entrepreneurship; and global trade. Of these ten, colleges in each region of the state worked with labor market data and with their partners to prioritize the industry sectors most important to their economy. This, in turn, allowed for better tailoring of resource allocation within each region along career pathways and CTE programming.

Regional Coordination

Collaboration and coordination among and between many organizations is key to a successful student journey, especially in geographies where colleges are densely located, as they are in urban areas. Prior to the CCC reforms, competition dominated institutional interactions and coordination across community colleges was irregular. A focus on regional collaboration changed this dynamic and enabled region-wide action and replications of innovations at regional scale. Smaller regions like South Central number seven community colleges with San Diego and North Far North at ten institutions. On the larger count, sixteen community colleges are in the Central Valley while twenty-four are in the Bay Area, and twenty-eight are in Los Angeles/Orange County. A tremendous amount of institutional capacity is parlayed when operating at regional scale.

Technical Assistance

Anticipating that colleges would struggle in a number of areas, such as learning to use new data tools and navigating curricular approval processes, the CCCCO established a network of technical assistance providers whose role it was to support and train the field in specific domains.

Identification of Further Reforms

During this period, further reforms were identified. The California Community College Association of Occupational Educators (CCCAOE), comprised of CTE deans and faculty, wanted the state's funding formula to address the higher cost of CTE programs. This could not be achieved through available administrative means. Making a change to the funding formula required new legislation paired with changes to the state budget. Both required extensive dialogue with a broad array of stakeholders, some of whom did not understand the value of workforce programs. While industry trade associations were supportive, the Academic Senate of Community Colleges, whose leadership was comprised mainly of faculty from the general education background, favored maintaining the status quo. General education faculty were unfamiliar with employers and labor market data, did not have to account for job placement nor wage gains, and often did not share the same urgency to undertake reforms.

To help policymakers better understand the earnings return on investment for students in associate and certificate level educational pathways, the CCCCO designed and launched an online tool, Salary Surfer, in 2013.³⁷ Salary Surfer visibly illustrated student earnings two years before and two and five years after completing a community college degree or certificate. Using unemployment insurance wage data from the California Employment

Development Division, Salary Surfer demonstrated that statewide students completing CTE programs averaged \$66,000 in earnings five years after completion compared to \$38,700 earned by students completing general education associate's degrees statewide. Notably, as Figure 3 shows, \$66,000 in earnings exceeded the living wage of \$60,771 needed to sustain a family of four in California (see Figure 3)³⁸.

2014–2015

In the fall of 2014, in order to broaden the stakeholder dialogue around CTE to enable further reforms, the California Community Colleges Board of Governors launched the Task Force on Workforce, Job Creation and a Strong Economy (later referred to as Strong Workforce Task Force) to provide recommendations for addressing the projected shortfall of one million middle skill workers needed by the State over the ensuing decade. The 36-member task force, comprised of stakeholders within and external to the system, produced 25 recommendations representing policy and practice changes necessary for increased production of industry-valued degrees and credentials.

The recommendations synthesized feedback from more than 1,200 stakeholders and 20 regional college conversations and townhalls during a nine-month period "to identify actions that c[ould] be taken to provide policy guidance, regulatory review, and legislative and budgetary actions with the goal of increasing the number of students obtaining CTE degrees and certificates—crucial for closing California's skills gap."³⁹ In September 2015, the Board of Governors unanimously adopted all 25 recommendations of the Task Force. The recommendations fell into seven areas: student success, career pathway, workforce data and outcomes, curriculum, CTE faculty, and regional coordination. Table 1 below lists the recommendation and corresponding type of action necessary to implement the recommendation⁴⁰ (see Table 1).

2016–2018

In June 2016, both the California Legislature and Governor echoed support for the Task Force's 25 recommendations and consequently bolstered CTE across California's community colleges with an infusion of \$200 million in recurring annual funds. These monies were specifically allocated to address the recommendations of the Strong Workforce Task Force. The "Strong Workforce Program" was enacted in law and realized in the state budget that summer. The magnitude of the budget allocation and ongoing nature of the funds was unprecedented for CCC's workforce mission.⁴¹

The Strong Workforce Program also set forth in law a requirement that 17 percent, or \$34 million, of the \$200 million be used to incentivize student completion and

employment outcomes, using measures aligned with the federal Workforce Investment and Opportunity Act (WIOA). Termed the *17% Positive Incentive*, the CCCCO commissioned four papers to inform the design of this performance-based funding model.⁴² The smooth rollout of the 17% Positive Incentive paved the way for later broad-scale adoption of performance-based funding by the CCC. In 2018, the Student-Centered Funding Formula (SCFF) affecting the system's \$7 billion in baseline funding was made into law. The SCFF emphasized disbursements based on student success.⁴³

Between 2016 and 2018, due to the reforms enacted during the tenure of DWM which included the Strong Workforce Task Force, the California's community colleges engaged in unprecedented levels of experimentation, innovation, and collaboration. Some notable examples are highlighted below.

Region-wide action: Curriculum approval was a time-consuming process involving local, regional, and state levels. As CCCCO encouraged community colleges to collaborate within their regions, community colleges began experimenting with economy of scale. A process breakthrough came in 2017 when the combined multi-regions of Los Angeles/Orange County determined that the labor market need for the Business Information Worker certificate was sufficiently large to warrant the first region-wide curriculum approval. This allowed all 27 colleges to offer the same curriculum without having to individually undertake regional approval. Other regions followed suit. By 2018, 86 community colleges in California offered the Business Information Worker pathway. Replication of a CTE program had never occurred at such a rapid pace before.

Deregulation of the state curriculum approval process through the application of Six Sigma principles: Encouraged by the CCCCO, ten community colleges of the North Far North region (which includes those within Sacramento to the northern borders of California) participated in a pilot to reexamine the traditional three level curriculum approval process by applying Six Sigma principles. The Six Sigma pilot sought to eliminate waste from administrative processes. The pilot illuminated ways to reduce the curriculum approval time from 6.6 to 3.2 months. Despite significant initial resistance to the pilot, participants saw the value of the process improvement and recommended to deregulate. Authority was given back to local colleges to approve curriculum without requiring state approval, which allowed for more timeliness in curriculum development.

New employer partnerships: Amazon Web Services experienced a shortage in cloud computing skills for itself and its customers. In the fall of 2018, Amazon Web Services, aided by the Los Angeles Economic Development Corporation, announced a partnership entitled California

Cloud Workforce Project. The project included all 19 community colleges in the Los Angeles area and sought to develop workers with cloud computing skillsets. Additionally, all the community colleges committed to working with their affiliated high schools to assist students with this industry-valued credential.⁴⁴ Apple Retail leaders saw the ten employability skills outlined in the system's newly developed New World of Work (NWOW) 21st Century Skills. Eight Apple stores were paired with eight community colleges' retail management pathways as part of Retail Ready California⁴⁵ to source local hires. Retail Ready is on track to expand to 53 community colleges by the end of 2019. Both Amazon Web Services and Apple Retail became partners with the community colleges as a result of reforms undertaken in DWM.

K-12 collaboration to facilitate early career exploration of industry sectors: The CCCCO joined with the Governor's Office on Business and Economic Development to generate awareness for cyber security careers amongst high school students. To groom early career exploration, the CCCCO supported the adoption and expansion of the CyberPatriot competition as a feeder program for pathways in cybersecurity. "The CyberPatriot competition puts teams of high school and middle school students in the position of newly hired IT professionals tasked with managing the network of a small company."⁴⁶ By 2018, 17 of the 28 middle and high schools that advanced to the national finals came from California.

Civic stewardship to strengthen community colleges: The CCCCO, California Forward, the California Stewardship Network, and the Stanford Educational Leadership Initiative joined in an initiative to strengthen the ability of California's civic leaders to serve as effective stewards of public investments in community colleges. The three organizations developed a seminar entitled "Civic Stewardship to Strengthen Community Colleges⁴⁷." The professional development sought to give civic leaders and college leaders a better understanding of ways to collaborate with each other, a novel approach for reinforcing reforms.

CTE rebranding throughout the state: The CCCCO launched a statewide branding campaign for CTE to address the Strong Workforce Task Force recommendation that students and employers lacked visibility into these programs. The objective of the campaign was "to raise awareness among students, their influencers, business, and others about the variety of workforce pathways available through CTE programs at community colleges" that led to good jobs.⁴⁸

Improved student experience: Pioneered through a collaboration between the CCCCO, Bakersfield College located in California's Central Valley, and Concentric Sky, Program Mapper was prototyped to help students visualize their education program options in the context

of expected employment outcomes along a guided pathway. “Featuring an interactive, pathway-based visualization of the traditional course catalog alongside easy-to-understand career data, Program Mapper facilitates a deeper understanding of the potential options a student might have as they explore the programs offered at a college.”⁴⁹ A patent has been filed for Program Mapper. To date, ten community colleges are in the midst of adoption with work underway to enlist the next ten. The California State University system has also begun conversations to extend the technology for its use.

Strong Workforce Stars: Driving workforce outcomes required an upgrade of the system’s top-level metrics to reflect the recommendations of the Strong Workforce Task Force⁵⁰ that acknowledged success looked different for CTE students. These metrics were incorporated in the rollout of the system’s new LaunchBoard data tool. The tool was designed to support CTE practitioners in decision-making by automating the data collection process and providing visual data for continuous improvement. In order to encourage a culture that welcomed data discussions rather than using data for punitive purposes, the CCCCO additionally sought to celebrate high-performing programs by launching an annual recognition of Strong Workforce Stars.⁵¹ These were CTE programs where students showed significant gains in factors important for advancing social mobility⁵² including a substantial increase in earnings, attainment of a living wage, and/or securing of a job closely matched with the field of study.

