TECHNOLOGY AND NATIONAL SECURITY

MAINTAINING AMERICA'S EDGE

Introduction by Joseph S. Nye, Jr., Condoleezza Rice, and Nicholas Burns
Edited by Leah Bitounis and Jonathon Price

CONTRIBUTORS INCLUDE:

Edward Alden, Nicholas Burns, Ash Carter, Jack Clark, Jared Cohen, John Deutch, John Dowdy, Diana Farrell, Walter Isaacson, Chandru Krishnamurthy, Katherine Mansted, Jason Matheny, Joseph S. Nye, Jr., Penny Pritzker, Condoleezza Rice, Eric Rosenbach, and Laura Rosenberger

The Aspen Institute
2300 N Street, N.W.
Suite 700
Washington, DC 20037

Published in the United States of America in 2019 by The Aspen Institute

Printed in the United States of America
ISBN: 978-0-578-42795-9
Wye Publication Number: 19/001

Cover design by: Steve Johnson
Interior layout by: Sogand Sepassi

aspen strategy group

Aspen Strategy Group Leadership

CHAIR EMERITUS

Brent Scowcroft
President
The Scowcroft Group, Inc.

CO-CHAIRS

Joseph S. Nye, Jr.
University Distinguished Service
Professor Emeritus
Harvard Kennedy School

Condoleezza Rice
Senior Fellow
Hoover Institution
Stanford University

DIRECTOR

Nicholas Burns
Goodman Family Professor of the Practice
of Diplomacy and International Relations
Harvard Kennedy School

DEPUTY DIRECTOR

Jonathon Price
Deputy Director
Aspen Strategy Group

SPECIAL PROJECTS OFFICER

Leah Bitounis
Special Projects Officer
Aspen Strategy Group

ASPEN INSTITUTE PRESIDENT

Daniel Porterfield
President and CEO
The Aspen Institute

MEMBERS

Madeleine K. Albright
Chair
Albright Stonebridge Group

Graham Allison
Douglas Dillon Professor of Government
Harvard Kennedy School

Christian Brose
Staff Director
Senate Armed Services Committee

Sylvia Burwell
President
American University

Kurt Campbell
Chairman and CEO
The Asia Group, LLC

Ash Carter
Director
Belfer Center, Harvard Kennedy School

James Cartwright
Inaugural Chairholder,
Harold Brown Chair
Center for Strategic and
International Studies

Eliot Cohen
Robert E. Osgood Professor of Strategic Studies
Johns Hopkins University's School of Advanced International Studies

Richard Cooper
Professor of Economics
Harvard University

John Deutch
Professor
MIT

Tom Donilon
Chairman
BlackRock Investment Institute

Diana Farrell
President and CEO
JPMorgan Chase Institute

Peter Feaver
Professor
Duke University

Dianne Feinstein
Senator (D-CA)
United States Senate

Michèle Flournoy
Co-Founder and Managing Partner
WestExec Advisors

Mike Green
Senior Vice President for Asia
Center for Strategic and International Studies

Stephen Hadley
Principal
RiceHadleyGates LLC

John Hamre
President and CEO
Center for Strategic and International Studies

Jane Harman
President, Director, and CEO
The Woodrow Wilson International Center for Scholars

David Ignatius
Columnist and Associate Editor
The Washington Post

Anja Manuel
Principal
RiceHadleyGates LLC

Jessica Mathews
Distinguished Fellow
Carnegie Endowment for International Peace

David McCormick
Co-CEO
Bridgewater Associates

Sam Nunn
Co-Chair
Nuclear Threat Initiative

Meghan O'Sullivan
Jeane Kirkpatrick Professor of the Practice of International Affairs and Director of the Geopolitics of Energy Project
Harvard Kennedy School

William Perry
Professor
Stanford University

Dina Powell
Managing Director
Goldman Sachs

Penny Pritzker
Chairman and Founder
PSP Partners

Tom Pritzker
Executive Chairman
Hyatt Hotels Corporation

Jack Reed
Senator (D-RI)
United States Senate

Susan Rice
Distinguished Visiting Research Fellow
School of International Service,
American University

Carla Anne Robbins
Faculty Director, MIA Program,
Baruch College
Adjunct Sr. Fellow,
Council on Foreign Relations

Susan Schwab
Professor of Public Policy,
University of Maryland
Strategic Advisor, Mayer Brown LLP

Smita Singh
Former Founding Director of the Global
Development Program
William and Flora Hewlett Foundation

Anne-Marie Slaughter
President and CEO
New America

Jim Steinberg
Professor
Syracuse University

Dan Sullivan
Senator (R-AK)
United States Senate

Fran Townsend
EVP, Worldwide Government,
Legal and Business Affairs
MacAndrews & Forbes, Inc.

Dov Zakheim
Senior Fellow
CNA

Philip Zelikow
White Burkett Miller Professor of History
University of Virginia

Robert Zoellick
Chairman
AllianceBernstein

Acknowledgements

This book is a product of many individuals who have generously given their time and talent to produce what you hold in your hands today.

First, we thank the chapter authors who not only wrote papers, but came to Aspen to share their ideas, have them challenged, edit their papers accordingly, and produce the stellar chapters that make up this book. They are the true experts, and we are grateful for their willingness to share their expertise with our group and the public.

Second, we thank our Aspen Strategy Group members who give their time every summer to come together in a nonpartisan manner to read, discuss, learn, and contribute solutions to the most pressing problems facing America. In a time of unrelenting partisanship, it is heartening to see that at the seminar table, politics does indeed stop at the water's edge. All our Republicans, Democrats, and independents are committed to finding the best way forward to ensure America's interests and secure its future.

Third, we thank the Aspen staff who spent many hours turning these papers into the latest edition of our policy book series. Gayle Bennett, our long-time editor of ASG publications, goes through every page of this book to ensure the highest quality possible—we appreciate your excellence. Steve Johnson and Sogand Sepassi do a masterful job on layout and design—thank you. And of course, our superb team of Scowcroft Fellows this summer—Brian Hopkins and Megan Lamberth—your dedication to this project demonstrates the very bright careers you have in front of you.

Fourth, we are extremely grateful to all of the Aspen Strategy Group's friends, sponsors, and funders who support our work to bring these ideas into the community. We are indebted to the Clermont Foundation, the Markle Foundation, Robert Rosenkranz and Alexandra Munroe, The Stanton Foundation, McKinsey & Company, the Margot and Thomas Pritzker Family Foundation, Robert Abernethy, the JPMorgan Chase Institute, Leah Joy Zell, The John Anson Kittredge Educational Fund, Lynda and Stewart Resnick, Dick and Gail Elden, Gail Engleberg, Adam Metz, and the Feldman Foundation. Without their support, this book would not have been possible.

Finally, we simply could not do what we do without our superb co-chairs Joe Nye and Condi Rice. Joe, who founded this group over three decades ago, and Condi, who was first brought to this group by Brent Scowcroft, moderate the conversations, ask hard questions, and serve as the very definition of what our public servants should aspire to be—we are grateful for your leadership.

As we navigate these uncertain times, it is reassuring to know that such a large community is committed to maintaining America’s edge and securing its future.

Contents

Introduction
Navigating Uncharted Territory in the Technological Era

Joseph S. Nye, Jr.
Co-Chair,
Aspen Strategy Group

Condoleezza Rice
Co-Chair,
Aspen Strategy Group

Nicholas Burns
Director,
Aspen Strategy Group

In August 2018, the nonpartisan Aspen Strategy Group (ASG) convened its thirty-fourth annual meeting in Aspen, Colorado. Over the course of three days, ASG members and invited experts from government, universities, think tanks, and the private sector debated the impact of dramatic technological change over the next decade on American national security. Our conversations covered a wide breadth of emerging technologies—artificial intelligence, machine learning, quantum computing, and biotechnology—and the challenges they pose to America's military, the intelligence community, U.S. economic power, and democratic institutions. Our group grappled with the central dilemma of how the U.S. government can harness these technologies—developed primarily in the private sector and research labs—to compete with China and other adversaries in the years ahead.

How do we reinvigorate America's innovation triangle of government, the private sector, and academia, which was so crucial in creating the country's global technological lead over the past half-century? How do we prepare the American workforce for the perpetual disruption of technology and the automation of a growing number of jobs? What steps must we take to modernize the U.S. defense apparatus and military to face future threats? How should the U.S. government respond to cyber and disinformation tactics perpetrated by foreign actors? Most critically, if the U.S. fails to respond to these dynamic changes, will it lose its place as the world's strongest military power in the years ahead?

These questions frame this book as they will our national debate on the future of America's power in the world. Our authors explore them further in the subsequent chapters. While these issues have been around for some time, the pace and scale of the change make it ever more urgent. Recognizing technology's tremendous influence on the key pillars of American power, we left Aspen with three overarching observations of how the U.S. government must respond to this transformative era.