Some 1,387 career technical education programs received Strong Workforce Stars designation in 2018.⁵³ Every community college in the system received at least one star. Community colleges were able to earn multiple stars since no cap existed on the number of programs eligible to receive the designation. In addition to acknowledgment by the Chancellor’s Office, some colleges received praise from their legislators, and most colleges issued press releases to share the news with their local community. In ensuing years, the expectation was for more CTE programs to receive Strong Workforce Stars designation, a marker of quality outcomes.

As mentioned earlier, one of CCC’s fundamental principles was viewing California as a several economic regions. This concept allowed the CCC to acknowledge each region’s unique CTE capacity and resources. Thinking in terms of regions was a more manageable way of coordinating programs, engaging stakeholders, and mobilizing assets. Stakeholders were more invested in the success of their region’s new CTE programs. Since CCC’s implementation of the concept of regions, the idea has spread from community colleges to involve other organizations and institutions with a role in a region’s workforce and economic development.⁵⁴

Part III: Six Lessons for Higher Education to Meet Workforce Needs in America

Higher education will continue to face disruptive forces and changing demographics. However, as the California Community Colleges example illustrates, cultivating innovation, experimentation, and collaboration offers six important lessons⁵⁵ for how colleges can think differently to meet the nation’s workforce needs.

1. *Regional coalitions of colleges, community service providers, and employers are essential to success. No individual institution can meet student needs and workforce gaps alone.*

Colleges can do many things within their footprint, but the needs of students are sometimes greater than the resources available at campuses. For example, many students struggle with food security.⁵⁶ This resulted in the need to partner with local food shelters to come onto campuses. The federally-funded America’s Job Centers and the USDA SNAP (Supplemental Nutrition Assistance Program) Employment & Training Program also provides food and other student supports but must be accessed through collaboration with county health and human services agencies. Braiding a variety of resources within a region provides comprehensive assistance to high need students as they navigate their learning journey.

Industries have more skill gaps than any one campus can address. In the healthcare industry, for example, hospitals face shortages of personnel ranging from nurses to sonographers to radiology technicians. Despite high need for these skilled workers, training for each occupation is expensive given the specialty equipment required. Colleges are reticent to host too many high cost programs. Adopting a regional lens allows for a coordinated, portfolio approach to workforce pipeline development across and allows division of specializations among participating colleges.

2. *Postsecondary education and training must be ongoing endeavors for students and workers.*

Community colleges see a wide diversity of students, not just recent high school graduates. Adults in the workforce increasingly need skilling, reskilling and upskilling to find, retain, and/or advance in their careers. Few colleges are able to deliver the affordable, convenient quality learning experience needed by their traditional base of high school graduates and demanded by displaced, employed, and underemployed adults.

Postsecondary education needs to reimagine the student journey for adults and redesign programs

and curriculum in ways that serve these students. In much the same way that vaccinations need booster shots to remain effective, students wishing to keep up with the evolving economy must pursue education and training infusions to remain competitive.

3. *Staying current with changing workforce needs means partnering with employers.*

Globalization and digitization are transforming business and workforce needs at an increasingly rapid pace. Colleges need to talk in earnest with employers to understand what they value and align their curriculum with what employers need. In order to stay current, higher education needs to improve workforce outcomes for students to help fuel a strong economy.

The CCC prioritized which industry sectors drove each regional economy as a foundational decision for aligning programs with labor market needs. Based on this prioritization, Strong Workforce Program resources were made available to ensure “more and better” CTE offerings.

As an example, employer advisors in the energy, construction and utility⁵⁷ industry gathered and pinpointed the HVACR (heating, ventilation, air conditioning and refrigeration) Excellence credential as the most valued within the industry. At the time, there were also 1,261 job openings projected across the state, while supply from CCC colleges was only 393. Faculty “cross-walked” their curriculum against the competencies outlined in the HVACR Excellence credential. Though this work, faculty found important gaps which they addressed and worked together on outreach and production of an enlarged talent pool.

4. *Students want their institution of higher education to solve the education-to-employment dilemmas.*

After earning their degree or certificate, the majority of students want to successfully enter the job market but often need help landing their first job. Career center and resume assistance are often inadequate. Colleges that actively support students through internships, work-study co-ops, and apprenticeships give students exposure to work in the “real world.” Students gain work experience and differentiate themselves from the plethora of entry-level candidates.

Work-based learning models are particularly helpful for students to practice employability skills such as teamwork, communication, and cultural competency while gaining exposure to the norms

of the workplace, like showing up on time and dressing appropriately. Models for experiential learning also allow students to apply technical and digital skills.

CCC’s California Apprenticeship Initiative grants⁵⁸ expanded apprenticeships from their traditional use in the construction industry into new and emerging industries across 80 grants. Grants were made to introduce apprenticeship training in industries like biotech, healthcare, early childhood education, and more because this model proved a reliable way to help build a strong workforce.

Germany, Switzerland, and Canada intentionally designed these work-based learning experiences into a student’s educational journey, while many colleges in the United States leave employment to happenstance. Colleges must think differently to help students adapt to the changing job market.

5. *Students demand personalization in their educational experience.*

In an era of Netflix and Amazon Prime, students increasingly want and expect everything to be personalized. One generation from now, students will want their educational journey styled the same way.

Students will want courses that offer relevant knowledge and skills, and that do not repeat what they already know and can do. They will want support services specific to their personal circumstance, recommendations for student clubs and activities based on their interest profile, and readily available classes at times and in formats that work for their schedules.

Colleges must be ready for a day when students unbundle and rebundle their college offerings. For example, students may take online coursework from online provider Coursera to earn a specialization even as they live on college campus. As the nation’s technology infrastructure allows for more options, students will demand more customization. Institutions need to upgrade data and technology infrastructure, as well as provide professional development to faculty and staff, to be ready to serve this new generation.

6. *Artificial Intelligence (AI) and machine learning will continue to shift skill sets in the future of work.*

AI is part of daily life. Most people rely on their mobile devices for everything from conducting business to mapping directions. As AI and

machine learning become more sophisticated, skillsets for humans to work alongside the machines (robots, computers and devices) or to work within the machine (virtual environments) will evolve. Demand for rote memorization will diminish. Instead, collaboration, creativity, critical thinking, and communication skills will be the transferable skills that students must develop for employability.

Conclusion

Thinking differently does not come easily to higher education institutions. There are decades of tradition to overcome and campus interests that favor preserving the status quo. However, change is essential if colleges want to be relevant in the 21st-century economy and have their students thrive in the workforce.

Change is never easy, but it is possible. From 2012 through 2018, the California Community Colleges ushered in a renaissance of its workforce mission. Public investments grew from \$100 million to \$1 billion in seven years for the 115 community colleges in the system. As a result, workforce education and training in California was transformed from an afterthought to a state policy priority.

¹ The author would like to thank James Timbie, Terese Rainwater, Mike Kirst, and Pete Weber for their advice and assistance with this paper, and express appreciation to CCCCO Chancellors Jack Scott, Brice Harris, Erik Skinner, and Eloy Oakley for their leadership during her tenure as an appointee of Governor Jerry Brown.

² "California Economy, California Economic Summit." About Us. Accessed April 2, 2019. <http://caeconomy.org/pages/about-us>.

³ Carnevale, Anthony P., Jeff Strohl, Neil Ridley, and Artem Gulish. *Three Educational Pathways to Good Jobs: High School, Middle Skills, and Bachelor's Degree*. Report. Georgetown University. 2018. Accessed March 1, 2019. <https://1gyhoq479ufd3yna29x7ubjn-wpengine.netdna-ssl.com/wp-content/uploads/3ways-FR.pdf>. Pg. 1.

⁴ «Lumina's Goal.» Lumina Foundation. Accessed March 01, 2019. <https://www.luminafoundation.org/lumina-goal#goal-issue>.

⁵ "Job Openings and Labor Turnover Summary." Economic News Release. March 15, 2019. Accessed April 2, 2019. <https://www.bls.gov/news.release/jolts.nr0.htm>.

⁶ "Job Openings and Labor Turnover Summary." Economic News Release. March 11, 2014. Accessed April 2, 2019. https://www.bls.gov/news.release/archives/jolts_03112014.htm.

⁷ «United States in Education at a Glance 2018: OECD Indicators.» OECD. 2018. Accessed March 01, 2019. <https://www.oecd-ilibrary.org/docserver/eag-2018-71-en.pdf?expires=1551328688&id=id&acname=guest&checksum=5AE59C4B846BA27BA39C9A5FFC68FDFF>.

⁸ <https://www.pewresearch.org/fact-tank/2017/01/18/u-s-still-has-a-ways-to-go-in-meeting-obamas-goal-of-producing-more-college-grads/>

⁹ "Lumina defines high-quality credentials as degrees and certificates that have well-defined and transparent learning outcomes which provide clear pathways to further education and employment." *Lumina Foundation's Strategic Plan*. Report. 2009. Accessed March 1, 2019. https://www.luminafoundation.org/files/wp-content/uploads/2011/02/Lumina_Strategic_Plan.pdf. Pg. 1.