I. The Necessity of Defense Modernization

One of the most pressing national security issues is military modernization and ensuring the U.S. defense structure is prepared for these dramatic future challenges. This modernization requires a unified, sustained vision that encompasses not only the U.S. Defense Department and military, but the Congressional leadership in both parties, the current and future Presidential administrations, and the private sector. The Administration must reevaluate the U.S. defense budget and procurement processes, rethink its human capital strategy, and face the looming ethical considerations of lethal autonomous weapons systems.

The U.S. Defense Budget and Procurement Process

In June 2018, the Senate approved one of the largest defense budgets in modern American history—nearly $716 billion. The major challenge for our country, however, is how to spend it wisely. The U.S. military continues to allocate much of its annual budget to building, operating, and maintaining large, manned, and expensive systems, while a minute percentage of the budget—less than 1 percent—is allocated toward science and technology research. Some of our participants argued that the defense budget must begin to reflect America's future military needs. Military modernization means shifting away from expensive, manned systems toward large quantities of smaller, dispersed systems that are unmanned, expendable, and as autonomous as possible. The Pentagon needs to incorporate emerging technologies, or it risks losing its advantage among near-peer competitors.

The defense procurement system must also be modernized to reflect the rapid pace of technological advancement. The U.S. cannot wait for years to acquire readily available technology in the private sector. The Pentagon must build a sustained partnership with the technology sector, including companies in Silicon Valley. Open dialogue is key to establishing trust between these two juxtaposed worlds. For the Defense Department, this relationship is imperative. The Pentagon must collaborate with industries to acquire the latest, most innovative cyber and artificial intelligence (AI) technologies.

The major tech firms in turn, must be willing to work with the Pentagon to ensure our military has the benefit of the most advanced technological break throughs in AI, quantum computing, and other areas.

Human Capital Strategy

The future needs of the U.S. military go beyond the procurement of advanced technology. The military does not yet have nearly enough technical specialists in its ranks to lead the Pentagon into the digital age.

The Defense Department must create pathways and incentives for young, talented technologists and engineers to join its ranks. The Defense Department could institute, for example, a talent exchange program with private sector technology firms, similar to the initiative former Secretary of Defense Ash Carter, one of our ASG members, proposed in 2016. This would give commercial sector technologists a sense of the Defense Department's mission and day-to-day operations.

This effort, ongoing in certain departments within the U.S. defense structure, must be urgent and widespread. The Defense Department needs to prepare its ranks to ensure the continued security of America's military personnel and the nation itself.

Ethical Considerations

The issue of maintaining a "human in the loop" is a more nuanced debate. The military already deploys systems with a high degree of autonomy. The Defense Department can and should work with private firms to ensure ethical considerations are embedded in the planning for unmanned systems. This process will be ongoing, and the U.S. government should continue to debate these ethical considerations as artificial intelligence and machine learning technologies advance and permeate every facet of America's defense structure.

II. The Future Character of Conflict

Our discussions in Aspen touched heavily on the nature of future conflict and the potential responses of the U.S. government to these emerging threats. Beyond the threat of cyberattacks from state rivals and non-state actors, the U.S. will continue to face an onslaught of devastating disinformation campaigns and hybrid warfare tactics of the type developed by the Russian government against the U.S. during the 2016 Presidential election and since then. The U.S. government should work with American allies and other international partners to counter information operations and cyber threats and defend against these multi-dimensional attacks.

Warfare in the Cyber Domain

To prepare for continued conflict in the cyber domain, the U.S. government must adopt a cyber strategy that includes both rules of engagement for offensive measures and strategies for defending the country's national security and critical infrastructure. This task is made more difficult by the rapid diffusion of technology. The U.S. must presume that whatever technology it develops will eventually become available to all, including less powerful nations, non-state actors, and even individuals.

As the prevalence of cyber conflict continues to surge, the U.S. must establish clear declaratory policies, which outline the country's measured response to cyberattacks and information operations, to deter aspirant cyber attackers.

China has stated its intent to be the leader in A.I. in the near future. Washington should adopt an ambitious program to respond to China's digital aspirations. Tensions between the two nations will endure, and America cannot afford complacency. The U.S. must establish a clear, long-term vision for AI and cyber development—a response to China's "Made in China 2025" strategic plan. This vision should encompass not only the Defense Department and military, but commercial technology firms and research institutions as well.