¹⁰ "United States in Education at a Glance 2018: OECD Indicators.» OECD. 2018. Accessed March 01, 2019. <https://www.oecd-ilibrary.org/docserver/eag-2018-71-en.pdf?expires=1551328688&id=id&acname=guest&checksum=5AE59C4B846BA27BA39C9A5FFC68FDFF>. Pg. 2. See also, Periscope. «A Stronger Nation 2019: Lumina Foundation.» Page Redirection. Accessed March 01, 2019. <http://strongernation.luminafoundation.org/report/2019/#nation>.

¹¹ «Lumina's Goal.» Lumina Foundation. Accessed March 01, 2019. <https://www.luminafoundation.org/lumina-goal#goal-issue>

¹² Podcast. Solutions for Higher Education with President Scott L. Wyatt Episode 47—Innovation: Education and Preparation for the Workforce. Accessed March 1, 2019. <https://www.suu.edu/presidentpodcast/47/>.

¹³ *Today's Student*. Report. Accessed March 1, 2019. <https://www.luminafoundation.org/files/resources/today-s-student.pdf>. Pg. 2-4.

¹⁴ Ibid, Podcast. See also, Carnevale, Anthony P., Jeff Strohl, Neil Ridley, and Artem Gulish. *Three Educational Pathways to Good Jobs: High School, Middle Skills, and Bachelor's Degree*. Report. Georgetown University. 2018. Accessed March 1, 2019. <https://1gyhoq479ufd3yna29x7ubjn-wpengine.netdna-ssl.com/wp-content/uploads/3ways-FR.pdf>. Pg. 1.

¹⁵ *Today's Student*. Report. Accessed March 1, 2019. <https://www.luminafoundation.org/files/resources/today-s-student.pdf>. Pg. 6.

¹⁶ The Georgetown University Center on Workforce and Education defines "a good job as one paying a minimum of \$35,000 for workers between the ages of 25 and 44 and at least \$45,000 for workers between the ages of 45 and 64. This results in 2016 median earnings of \$56,000 for workers without a bachelor's degree, up from \$55,000 in 2015; median earnings of \$75,000 for workers with a bachelor's degree or higher; and overall median earnings of \$65,000 for all good jobs." Carnevale, Anthony P., Jeff Strohl, Neil Ridley, and Artem Gulish. *Three Educational Pathways to Good Jobs: High School, Middle Skills, and Bachelor's Degree*. Report. Georgetown University. 2018. Accessed March 1, 2019. <https://1gyhoq479ufd3yna29x7ubjn-wpengine.netdna-ssl.com/wp-content/uploads/3ways-FR.pdf>. Pg. 1.

¹⁷ Ibid Carnevale, Anthony P., pg. 10.

¹⁸ Ibid Carnevale, Anthony P. pg. 11.

¹⁹ Ibid Carnevale, Anthony P. pg. 4.

²⁰ Ibid Carnevale, Anthony P. pg. 18.

²¹ Obama, President Barack, and Jill Biden. «Remarks by the

President and Dr. Jill Biden at White House Summit on Community Colleges.» White House Summit on Community Colleges, White House, Washington, DC, October 5, 2010. <https://obamawhitehouse.archives.gov/the-press-office/2010/10/05/remarks-president-and-dr-jill-biden-white-house-summit-community-college>.

²² «Obama Administration Announces New Regulations to Strengthen Employment and Training Opportunities for Millions of Americans.» News release, June 30, 2016. The White House. Accessed March 1, 2019. <https://www.ed.gov/news/press-releases/obama-administration-announces-new-regulations-strengthen-employment-and-training-opportunities-millions-americans>.

²³ *Analysis of the White House National Job-Driven Action Plan*. Report. August 2014. Accessed March 6, 2019. https://m.nationalskillscoalition.org/resources/publications/file/2014-08-07_WhiteHouseActionPlan-NSCAAnalysis.pdf.

²⁴ Ibid Carnevale, Anthony P. Pg. 11.

²⁵ *Empowering Community Colleges to Build the Nation's Future: An Implementation Guide*. Report. 2014. Accessed March 1, 2019. http://www.aacc21stcenturycenter.org/wp-content/uploads/2014/04/EmpoweringCommunityColleges_final.pdf. pg 3.

²⁶ The author acknowledges the need for seamless pathways from “cradle to career.” The silos that exist between and among K-12, postsecondary education, and employers often make a student’s academic and career journey unnecessarily challenging. However, further exploration of this topic is beyond the scope of this paper.

²⁷ Freeman, Jennifer. “Know the Difference: Career Pathways vs. Guided Pathways.” JFF. January 22, 2018. Accessed April 24, 2019. <https://www.jff.org/points-of-view/know-difference-career-pathways-vs-guided-pathways/>. “Guided pathways are a framework for redesigning an entire community college to improve the student experience for everyone, from entry through graduation. Guided pathways initiatives involve: reforms in advising and registration processes to help all students identify their goals early in their academic careers, improvements in how students are supported throughout their courses of study, clear «maps» for each program of study that guide students as they select courses, and programs that, ideally, are aligned with career opportunities in the labor market.”

²⁸ “Energy Industry Competency Model: Generation, Transmission and Distribution.” Industry Solutions - Regional Implementation. Accessed April 23, 2019. <https://www.cewd.org/curriculum/downloads/EnergyCompetencyModelSeptember2010.pdf>.

²⁹ “Welcome to the California Community Colleges Chancellor’s Office.” California Community Colleges. Accessed April 2, 2019. <https://www.cccco.edu>.

³⁰ Stein, Tristan. *The Master Plan for Higher Education in California and State Workforce Needs: A Review*. Report. December 2018. Accessed March 6, 2019. http://opr.ca.gov/docs/20181226-Master_Plan_Report.pdf. Pg. 2.

³¹ Ibid, Stein. Pg. 1-4.

³² Baldassare, Mark, Dean Bonner, Sonja Petek, and Jui Shrestha. “PPIC Statewide Survey: Californians and Their Government.”

Public Policy Institute of California. March 2011. Accessed April 2, 2019. <https://www.ppic.org/publication/ppic-statewide-survey-californians-and-their-government-march-2011/>.

³³ CTE is used here to refer to career technical education

³⁴ Institute for Higher Education Leadership and Policy *Career Opportunities: Career Technical Education and the Completion Agenda* <http://edinsightscenter.org/Portals/0/ReportPDFs/career-opportunities-part-1.pdf?ver=2016-01-15-155401-970>

³⁵ California Community College Chancellor’s Office data. From the March 2013 CCCAOE presentation.

³⁶ “Doing What Matters for Jobs and the Economy.” EUpdate. July 2012. Accessed April 2, 2019. https://doingwhatmatters.cccco.edu/Portals/6/docs/July2012eUpdate_TheDoingWhatMattersFramework.pdf.

³⁷ “Welcome to Salary Surfer.” 2018. Accessed April 2, 2019. salariesurfer.cccco.edu. See also, “New Salary Surfer Website Shows Median Earnings for Those Who Earn Certificates and Associate Degrees at California Community Colleges.” Press Release. June 19, 2013. Accessed April 2, 2019. https://californiacommunitycolleges.cccco.edu/Portals/0/DocDownloads/PressReleases/JUN2013/PRESS_RELEASE_SalarySurferWebsiteLaunch_FINAL_061913.pdf.

³⁸ *Task Force on Workforce, Job Creation, and a Strong Economy*, presentation by California Community Colleges (version as of November 10, 2014), <http://doingwhatmatters.cccco.edu/portals/6/docs/SW/2-6-15%20BOG%20WF%20Task%20Force%20Regional%20College%20Slides.pdf>

³⁹ *BOARD OF GOVERNORS Task Force on Workforce, Job Creation, and a Strong Economy Report and Recommendations*. 2015. Accessed April 2, 2019. https://doingwhatmatters.cccco.edu/portals/6/docs/sw/BOG_TaskforceReport_BW_V1.pdf. Pg. 8.

⁴⁰ For the purposes of Table 1, “legislative” means action that requires enactment of a law. “Budget” refers to the incorporation of language into the California State Budget. “Regulatory” refers to policies within the authority of the Board of Governors of the California Community Colleges to change or originate. “Administrative” means that the CCCCO leadership could initiate the recommendation, usually through a change in practice or by issuing guidance to the field. Of these four, Legislative and Budget actions are the most complex to attain. “Strong Workforce Project Plan: Implementation Details.” 2019. Accessed April 2, 2019. <http://doingwhatmatters.cccco.edu/StrongWorkforce/ProjectPlan.aspx>.

⁴¹ The next largest workforce categorical for community colleges in the California State Budget was \$22.9M in-going funds for Senate Bill 1402 Economic and Workforce Development.

⁴² “17% Committee.” 2019. Accessed April 2, 2019. doingwhatmatters.cccco.edu/StrongWorkforce/17PercentCommittee.aspx.

⁴³ “Overview of the Student Centered Funding Formula.” 2018. Accessed April 15, 2019. https://extranet.cccco.edu/Portals/1/CFFP/Fiscal_Services/Budget_Workshop/2018/FundingFormula36.pdf

⁴⁴ “Amazon Web Services, 19 Los Angeles Community Colleges Collaborate to Launch a Regionally-Recognized Cloud Computing Certificate.” 2019. Accessed April 2, 2019. <https://laedc.org/2018/08/09/amazon-los-angeles-colleges-cloud/>.