Collaborating with International Partners

To adapt and thrive in this challenging new environment, the U.S. must rely on its wealth of international partners and allies. By cooperating and sharing data, the U.S. and its allies can begin to learn from and counter cyberattacks perpetrated against them. The U.S. should lead NATO in developing strategies for deterring cyber and information operations, ensuring not only its own security, but also the security of its allies. NATO must ensure that each member country fulfills the objectives set out in Article III—that all nations will "maintain and develop their individual and collective capacity to resist armed attack." This resilience must include the threat of cyberattacks and hybrid warfare tactics.

NATO allies should also continue to discuss the Article V threshold in the cyber realm. At what point will NATO member countries consider a cyber attack against one of them to be an attack against all of them? These questions are not theoretical. Hybrid warfare, executed by countries such as Russia and China, will continue to evolve and advance, pushing the limits and boundaries established by NATO member countries.

III. Reinvigorating the Innovation Triangle

Walter Isaacson's chapter in this book outlines the history and strength of America's innovation triangle. This alliance of government, universities, and the private sector was essential to the country's technological advancement since World War Two. In the midst of dramatic technological shifts, America must rely once more on this whole-of-nation approach.

Why is this triangle not functioning as it has in the past? The problem is not a lack of innovation in the U.S., but a failure of meaningful and sustained collaboration among the three pillars of the triangle. Deep mistrust permeates the relationship between Washington and Silicon Valley. The first step is to bridge these two disparate worlds. The second step is reinvigorating federal funding for research and development. The government must increase federal funding for university research labs to ensure continued American innovation and basic science research. Third, and most critically, the U.S. government must propose a strategic vision for cyber and emerging technologies that encompasses the other two pillars of the triangular alliance. This vision should include input and ideas from government officials, Congressional leaders, leading commercial sector technologists, academics, and researchers. A whole-of-nation approach requires a whole-of-nation strategy.

Future of Work

The innovation triangle is also pivotal in preparing the American workforce for continued economic disruptions caused by technology and automation. An increasing number of jobs will be subjected to automation, and Washington must work with the private sector and research institutions to help Americans adapt and thrive in this evolving workforce. Retraining workers for the digitized economy is vital, and must be vigorously endorsed by federal and state governments and widely implemented by employers. This long-term challenge goes far beyond just the economic ramifications—it poses a threat to the health and security of our civil society. An alliance of government, the private sector, and research institutions is necessary to prepare the country and its workforce for this future.

Maintaining America's Edge

This is still the dawn of the digital era. While cyber technology and artificial intelligence are in full public view, they are merely examples of what is to come. It is imperative that the U.S. government respond to the changing nature of military technology. Military modernization must be a priority, and the defense budget must reflect the military needs of the future, both in equipment and personnel. This sustained challenge was summarized by Chairman of the Joint Chiefs of Staff Joseph Dunford, who warned Congress: "In just a few years, if we do not change our trajectory, we will lose our qualitative and quantitative competitive advantage."

In the essays that follow, technologists, career government officials, private sector leaders, and scholars dive deep into these critically urgent issues. While each chapter touches on a unique piece of the technological revolution, the central message is clear: to maintain America's technological edge for the future, we must act now.

Ash Carter's Ernest May Memorial Lecture opens this volume with historical examples of disruptive technological change, before exploring possible solutions to the dilemmas we face today.

Jason Matheny and Jack Clark's essays introduce specific digital and biological technologies that are shaping important instruments of American power and the capabilities and risks associated with these perpetually evolving landscapes.

While we strive to create the environment for successful technological innovations, we also need to understand how our home front must adapt to resolve the problems innovation creates as digitization, machine learning, and A.I. affect American lives. Diana Farrell, Edward Alden, and Penny Pritzker dedicate their pieces to understanding how these changes are affecting the Future of Work so that we can create inclusive economic growth and a sustainable labor market.

John Dowdy and Chandru Krishnamurthy delve into the future nature of conflict and how emerging technologies are fundamentally altering core defense missions and challenging existing thinking on autonomous systems. They argue for urgent modernization of our defense procurement system and a more effective partnership between the private sector and the Federal government.

Beyond the battlefield, technologies also pose serious threats to democracies and institutions around the globe. Malicious cyber and disinformation campaigns are creating a new level of hybrid warfare and "digital war zones" that we must also tackle, as Jared Cohen outlines in his chapter. Laura Rosenberger, Eric Rosenbach, and Katherine Mansted offer their proposals on the best ways for democracies to counter and survive this age of disinformation.

It is necessary to look at the larger context of the global balance of power when discussing America maintaining its technological edge. John Deutch's chapter examines America's greatest competitor in this space—China, with its well-funded, cohesive, and long-term innovation strategy *China2025*. His chapter offers insights on how the U.S. should reimagine a clear, multifaceted innovation plan to remain competitive without undercutting our own openness and values.

Finally, Walter Isaacson argues that the underpinnings of our new strategic long-term innovation policy can be found by remembering how America first achieved its technological edge—through strong working relationships between the U.S. government, universities, and the private sector. His chapter reminds us that for the U.S. to remain on the forefront of technological developments, it must prioritize significant investments in the ground-breaking research that this requires.

Our hope is that the ideas and policy recommendations presented in this book will amplify the lessons we learned during our time in Aspen, moving them beyond our seminar table and into the classrooms, labs, offices, and homes of those dedicated to maintaining America's national security technological edge. Our future depends on it.

“The right decisions will not be made without strong input from technologists themselves.”

–ASH CARTER

The Tenth Annual Ernest May Memorial Lecture
Shaping Disruptive Technological Change for Public Good

Ash Carter
Director
Belfer Center for Science and International Affairs
Harvard Kennedy School

Editor's Note: Ash Carter presented the annual Ernest R. May Memorial Lecture at the Aspen Strategy Group's August 2018 Summer Workshop in Aspen, Colorado. The following are his remarks delivered at the meeting. The lecture is named for Ernest May, an international relations historian and Harvard John F. Kennedy School of Government professor, who passed away in 2009. ASG developed the lecture series to honor Professor May's celebrated lectures.

I am very pleased to give the Ernest May lecture for two reasons. The first is that I urged Joe Nye, Condi Rice, and Nick Burns to make technology the theme of an Aspen Strategy Group gathering, even recognizing the equally momentous international and political currents of interest to this group today.

The second reason is that Ernie May was a colleague and a friend—and a historian. As you'll see, I come at this subject from the perspective of my own origins as a scientist. But it's unrealistic to expect leaders in the real world to use their own knowledge of science very much in policy making—or to use economics, political science, or even philosophy. Instead the dominant mental methodology of real policy makers is historical reasoning. Ernie emphasized as much in his seminal book with Dick Neustadt, *Thinking in Time: The Uses of History for Decision Makers.* He would, therefore, have approved of a lecture that begins with an effort to use history to illuminate the seemingly very ahistorical topic of disruptive technological change.

First let me say that I use "disruptive" in both its good and bad connotations. Disruptive scientific and technological progress is not to me inherently good or inherently evil. But its arc is for us to shape. Technology's progress is furthermore in my judgment unstoppable. But it is quite incorrect that it unfolds inexorably according to its own internal logic and the laws of nature. My experience and observation is that this is true only directionally. Which specific technologies develop most quickly is heavily shaped by the *mission* that motivates and rewards the innovators: improving health, selling advertising or some other service, cheap energy, education, or national defense, for example. Making "disruption" more good than bad is the topic, as I understand it, of this year's Aspen Strategy Group.

A little personal history leads into what I want to say about technological history. I began my career in the field of subatomic physics, and the elders in that field were all of the Manhattan Project generation or the very immediate aftermath. Mentors of mine were Sidney Drell, Edward Teller, Hans Bethe, Richard Garwin, and others. It was their example, in fact, that made me interested in the consequences of science for public purpose. The culture that they inculcated in those of my generation stressed that along with the great ability to make change came great responsibility. It was that culture of science at the time of my upbringing that ultimately drew me into the service of national defense, and finally to secretary of defense. After thirty-seven years of continuous service of one kind or another with the Department of Defense during the administrations of presidents of both parties, I was thinking last year about what to do next. I decided that aligning technology with public purpose and solving some of the dilemmas I'll describe in this lecture was the most consequential issue of our time, second only to protecting ourselves and creating a better world for our children, and that that's how I would spend my time.

That Greatest Generation was proud to have created a "disruptive" technology: nuclear weapons. It had ended World War II and deterred a third world war through almost fifty years of East-West standoff. But the flipside of that coin was an existential danger to humanity. Recognizing both bad and good, those same scientists—coming at it from various ideological and political directions—accordingly devoted themselves in the years following to developing arms control and nonproliferation as new fields of innovative endeavor, to missile defense and civil defense, to making strong contributions to the intelligence systems that were needed to monitor arms control agreements, and to reactor safety to make the accompanying revolution in nuclear power safer. This is the culture that I knew.