⁴⁵ Wells, Shannon. "21st Century Skills Badging: A Bridge from Education to the Workforce." October 20, 2017. <https://medium.com/read-write-participate/21st-century-skills-badging-a-bridge-from-education-to-the-workforce-350f8ac688c4>.

⁴⁶ "CyberPatriot competition." Accessed on April 20, 2019. <https://www.uscyberpatriot.org/>.

⁴⁷ "Civic Stewardship to Strengthen Community Colleges." Graduate School of Education. March 2018. Accessed April 2, 2019. <https://ed.stanford.edu/events/civic-stewardship-strengthen-community-colleges>.

⁴⁸ "CTE/CE Rebranding Campaign." Information Resource Hub for Career Technical Education Faculty. 2019. Accessed April 2, 2019. <https://doingwhatmatters.cccco.edu/ForCareerEducationFaculty/CTERebranding.aspx>.

⁴⁹ "Program Mapper." California Community Colleges. 2018. Accessed April 2, 2019. <https://www.concentricsky.com/work/detail/program-map>.

⁵⁰ Booth, Kathy. "Moving the Needle: Data, Success, and Accountability for Workforce Programs." WestEd. March 2015. <http://doingwhatmatters.cccco.edu/portals/6/docs/SW/Task%20Force%20on%20Workforce%20Data%20and%20Accountability%20White%20Paper.pdf>.

⁵¹ «Career Education 2018 "Strong Workforce Stars" Programs Named as Enhanced Salary Surfer Website is Released to Help Students Make Decisions about College and Career." California Community Colleges Chancellor's Office. June 20, 2018. <https://careered.cccco.edu/wp-content/uploads/2018/07/PR-Strongworce-Stars-6-20-2018-FINAL.pdf>

⁵² "Career education programs can receive Strong Workforce Stars designation when meeting one or more of the following numerical thresholds.

- Fifty percent or more increase in student earnings, determined by comparing students' earnings one year before and one year after exiting the California Community Colleges system. This data is based on a match to the state's wage file.
- Seventy percent or more of students attain a regional living wage. This percentage, aligned with a threshold set by the California Workforce Development Board, analyzes the proportion of graduates and skills builders (defined as students whose goal is not to complete a program, but to take classes to increase job skills) who attained the living wage for a single individual in the college's region. This analysis is based on a match to the state's wage file and a comparison with data from the Insight Center for Community Economic Development.
- Ninety percent or more of students report that their current job is close or very close to their field of study. This is based on responses to the system's CTE Outcomes Survey, which is sent to community college students after they stop taking courses."

"Standouts in Career Education Earn Accolades as California's Community Colleges Advance Social and Economic Mobility". California Community Colleges Chancellor's Office. July 27, 2017.

<https://careered.cccco.edu/wp-content/uploads/2017/07/PR-STRONG-WORKFORCE-07-24-2017.pdf>.

⁵³ 2018 Strong Workforce Stars in the eShowcase. California Community Colleges Chancellor's Office. Accessed April 20, 2019. http://www.dwmshowcase.com/workforce_stars.asp.

⁵⁴ For example, see Scott, W. Richard, and Michael W. Kirst. *Higher Education and Silicon Valley: Connected but Conflicted*. Baltimore: Johns Hopkins University Press, 2017. See also, *LA Area Community Colleges at a Crossroads: A Market Assessment & Call to Action College Presidents' Preview*. Report. Los Angeles, CA: Los Angeles County Economic Development Corporation (LAEDC), 2019. 1-38.

⁵⁵ These lessons are adapted from the author's blog: <https://medium.com/@WorkforceVan/six-ways-colleges-can-think-differently-to-meet-the-workforce-needs-in-america-2499018cc13b>

⁵⁶ Broton, Katharine M., and Sara Goldbrick-Rab. "Going Without: An Exploration of Food and Housing Insecurity Among Undergraduates." *Educational Researcher*. December 07, 2017. Accessed April 2, 2019. <https://journals.sagepub.com/doi/abs/10.3102/0013189X17741303?ai=1gvoi&mi=3ricys&af=R>.

⁵⁷ Campbell, Anna. "Curriculum Alignment Planning For Heating, Ventilation, Air Conditioning and Refrigeration In California Community Colleges." May 26, 2016. Accessed April 2, 2019. <https://www.ecusectordwm.com/uploads/1/0/1/6/101615310/hvacr-curriculum-review-findings-v2-1.pdf>.

⁵⁸ "Apprenticeships in California." California Apprenticeship Initiative Hub. 2019. Accessed April 2, 2019. <https://caihub.foundationccc.org/Apprenticeships-in-California>.

Van Ton-Quinlivan led the workforce mission of the California Community Colleges, the largest system of higher education in the United States, from 2011 to 2018. She served initially as Vice Chancellor of Workforce and Economic Development and then Executive Vice Chancellor of Workforce and Digital Futures.

Supporting Data

Figure 1. Energy Industry Competency Model²⁸

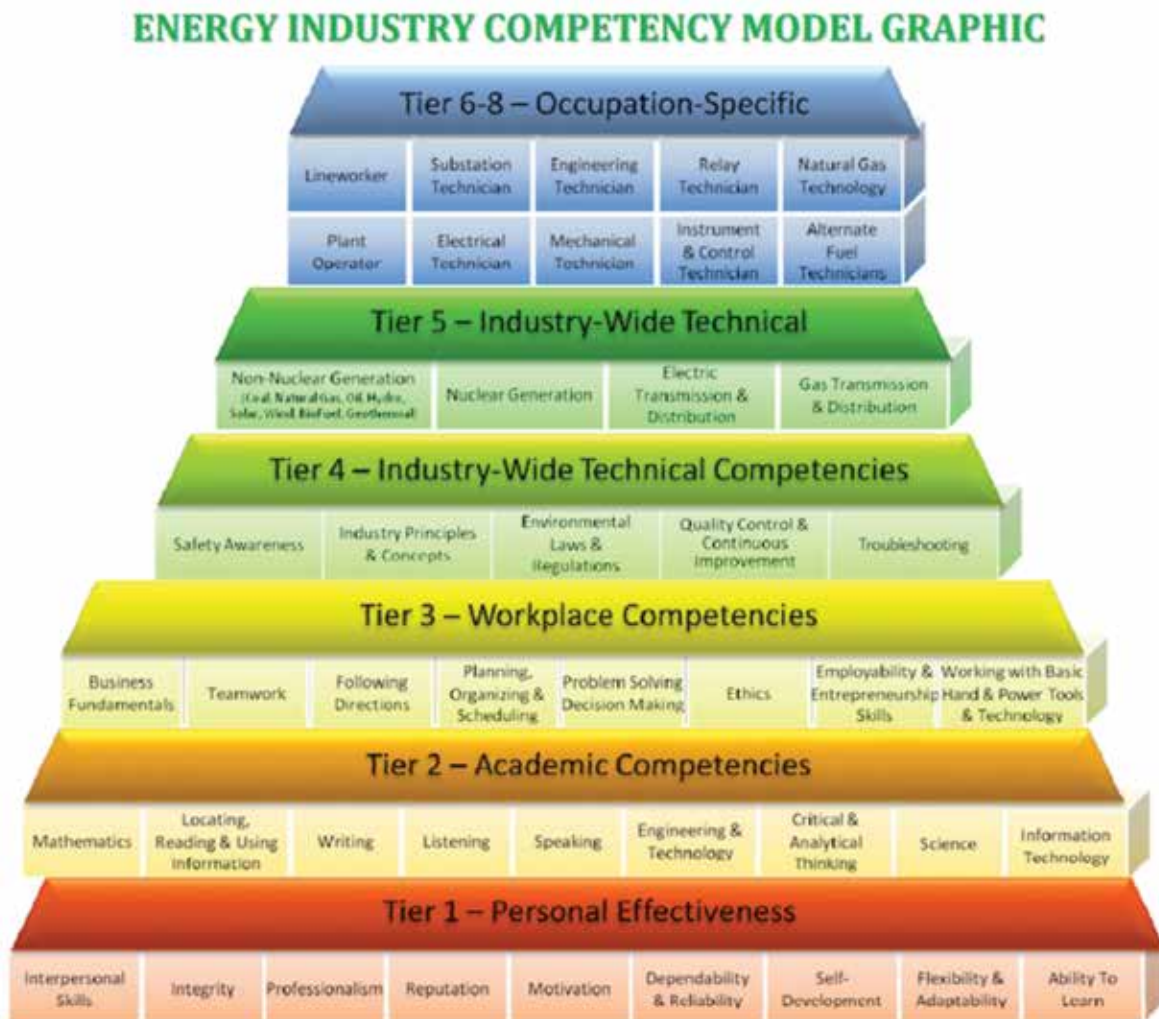


Figure 2. 10-Year Trend in California Community College's Career Technical Education (CTE) Portfolio (2001–2011)³⁵