The generation of leaders that came shortly thereafter was very, very different. The tech culture, including what is most associated with Silicon Valley but is actually pervasive in digital tech (though not the rest of tech), grew out of the hippie and counterculture movements. This is a very different kind of social impulse. It is inherently distrustful of government and believes that public good and public purpose will somehow emerge through a popular and supposedly freer mechanism.

I won't pretend to understand or share this ethos, but it is still the prevailing one among not only the founders, but many of the employees of the tech companies today. It shows in some of the troubling dilemmas that I'll turn to later.

Another characteristic of that long-gone era of my upbringing is that in those days most technology of consequence arose in America, and most of that from or within the walls of government. Neither is true anymore. Technology today is commercial and global. That creates an entirely different context for the pursuit of public purpose. (I prefer to use the term public *purpose* instead of public *policy*, because public policy suggests actions of government. In matters of technology today, as in the atomic age, solutions require unified effort of the tech community and government.)

A consequence is that some of the moral guidance to steer us to a good technological future will need to come directly from entrepreneurs and companies. The right decisions will not be made without strong input from technologists themselves. That is what originally convinced me to work on defense problems. I realized that many of the key issues during the Cold War had a strong technological component, and they could not be addressed well without the input of people like me. Big issues and a chance to see your training make a difference are a powerful attractive combination to a young technologist.

This being the case, I'm happy to say that today's generation is very different from the second generation I described. I see it every day at Harvard and MIT and did likewise during my time at Stanford. There's a strong demand for instruction and guidance on how to contribute to public purpose. Many of these young people are not looking at going into government, but they are looking to do something more consequential than get people to click on ads.

I discovered that I was able to tap into this same reservoir as secretary of defense. I always said that as secretary of defense I was "the secretary of defense of today" and also "the secretary of tomorrow." Secretary of today meant standing strong against Russia and China; deterring and defending ourselves, allies, and friends from North Korea and Iran; and destroying ISIS and other terrorists in Iraq, Syria, Afghanistan, and around the world. Secretary of tomorrow meant ensuring we had the people, strategies, and technologies to continue making ours the world's finest fighting force. I wasn't sure I could succeed when I embarked on my so-called outreach to the tech community, beginning by founding a Pentagon outpost in Silicon Valley, the Defense Innovative Unit-Experimental (DIU-X), which we subsequently replicated in both Boston and Austin. I would have established additional outposts in more tech hubs were I still secretary of defense, and I hope Jim Mattis does. Despite the Snowden hangover, I found that there was a hunger among most of the tech company employees to be part of something bigger than themselves and their firms. I found great uptake through DIU-X and also through the Defense Digital Service, which allowed technologists to come and go right in the halls of the Pentagon with their hoodies on and aviator glasses on their foreheads. I am particularly proud of the Defense Innovation Board I instituted, which included senior leaders like Eric Schmidt (to whom I'm grateful for serving as chair), Jeff Bezos, Reid Hoffman, Jen Pahlka, and others. All this reflected my principle that technologists and the tech industry were essential to achieving the important public purpose of national security.

This outreach to the wider technology community was an essential complement to the big funding impulse we gave to the DoD research and development budget in the so-called "third offset," and the huge strategic reorientation we were making from fifteen years of counterterrorism and counterinsurgency to the big-ticket, full-spectrum threats associated with Russia and China. At some $80 billion per year, DoD's R&D effort is more than twice Google's, Microsoft's, and Apple's R&D combined.

The other defining experience for me were the wars. I was undersecretary for acquisition technology and logistics during the big Afghan surge of 2010, and I found that alongside the F-35 Joint Strike Fighter, the KC-46 aerial refueling tanker, and the other big traditional programs I had to manage, my daily preoccupation was making sure that the troops had everything they needed to win and protect themselves. That meant new kinds of Mine-Resistant Ambush Protected (MRAP) vehicles for Afghanistan, persistent surveillance like aerostats, all kinds of techniques to counter improvised explosive devices (IEDs), and things that you may not associate with the "weapons czar": buying dogs, ballistic underwear, and so on. Nothing was too small, nothing too inconsequential. Every day the wars were Job 1, and I make no apologies for that. Whatever you think of the wars, when the kids are out there, you have to be all in.

It wasn't enough during war to carry out the usual ten-year defense program. You had to do a ten-week program, even a ten-day program. I began thinking about how we needed to change, not only to serve the wars that are, but the wars that *might* be at any moment. We need to make sure that we don't have regrets if we get in a dustup with Iran, for example—we need to give them a bloody nose and make sure they don't give us a bloody nose in the first few days. We don't want to look back and say, "I wish I had done something; I wish we had done something that we could have done but that we didn't do because we were on the old Cold War tempo."

• • • • •

I described the post–World War II technology cultures in the US. Going back even further in history, the great transition that everyone, especially economists, loves to study is the farm-to-factory migration. This is often described as a success story, and, in retrospect, it surely must be so regarded. Hundreds of millions of people changed fundamentally their way of life while the means of production moved from individual artisanship to collective mechanized effort, and, for the most part, their lives were much better in the end. At the same time, it looks better in the rearview mirror than it must have at the time. Don't forget that the farm-to-factory movement took decades to sort out. It is not clear that the pace of change today will give us that kind of time to make momentous technologically driven adjustments.

The farm-to-factory migration was also pretty rocky if you think about the rise of communism, the formation of urban ghettos, and other speedbumps that were less than minor, and only if you forget also that the transition failed miserably in some countries, notably Russia.

Above all, the success was not at all automatic—far from the work of the invisible hand. In the United States and Britain, there emerged from the bleakest period of the Industrial Revolution the Progressive and Chartist movements, which by introducing regulation of commerce, foods, and medicines made large-scale, widespread, anonymous, non-artisanal production and distribution of goods acceptable, since it was no longer possible for a person to know who made something they consumed or where it was coming from. The list goes on: child labor laws, compulsory public education, boards of public health, the Sherman Antitrust Act of 1890, muckraking journalism, labor unions, and so on. In short, the farm-to-factory transition was paralleled and made a success in this country not by laws of technology or economics alone, but by a host of non-technical innovations that set the conditions for overall public good.

One way to pose the week's topics is therefore: How do we set the conditions for today's disruptive changes to redound to the overall good of humankind? How might the tech communities contribute to solving some of the big dilemmas of today's looming disruptive change in the three big categories: digital, biotech, and jobs and training?

• • • • •

There are so many digital dilemmas: offensive and defensive cyber, big data, augmented reality, quantum computing, internet of things, and others, but let me touch on two: social media and artificial intelligence.

Social media are wonderful enablers of commerce and community, but also of darkness, hatred, lies, and isolation; invasion of privacy; even attack. I, therefore, had much higher hopes for the Facebook hearings before Congress, featuring CEO Mark Zuckerberg. Hearings are a way of calling the public's attention to the tech dilemmas and paving the road to a solution. In the case of the Facebook hearings, there was no need to call attention: 91 percent of Americans, according to a recent Pew survey, feel that they've lost control of how their personal data is collected and used, and two in three think current privacy laws are not sufficient.

In terms of leading to solutions, however, the hearings laid an egg. They missed entirely what was a historic opportunity to devise what everybody seemed to acknowledge is needed: a mix of self-regulation by tech companies and informed regulation by government. Zuckerberg, for his part, gave an account of his company's ethical conduct that sufficed for one news cycle, but will not, I fear, suffice at all for the great arc of history. As for the quality of the congressional questioning, well, all I can say is that I wish members had been as poorly prepared to question me on war and peace in the scores of testimonies I gave as they were when asking Facebook about the public duties of tech companies! But make no mistake, we need to land this plane.

Ernie May might have advised us to look back a little bit upon history's analogous dilemmas. How might the members have been better informed to prepare their own way in the Zuckerberg hearing? It's not that this issue, or any of the ones I'm discussing in this lecture, has become particularly partisan. Who would have thought there was another form of gridlock in Washington?

One of my early Washington jobs was for an organization called the Office of Technology Assessment (OTA). It was the fourth congressional support agency next to the Library of Congress's Congressional Research Service, the General Accounting Office, and the Congressional Budget Office. OTA did high-quality work for members on exactly subjects like this. It would have prepared a report in consultation with Facebook, other media companies, tech experts, lawyers, lobbyists, and so on and tried to put together options for that combination of self-regulation and regulation that was the underlying consensus solution in the hearings. OTA was eliminated during the Gingrich revolution as part of the effort to downsize government. The other three congressional support agencies were big and powerful and could defend themselves, but little OTA got the axe.