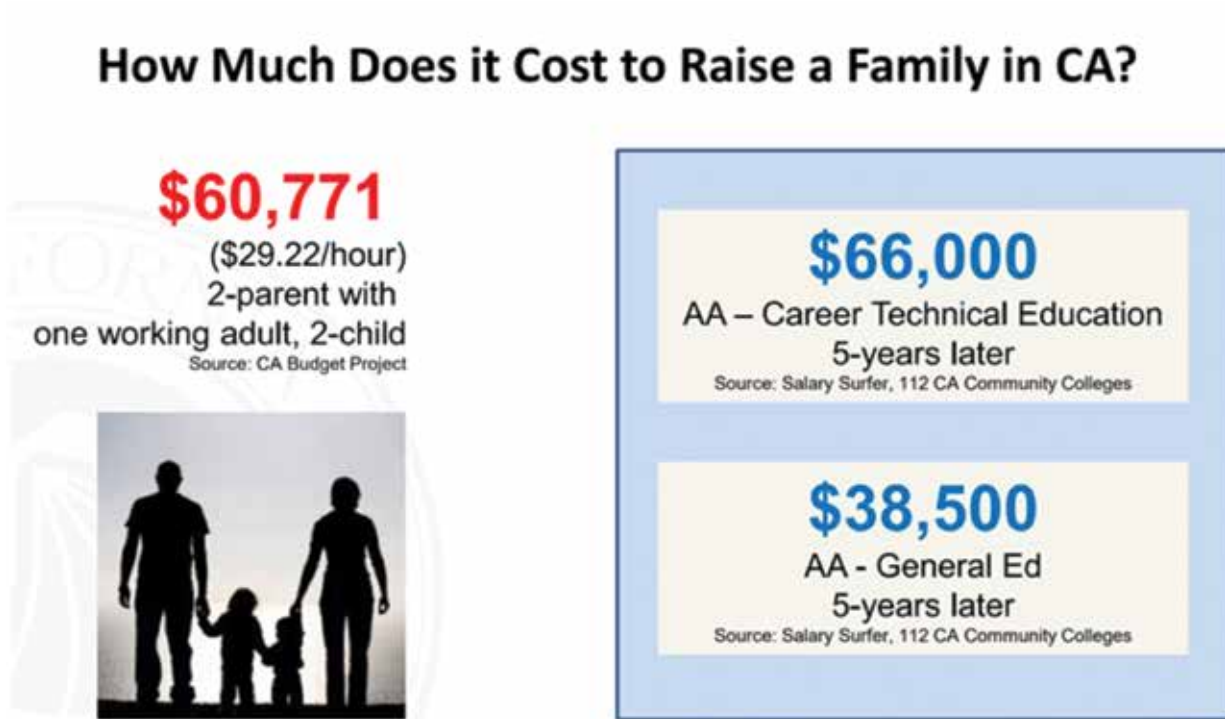


Figure 3. Cost of Raising a Family in California³⁸



Source: CCCC Strong Workforce Town Hall, <https://bit.ly/2Gg8GzQ>

Table 1. Task Force on Workforce, Jobs Creation, and a Strong Economy Recommendations

STUDENT SUCCESS	Type of Action
1. Broaden and enhance career exploration and planning, work-based learning opportunities, and other supports for students.	Administrative Budget Legislative
2. Improve CTE student progress and outcomes.	Administrative Regulatory
CAREER PATHWAY	
3. Develop and broadly publicize industry-informed career pathways that prepare students for jobs needed within the regional labor market.	Administrative
WORKFORCE DATA & OUTCOMES	
4. Create common workforce metrics for all state-funded CTE programs and expand the definition of student success to better reflect the wide array of CTE outcomes of community college students.	Administrative Legislative
5. Establish a student identifier for high school students and those enrolled in postsecondary education and training programs to enable California to track workforce progress and outcomes for students across institutions and programs.	Administrative Legislative
6. Improve the quality, accessibility, and utility of student outcome and labor market data to support students, educators, colleges, regions, employers, local workforce investment boards, and the state in CTE program development and improvement efforts.	Administrative
CURRICULUM	
7. Evaluate, strengthen, and revise the curriculum development process to ensure alignment from education to employment.	Administrative Regulatory
14. Consider options for meeting minimum qualifications to better integrate industry professionals who possess significant experience into CTE instructional programs.	Administrative Regulatory
15. Enhance professional development opportunities for CTE faculty to maintain industry and program relevance.	Administrative Regulatory
16. Explore solutions to attract industry professionals in high-salaried occupations to become CTE faculty in community colleges.	Administrative

REGIONAL COORDINATION	
17. Strengthen communication, coordination, and decision-making between regional CTE efforts and the colleges to meet regional labor market needs.	Administrative
18. Clarify and modify, as appropriate, state regulations to allow colleges to regionalize course articulation along career pathways utilizing regional or state curriculum models.	Regulatory
19. Develop regional leadership and operational partnerships among community college, industry, labor, and other workforce and economic development entities to improve the delivery of all CTE efforts.	Administrative Budget
20. Develop robust connections between community colleges, business and industry representatives, labor and other regional workforce development partners to align college programs with regional and industry needs and provide support for CTE programs.	Administrative
FUNDING	
21. Create a sustained, public outreach campaign to industry, high school students, counselors, parents, faculty, staff, and the community at large to promote career development and attainment and the value of career technical education.	Budget
22. Establish a sustained funding source to increase community colleges' capacity to create, adapt, and maintain quality CTE courses and programs that are responsive to regional labor market needs.	Budget
23. Create a predictable, targeted, and sustained funding stream that leverages multiple local, state, and federal CTE and workforce funds to support an infrastructure for collaboration at the state, regional and local levels; establish regional funding of program start-up and innovation; and develop other coordination activities.	Budget
24. Review, analyze, and modify, as needed, laws and regulations related to student fees for disposable and consumable materials and CTE facilities.	Budget Legislative

Observations from the Roundtable



Classical approaches can work. That was the message delivered by discussants at our roundtable on the interaction of emerging technologies with the domestic economy. Education, migration, and responsive regulatory policy were all offered as examples of policies that have worked before to help the United States economy take advantage of rapid changes while mitigating their disruptions. It's tempting to frame rapid technological change as an unprecedented challenge for this country, and one requiring unprecedented forms of governance. Similar arguments were, for example, to try to deal with the unexpected inflation of the early 1970s through "new methods" such as draconian economy-wide wage and price controls. Those failed spectacularly and sent the U.S. economy on a decade-long spiral. Our discussants therefore warned against throwing out orthodox policies for untried alternatives, as the result of doing so would be to replace one set of uncertainties—the complexity of the coming change itself—with two.

Productive Work

Contributor Erik Brynjolfsson described machine learning as potentially the most important technology of the generation. Three things drive its rapid ascent: 1) the mass digitization of data throughout the economy, more closely linking the data sets of our computers with the environment of people and our daily activities; 2) significantly better computing power that reduces the time needed to run computationally-expensive machine learning decisions by orders of magnitude (and more closely matching human decision timeframes); and 3) better algorithms that give better or faster results given some set of data. These parallel changes allow machines to effectively share work with humans, with each assuming the tasks it does best.

Brynjolfsson quoted Stanford AI pioneer Andrew Ng in expecting, "Anything today which can be done by a human in less than one second is well-suited to be done better by a machine." Since 2015, for example, machines surpassed (a general population of) humans in image recognition. Since 2017, voice-recognizing machines have approximately equaled human abilities on call switchboards. Applications are proliferating. When eBay in 2014 introduced fully automated, machine learning-based language translations to its Latin American marketplace listing titles, U.S. exports to the region's buyers increased by 11%, essentially overnight. So machine learning can improve the functioning of existing markets.

While some applications have burst on the scene very quickly, raising governance questions in the process, Brynjolfsson explained why potentially more fundamental changes will actually take some time. And the brake is not so much the technology as the need to redesign jobs, as some tasks in most occupations are suitable for machine learning, while others will continue to require human labor.

To that end, Brynjolfsson's paper aimed to tease apart one of the key governance questions around emerging technologies: the impact on employment. He described machine learning as being different from earlier types of automation, and it is not possible simply to extrapolate earlier experience to understand this new field. He described how machine learning would cut across wage and skill levels, with lower wage jobs disproportionately affected. But most jobs will be reinvented rather than eliminated.

First, Brynjolfsson argued the importance of seeing today's jobs as bundles of activities and tasks. A radiologist, for example, can be said to perform 27 distinct tasks as part of her work, some of which can be done better by machine learning, some of which cannot. A minority of radiologists of the future may spend more time teaching machines to do those particular tasks well. Most radiologists will spend less time on those tasks compared to other important parts of their jobs, such as patient interaction. Redesigning jobs will be key to machine learning productivity gains.

Secondly, employment is not a matter of slicing up a pie. The supply and demand of jobs and their tasks are dynamic. While machine learning and artificial intelligence can substitute for some human activities, it can also augment. Returning to the radiologist example, consider that a doctor using machine learning to augment her diagnoses could potentially diagnose more patients in a day, in fact making her a more valuable (that is, productive) employee than she was previously. Supply and demand elasticities could play a role as well, with hospitals able to offer more radiological services at lower prices, and consumers then deciding to take more CAT scans. Finally, there is the potential for new, utility-enhancing tasks that might emerge through invention and reengineering—using radiological

machine learning to perform wholly new types of diagnosis and monitoring for example, or automated remote care that simply wasn't available before, and so on.

Finally, Brynjolfsson argued that machine learning job impacts can be predicted, or at least, understood, by using a skills-based framework. This breaks down the problem into manageable pieces: we can enumerate the skills required in each job, and we can separately evaluate the actual progress of machine learning or other potentially disruptive technologies on each of those skills. This does not dictate a single policy response, but it does create a useful map to guide good governance efforts for policymakers, who are understandably concerned about impacts to their constituents and want to prioritize their efforts where most needed. And it should give the confidence to freely encourage the productivity-enhancing aspects of these technologies, rather than taking a defensive crouch in an attempt to prevent the effects from arriving on their own.

It's important to remember that future upside, which does not otherwise get to advocate for itself in today's policymaking.

A country's wealth is directly linked to the growth of its workforce multiplied by the growth in its worker productivity. Through this project we have seen how U.S. demographic trends are relatively healthy compared to other world powers. But troublingly, growth in worker productivity has stagnated since the 2008 recession. Emerging technologies such as machine learning or additive manufacturing are promising antidotes to this, but they must first be applied in the existing economy. To that end, our discussants described how U.S. business reorganizations, which are an indicator of firms becoming more efficient to make use of new technologies or market conditions, are actually happening slower today than they did 20-30 years ago. And machine learning may be no different, with an estimated ten-to-one ratio for in-firm worker reskilling and reorganization costs versus actual investments in machine learning IT. Upfront costs to such investments are high, with productivity gains following later. Such trends are not new: electricity, for example, was an obviously excellent technology, but it took 30 years for factories to see productivity gains from it given the need to strand existing (often steam-powered) assets while developing new skills, methods, and business models to fully take advantage of the benefits of the new technology.

This points to the need to encourage not just entrepreneurship, which is always welcome, but also an efficient labor market more broadly.

How to do this? Updating worker skills will be important, and applied education is a topic to which we will return. But part of an efficient labor market is the matching of skills and capacities to the needs of employers. As architect and urban planner Alain Bertaud has observed, in prison everyone has a job, and a short commute at that, but you could not argue that convicts are living their most productive lives. Flexible markets create the structure for such matching to occur naturally, through an emergent order that makes use of individuals' knowledge about themselves, their capabilities, and their preferences and the workforce needs of industry in ways a government planner never could. Discussants offered a few concrete recommendations to drive improvements in U.S. labor markets given coming disruptions.

One was for a reversal of the explosion in occupational licensing, which has grown to cover 25% of all American workers, up from just 5% in the 1950s. Extensive licensing requirements for trades, many of them in the service sector—everything from yoga instructors to physicians to fruit pickers—began with the justification of protecting public health and safety. Amid the decline of U.S. union membership, however, licensing is now used by trade lobbies of existing workers in a good profession as a means to explicitly block qualified new workers (whom they fear might over time put downward pressure on wages) from entering their field. Licensing makes it harder for someone to start a new line of work without unnecessary time and expense. Until recently, for example, Maryland required 1,500 hours of training to work in a blow dry salon. Arizona required 1,000 hours. Since this licensing is generally state based, it can also restrict worker mobility, effectively segregating the great diversity and expanse of the United States, one of our unique attributes, into 50 disjointed labor markets.

Improving that mobility across jobs, especially across geographies, will be another key challenge. Americans have long been a mobile people, and the idea of setting out for new fortunes in a new town or state is part of our psyche. And Brynjolfsson's work suggests that machine learning will have varying geographic affects across the country. But Americans have become less mobile. Since the middle of last century, the share of Americans who had moved in the previous year to a different location within the same county fell by half, in a steady, secular decline. The share that had moved from one county to another fell by more than one-third. And the percent of young adults—generally the

most mobile segment of the population given the need to establish households and establish careers—who have moved at all in the past year also fell by one-third since the 1960s. Too many Americans are stuck in areas or jobs that are not using their full potential, or in no job at all.

Our discussants described some of the theories of why this reduction in mobility may be happening, and potential governance strategies to address it.

Issues included an over-subscription to government welfare programs, such as disability-based social security, which create perverse incentives to stay on the programs instead of seeking new employment. Rolls for disability in the United States grew from 2.7 million in the mid-1980s to nine million in 2014 alongside an ageing population and a widening definition of eligibility, despite workers on average reporting feeling healthier and jobs becoming less physically-demanding. That gives a receiving individual less incentive to find new work given the “effective marginal tax rate” they would incur in doing so—and the potential loss of existing benefits, including monthly payments and Medicaid. Reforms to more strictly define beneficiaries to those who are unable to do work of any kind—as opposed to being unable to perform their existing profession—will be important as the entire U.S. population shifts towards more dynamic employment environments. Private employer benefits can also be a barrier to those who may wish to switch jobs but fear losing health insurance, particular physicians, or other built-up entitlements in doing so, pointing to the need for better benefit portability.

Other impediments to increased mobility and more efficient U.S. labor markets are less obvious. High student loan debt, for example, may be keeping even well-educated young people at home living with parents. Census results show that the share of young adults living with their parents stayed steady from 1990 until the mid-2000s but exploded thereafter, from 11.6% in 2005 to 22.0% in 2017. This is a bid to reduce cost of living, but it may also reduce their potential to gain income from finding the best jobs available to them, with career-long earning impacts. Another factor may be the prohibitive cost of real estate in the nation's most productive urban regions, such as the San Francisco Bay Area, Los Angeles, Seattle, Boston, Washington, D.C, or New York City, which now disproportionately produce jobs in excess of the suburban or rural areas, which once dominated. Many vibrant U.S. urban areas (with the notable exception of Houston) suffer from decades of housing development restrictions and high construction costs that have suppressed the supply of new housing to well-below job growth. This drives up rents and drives down home-ownership (down by one-sixth among young adults since the year 2000). Some have argued that it also drives down marital rates (down to 40% from 55% over the same period) and, eventually, fertility, with long-term demographic implications. One recent study by economists Chang-Tai Hsieh and Enrico Moretti estimated that such restrictions since 1964 have reduced the size of the U.S. economy by one-half. Now, one-half is probably an extreme number, but it points to the foundational importance of a well-functioning labor market. This is also why it is no bad thing that the salient threat of disruptions from artificial intelligence and other emerging technologies now prods policy makers to publicly revisit what may otherwise be considered an esoteric topic.

And of course all of this process of technology implementation will take human creativity. Amazon.com reinvented the bookstore (and other stores) not by substituting machines for humans in shelving and checkout, but by changing everything from the supply chain process to the end consumer, down to the location and operation of warehouses that could offer millions versus thousands of products. Going forward, the process of creative destruction will require rethinking entire businesses around taking advantage of machine learning technology. Brynjolfsson offered a startling prediction to technology bears: even if machine learning technology's progress froze today—no more announcements from Google's research teams or from Stanford's computer science labs—the U.S. economy would still see decades of innovation on business practices that would improve aggregate productivity. Put another way, the much-publicized spread of artificial intelligence is in fact firms applying existing machine learning technology to an ever-widening expanse of industries and problems. This innovation will require lots of intangible investments. Consider that Americans across a number of firms have already spent \$100 billion on the development of vehicle self-driving technology, even though no driver has yet been replaced. Machine learning is not a loaded gun. In fact, the bullet has already been shot.

A word on social bias in emerging technologies, the subject of substantial roundtable discussion and disagreement. Discussants agreed that machine learning and other algorithmic decision-making systems of the sort now regularly used by both internet companies and other institutions, including the government, are biased, reflecting bias in the data used to train the systems, in the design of algorithms, and in the interpretation of results. Internet companies, for example, collect personal information, create and refine behavioral profiles of individuals, and develop algorithms

that curate content, all with the objective of maximizing attention and commercial profits. So it is often in the firms' economic interest to discriminate. And discussants readily accepted that such decision-making systems can perpetuate or even introduce new bias. Some forms of bias are illegal, such as when ad buyers target housing options in certain locations to one race and in other locations to another race. Legal discrimination might include showing quality investment opportunities to one class of users and shady schemes to another. And bias can also be an issue in many other contexts as well outside of internet advertising or content curation, such as machine learning systems that support decisions on hiring, mortgages, insurance availability and rates, and sentencing guidelines.

The disagreement was on the scale and importance of such bias within an overall system, what might be done to effectively counter it, and the collateral costs of doing so. While some argued that any amount of unintentional or illegal bias is problematic, and saw potential for machine learning decision-making systems without the biases common to human decision-making, others pointed to the known biases of humans in our own decision-making, which suggests comparing machine learning systems to a status quo (flawed) human baseline rather than an idealistic one of perfection.

And whereas some participants suggested (federal) government agency regulation or Congressional lawmaking as ways to provide strong enough enforcement that powerful technology companies would actually comply in removing bias, others warned that involving government agencies, who have less information and technical expertise available to them than private firms, could actually end up making bias problems worse in a fast-changing technological environment. These experts instead offered "the regulation of the marketplace" as a preferred alternative, whereby if one firm were to provide poor or otherwise undesirable services, they could naturally lose users to rival firms that did better. The technological landscape is littered with such corpses. Hewlett Packard, once a byword for dominance in scientific and engineering hardware, as well as high-risk research, rapidly fell amid the explosion of the low cost "PC clone" business in the 1990s. IBM, Kodak, Sun Microsystems, AOL, Yahoo, Blackberry, MySpace: each enjoyed seemingly unassailable product positions, often up until the point that those positions became obviously doomed as they were outflanked in unpredictable ways.

Finally, and more generally, is concern about unintended consequences. Those brought up in Silicon Valley have an innate sense that while the incredibly innovative technologies that were developed and popularized there were not necessarily done without government help, they certainly benefitted from a being largely "left alone" to do business, and repeatedly change how they do business, amid an ever-shifting technological landscape. The result of that has been a mix of good and bad, as with every industry, but the business and consumer products used around the world every day to great utility and at low cost argue that the good has been overwhelming. There is a fear then that a jump to government intervention could break that system, either by actually serving to better entrench the positions of those tech incumbents through regulatory capture, or by weakening the overall ecosystem's attractiveness for investment. Our discussants furthermore observed how the reach of attractive internet technologies to willing users has proved longer than the reach of any one nation's domestic regulatory arm. The American parents of smartphone wielding middle schoolers may be concerned about the content of their Facebook newsfeeds, or Snapchat advertising, and press regulators to do something about it—but what of the China-based, adolescent-targeting short video app TikTok, which has grown to 500 million users worldwide, including 100 million U.S. downloads, since its release two and half years ago? This was probably not the sort of trade that those with real concerns on privacy or freedom of expression have in mind. If we are banking on new technologies to enable broad productivity gains in an emerging economy, then we should at every step consider the social costs of limiting that against any expected benefit. Monitoring and careful deliberation is in order.

Technical Education

The choices that American students make about acquiring skills through their educations, and that workers make about learning new skills while on the job or between them, underpins what bundles of "tasks" they will be able to productively perform in a changing economy. 65% of current U.S. job openings require some level of post-secondary skills, and our discussions of state and local community colleges as increasingly important institutions for providing applied education to a diverse spectrum of Americans supported that. Former vice-chancellor of the California Community Colleges Van Ton-Quinlivan reported that five years after completing a 2-year "career technical education" program at a California community college, a worker makes an average salary of \$66,000. Five years after completing a 2-year general education associates degree the average salary is just \$38,500. At the same time, employers report seeking a broad array of general skills in addition to occupation-specific or technical skills—critical thinking, problem solving, language and effective communication, teamwork.

Education in an emerging new world will not be a matter of funneling students into today's "recession-proof" jobs (which may see novel challenges from emerging technologies), or of focusing on STEM education. Rather, the goal should be to produce graduates who have specific skills that meet the needs employers are looking for today, and a broad enough framework for overall learning that they can successfully return to the education system again and again throughout their careers to quickly acquire new skills as the task bundles change.

Why our interest in two-year community colleges? First, they already exist and go relatively unnoticed in policy dialogues that jump between the deep dysfunction of the American K-12 system, and this country's relatively high performing—but very expensive—four-year university system. California's 115 school community college system, our discussants noted, is likely the largest higher education system in the country. Of the state's two million unfilled jobs, half require a four-year college degree, but half need less than that. Second is that in federal, state, and local governance environments, where budgets are likely to be increasingly crowded out by compulsory spending items such as health care entitlements and pensions, community colleges remain focused and cost-efficient. They are often located close to home with yearly price tags of \$3,000-\$6,000 versus many multiples of that for a longer (and often residential) four-year option. This is good for government budgets, and it is good for students who can better avoid loan debt traps. Third is their track record at educating a diverse range of students. In California's two-year institutions, for example, 60% of students are women, 37% are 25 or older, 24% have children or other dependents, 31% are from families in poverty, and 64% work—40% of those full time. Their customer base more closely matches the profile and needs of mid-career students of the future.

Today's students seek a return on their educational investment—good jobs to support themselves, provide for their families, pay their loans, and save for retirement. And community colleges with career technical educational offerings are attractive to them for their ability to make available occupations that are known to be productive and locally in-demand: nurses, emergency medical technicians, welders, utility line workers, plant operators, or maintenance and repair technicians, for example. 80% of California graduates stay within their own region to find work, and local community colleges are able to partner with nearby employers, individually or perhaps more effectively in regional coalitions, and in doing so keep pace with their evolving needs by mixing and matching task-specific modules with general skills such as English, math, and social reasoning.

Our discussants noted the increasing importance of two-year community colleges remaining nimble as emerging technologies drive acceleration across the economy and society itself. California's system provides examples of agile programming, such as pooling resources across smaller colleges to build effective collaborations with regional employers such that they see it in their direct interest to interact with students through teaching, internships, and curriculum development, or minimizing the bureaucracy that can slow the roll-out of new curricula.

We also identified the opportunity for community colleges—with their relatively short "business cycles"—to more directly engage with both the employers and high schools, who will be providing their next crop of students, so as to reduce friction in the handoff. Employers can increase graduation rates, for example, by front-loading tuition reimbursement instead of paying employees back after they have incurred the costs. Or community colleges can expose their curricula and major options to high schoolers and their teachers to telegraph future career options. The latter idea was described as being particularly important for minority students who, once in the community college system, tend to select familiar but generally less-productive (and lower-earning) areas of study: in California, the top major for Latinos is early childhood education, and for African Americans it is social work. Ultimately, the country's emerging workforce needs will be met through self-responsibility as students and workers are exposed to incentives to learn, and community colleges will be a key infrastructure in enabling them to execute on those choices.

Of course reforming America's K-12 education itself is of paramount importance. Discussants were proud of the fact that America does have great primary and secondary schools. But they lamented that it also has terrible schools that hold back their graduates' achievements and earnings for life. And their quality is based largely upon their zip codes. For students who are driven, community and four-year colleges find themselves completing the remedial teaching that high schools failed to deliver. Again and again our project has identified the importance of a strong basic education in reaping the benefits of an emerging new world while avoiding the worst of its pitfalls. We will revisit this topic at another session in more detail.

New Arrivals

Since the 2008 recession, there has been much popular and academic discussion of economic inequality in the United States, with the spotlight on "the 1%" and billionaires like Jeff Bezos or Mark Zuckerberg. But focusing on the

top misses the point. Jeff Bezos helped to create new growth that benefitted others, too. Instead, if you are really concerned about people's well-being, you should focus on the bottom (How are the least well-off Americans doing over time? Are they improving?), and on the middle (Are wages of those in the middle of the income distribution going up or down or sideways?).

And to that question, history would suggest that rapid economic growth rates are among the best ways to benefit Americans at the middle and lower rungs of society. When unemployment rates fall to their lowest levels and labor markets tighten, social groups with generally higher unemployment rates—women, minorities, the elderly, the less-educated, those with criminal records—are disproportionately pulled into the workforce or are able to upgrade from existing jobs. The converse is also true during a downturn: “last hired, first fired.” During the depths of the 2008 recession, black employment rates were falling at 5-6% per year, nearly double the 3-4% rates for whites. But since 2012, black employment rates have recovered at 3-4% per year, versus approximately 1% for whites. Similarly, unemployment rates for workers with a college degree were 10% lower than those for high school dropouts in 2009, but after nearly a decade of continuous economic expansion that spread had fallen to just over 3%. So as a general rule this points to the importance of economic performance as a key enabler of social equality.

In an emerging world, however, our concept of labor markets is expanding. Globalization set off the freeing of capital flows around the world, allowing for huge cross-border investments and speculation. And it ushered in the global movement of goods as complex multi-country supply chains were established to take advantage of beneficial regional attributes and trade soared. Outside of elites, however, globalization did not fully unlock the third leg of the economic stool: labor. Cross-border services can be considered a form of labor mobility, and the value of trade in services is gradually rising versus that of trade in merchandise, particularly in advanced countries like the United States. But what of the movement of people themselves? We know a strong and dynamic economy pulls in marginalized workers in our own country's labor market—and increasingly we see this now pulling in workers from other countries as well, both within or outside of the legal frameworks in place for that.

Our roundtable discussants considered the history of and continued role for immigration in the United States from this economic perspective.

In 2017, about 258 million people around the world (3.4% of the population) lived outside their country of birth. And countries consider multiple factors as they consider how many immigrants to admit, with what skills, with what status. These include security, cultural and ideological concerns, economic interests, and rights. Migration expert Jim Hollifield offered three such frameworks across the unique U.S. migrant history. The “Massachusetts model,” which welcomed immigrants on the basis that they assimilate to the host culture. The “Pennsylvania model,” which treated newcomers essentially equally given a baseline respect for local law and basic values. And the “Virginia model,” which focused on bringing immigrants for labor.

The United States has applied those models to varying degrees across four major waves of immigration, and the sources of new arrivals has changed as well through those waves. Before 1820, English and Scots dominated—and Africans were forcibly immigrated through slavery. Between 1840 and 1870, economic motivations drove Irish, German, and Scandinavian migrants, including many Catholics. From 1880 to 1914, Chinese migrants went to the western United States for work and to escape upheaval at home, while southern and eastern Europeans went to the East Coast, Midwest, and Southwest. Finally, from the 1970's to the present, migration has included both low-skill and high-skill workers from Mexico, Central America, and Asia. At each step, immigrants reshaped American society, and they played an increasingly important role in the economy. And each wave drew a reaction from other Americans, some of whom were directly economically impacted by their arrivals, and others who may have only considered themselves to be.

The foreign-born population in the United States today has grown by almost five times since 1970, and it has regained its peak share of nearly 14%, last seen at the turn of the 20th century. Discussants described how immigrants are increasingly important for economic growth, providing both labor and human capital. Today's immigrants, for example, are going to the states where the highest economic growth is. High-skilled immigrants such as engineers and scientists, nurses and doctors are a boon to the economy. And importantly, immigration now provides 30% of U.S. population growth—without immigration the U.S. population would have already stagnated and started to decline, like so many other countries an emerging world.

The most pressing immigration issue today is of course the 12 million people living in breach of U.S. law, half of whom have entered illegally and half of whom have overstayed visas in this country. An inability to enforce immigration

laws poses challenges to state sovereignty and security, the legal system, and civil society broadly. At the same time, the country cannot deport millions of people who wish to be here, some of whom have already undertaken deep struggle or sacrifice to try to become Americans. So this creates a paradox. To maintain economic competitiveness, the United States should keep its economy open to trade, foreign investment, and immigration, but it also needs to control its borders so as to not undermine the social contract and rule of law. Clearly there is no black and white answer to this, and it calls for balance.

Our discussants sketched out what an updated immigration strategy might look like, arguing that immigration policy can be both compassionate and “greedy.” Currently, for example, 80% of legal immigrants in this country come through the “family channel”—so-called “chain-migration”—where one legal immigrant is able to easily sponsor other family members abroad, some of them distant, to later immigrate to the United States as well. American public opinion is relatively forgiving of this practice, and it has been credited with helping ensure a social safety net for new immigrants who may otherwise struggle in a foreign society. But it does not necessarily do a good job of meeting the host country’s goals in terms of desired skills or attributes, and it also is not fair to other desiring immigrants, who may actually have a stronger benefit in coming to this country but no family member to help them.

Canadians deal with this issue through a multifaceted point system to balance this with other interests and approximately 60% of Canadian immigrants are now considered to be “economic” versus 30% family. The United States could do something similar—perhaps focus the family channel on the nuclear family only—while also considering members with certain demographics or capabilities. This would be similar to President George W. Bush’s idea of “matching willing workers with willing employers.” While experiences then and today demonstrate the political difficulty of getting a sweeping deal through, our discussants were optimistic that with the right leadership, and perhaps by staging reforms into smaller pieces to get the ball rolling, the underlying fundamentals support good prospects for a deal. As one of us has said about Washington, “Sometimes when you know that you are right about something, but everyone tells you it can’t be done, you just keep on with it and eventually the fundamentals will prevail.”

And the upshot is that this is not new territory. The history of U.S. immigration has never been clean. Yet we remain an immigrant nation in actuality and in self-image. And while those who watch acrimonious arguments over this playing out daily on cable news may find it a surprising claim, we should not have a crisis of confidence in our ability to handle immigration issues. The fact is that this is difficult, but we in the United States—alongside similar immigrant nations like Canada (with a 22% foreign-born population) and Australia (24% foreign-born population, not including the UK)—are actually the best globally at handling this issue because of our unrivalled experience. Globally, migration pressures will grow in an emerging world. Ubiquitous information and communication flows will make information about and interactions with foreign countries and citizens easier. Automation and advanced or additive manufacturing can untangle supply chains and their workers. Changing climates may set people alight from their existing homes or work, fleeing unsuitable areas or seeking new opportunities. Military conflicts through new weaponry may do the same. These flows will need to be managed. More importantly, they will need to be governed as they integrate into the broader society.

Our existing American diversity is a boon to this integration as it offers a “golden dome” under which new arrivals can find any number of suitable ways to live their lives as Americans. It also means that our governance institutions and procedures were designed from the beginning to encourage counter-balancing expressions of diversity around a common stake in the American creed.

In fact, the United States at its founding was probably the most conscious historical effort to set up a government over diversity. The historical background is instructive. Emerging from the Revolutionary War, the country faced a variety of internal regional challenges and remained more “states” than “united” under the loose Articles of Confederation. At the same time, the populace was heavily composed of immigrants and lacked a strongly hegemonic social default. James Madison and other founders in their contributions to the *Federalist Papers* wondered how the destabilizing tendency of various factions “united and actuated by some common impulse of passion, or of interest” would be managed in the hard-won Union.

In his review of the subject, Madison, for example, saw diversity—“a division of the society into different interests and parties”—as “sown in the nature of man.” As a part of human nature, it would therefore be impossible to remove its causes, whether through oppression or through consensus. In any case, the young state would not be powerful enough to do so even if it wanted to. Instead, the American answer to managing diversity—as diversity is inevitable—would be more diversity:

The diversity in the faculties of men, from which the rights of property originate, is not less an insuperable obstacle to a uniformity of interests. The protection of these faculties is the first object of government. [Federalist 10]

This argued for creating a framework to allow as much diversity as possible. Not redistribution, not seeking of a common denominator, not compensation, but protection of the abilities of its citizens to express diverse interests. Allowing as many diverse interests as possible within a large and expanding Union would naturally create a political and social system more robust to dominance by any single faction or against the spread of extremism. Were the state to limit any one interest, however, it risked unleashing another. The development of a Constitution (and Bill of Rights) was therefore their way of distributing power so that this diversity could be recognized. Moreover, over time, and not without misstep, the country has learned how to effectively govern over that diversity.

The very structure of a limited federal government based on checks and balances and the ability of state and local governments to have regional authority over matters closer to home, which affect peoples' lives most directly, allowed the nation as a whole to maintain and represent a diversity of opinion, resilient against domination by any one geographic interest or extremist fad. In addition, through today, these "laboratories of democracy" encourage experimentation in governance, the voluntary formation of *ad hoc* relationships, and help to improve government performance through a form of regional competition. Of course, protections of religious freedoms was also central to the early identity of the country at a time when the Church of England's monopoly on faith demonstrated the impossibility of forced commonality. As Thomas Jefferson reflected in his own letters, "Divided we stand, united we fall."

The country did not always uphold these values. Women's rights and later civil rights for blacks nearly split the nation. Even then, parts of society continued to try to keep the lid on and oppress the interests of millions of U.S. citizens for decades. When the pressure that built up eventually boiled over in the 1960s, it served as a stark lesson of the continual failure to effectively govern over this diversity. Though this has been recognized in the years since and was enshrined in the Constitution, it is worth considering the years of opportunity lost for not only the oppressed but also the nation as a whole had the value of this diversity been earlier enabled.

To return to our topic at hand, a diversity mindset can also be argued to have applied to the development of the U.S. economy. Acknowledging that some will be more economically successful than others, and allowing them to personally benefit from the value they create (while at the same time protecting the least well-off in society through a safety net), has meant that Americans have long had the chance to be rewarded for their own risk-taking entrepreneurship. Unlike some other modern societies, the existence of wealth in the United States—and a shared opportunity to realize it—is generally regarded as beneficial and not something morally corrupt to be appropriated by the state or stamped out through excessive redistribution. This underlying sense of responsibility has preserved a strong incentive for self-betterment across a diverse society. And it will help Americans at all strata of society to find creative ways to take advantage of this century's emerging new technologies for their personal and community benefit.

The United States of course remains a nation of immigrants, and the foreign-born share of the population since records were kept in the 19th century has stayed in a consistent 11–14% range. As an example close to home here in the San Francisco Bay Area, nearly as many residents of the region were born outside the United States as were born in California, and English is the primary language at home of less than half its residents. It may be surprising to learn that native-born Americans make up less than one-third of Silicon Valley tech workers. But somehow, it all works. Regional labor productivity now exceeds \$200,000 per employee and has grown 50% faster than the U.S. average since the turn of the century, and its share of U.S. patents granted has doubled. Other countries with less modern experience in this realm are now figuring out how to get there. The going can get rough, and U.S. immigration today has problems, too. But to take the longer viewpoint, this is in our gut. America's ability to incorporate newcomers while maintaining the diversity they brought with them is unrivaled, and that will help us to make the best of—and provide needed global leadership in—what may be an even more chaotic global migration landscape going forward.

About

New and rapid societal and technological changes are complicating governance around the globe and challenging traditional thinking. Demographic changes and migration are having a profound effect as some populations age and shrink while other countries expand. The information and communications revolution is making governance much more difficult and heightening the impact of diversity. Emerging technologies, especially artificial intelligence and automation, are bringing about a new industrial revolution, disrupting workforces and increasing military capabilities of both states and non-state actors. And new means of production such as additive manufacturing and automation are changing how, where, and what we produce. These changes are coming quickly, faster than governments have historically been able to respond.

Led by Hoover Distinguished Fellow George P. Shultz, his Project on Governance in an Emerging New World aims to understand these changes and inform strategies that both address the challenges and take advantage of the opportunities afforded by these dramatic shifts.

The project features a series of papers and events addressing how these changes are affecting democratic processes, the economy, and national security of the United States, and how they are affecting countries and regions, including Russia, China, Europe, Africa, and Latin America. A set of essays by the participants accompanies each event and provides thoughtful analysis of the challenges and opportunities.



The publisher has made an online version of this work available under a Creative Commons Attribution-NonCommercial license 3.0.

To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0>.

Efforts have been made to locate the original sources, determine the current rights holders, and, if needed, obtain reproduction permissions. On verification of any such claims to rights in the articles reproduced in this book, any required corrections or clarifications will be made in subsequent printings/editions. Hoover Institution assumes no responsibility for the persistence or accuracy of URLs for external or third-party Internet websites referred to in this publication, and does not guarantee that any content on such websites is, or will remain, accurate or appropriate.

Copyright © 2019 by the Board of Trustees of Leland Stanford Junior University

**Hoover Institution,
Stanford University**
434 Galvez Mall
Stanford, CA 94305-6003
650-723-1754

Hoover Institution in Washington
The Johnson Center
1399 New York Avenue NW, Suite 500
Washington, DC 20005
202-760-3200