I also remember the Senate Arms Control Observer Group when I worked for Paul Nitze, who was President Reagan's chief arms control advisor. They not only met with the administration regularly, but also had a panel of experts from that post-atomic era of scientists who advised them on technical matters such as verification of arms control agreements, nuclear effects and civil defense, antiballistic missile systems, and survivable-basing modes. Then there has long been a committee of scientists who advise the select intelligence committees on the super-secret optical, radar, signal, infrared, and other satellite programs used for intelligence purposes.

So, in short, Ernie might say that there is in living memory the idea of bipartisan outreach to, and reliance on, tech expertise.

As we think, in the manner of Ernie May, about precursors that may be models for new institutions to join tech and public purpose today, another one that comes to mind is the National Security Telecommunications Advisory Council (NSTAC). In my very first job in the Pentagon in 1981, one of my office's duties was to lead Cap Weinberger's battle against the breakup of AT&T. In retrospect, we resembled the last Japanese soldier on Saipan in World War II, still charging about the jungle unaware that the emperor had surrendered. NSTAC was established, in part, at Weinberger's insistence once the AT&T breakup became inevitable to make sure that the deregulated system continued to serve the public interest.

The breakup of AT&T was, in fact, an episode in a long history of communication and information system regulation. This history begins with the US Postal Service, a natural communications monopoly that the government managed. When the telegraph came along, the US government decided against absorbing it into the postal service as most European governments did. There followed a period of vigorous competition, which ended in a Western Union monopoly—as it had to for a natural monopoly—and was regulated accordingly.

Western Union remained a regulated monopoly, but its fear of more regulation probably was a factor in discouraging it from getting into telephone communications, leaving that field to AT&T, which also drifted toward natural monopoly. Some of the same concerns may have applied to AT&T's decision not to move into radio but instead content itself with carrying its programs over its long-haul lines. When NBC Radio became too big, it was forced to split into NBC and ABC. The Nixon administration gradually relaxed strictures on cable, in turn, to challenge the major broadcast networks.

So, there is an abundant history of antitrust or other government regulation applied to natural monopolies of information and communication. Ernie might remind us that this history could have inspired some kind of productive output from the Zuckerberg hearings, such as regulation based loosely on antitrust to handle Facebook's monopolization of its form of social media. Some economists argue that since Facebook and Google are free, no economic harm can be shown to the consumer by the government, and, therefore, the government has no antitrust authority. This interpretation would be alien to both Senator Sherman, of the Sherman Antitrust Act, and Justices Brandeis and Douglas, who wrote the early opinions. They repeatedly stressed that the government's interest was in the general public good and was not confined to price gouging.

Here is Justice Douglas: "The philosophy and the command of the Sherman Act is founded on a theory of hostility to the concentration in private hands of power so great that only a government of the people should have it."

Here is Justice Brandeis: "… the maintenance of competition does not necessarily involve destructive and unrestricted competition, any more than the maintenance of liberty implies license or anarchy."

So with a little of Ernie May's history in mind, and returning to technology, join me in what Einstein called a thought experiment. What would be different algorithmic approaches to social media curation and delivery, and how might they reflect the public good? You can imagine a number of them. One algorithm would organize digital platform content by maximizing advertising and platform revenue. This is essentially the prevailing model. A second would reflect individual choice, offering what you seem to want based on your past patterns. There is some of this in Facebook and other feeds in order to promote the ends of the first algorithm.

A third algorithm would stress the crowd, that is, what everybody else seems to be watching, what is "trending." A fourth might be profit-based, but share profit with the owner of the data in another form of subscription-free service. A fifth channel you might dial would have content curated by professional journalists—the elusive Campbell Brown at Facebook. Another possibility is multiple demonopolized competing platforms that use whatever model they choose. My concern about a simple breakup of Facebook into smaller Baby Bell-type offspring is that they will only end up competing to represent the lowest common denominator, and we will have a worse outcome than we do now.

The best world to me would be one where there are multiple channels representing these different algorithmic models, and the consumer could simply switch from channel to channel and shop, compare, and pay accordingly, with the content of all subject to some rules written by a public commission that went beyond simple strictures on terrorism, child pornography, and the like. When I watched *I Love Lucy* as a child and Lucy and Ricky prepared to go to sleep at night, they got into twin beds separated by a nightstand with a lamp on it. That was regarded as appropriate to protect decency and children.

So thinking historically and conceptually, there are a number of possibilities and mixes that might have emerged from the Zuckerberg hearings. But nothing did. Ernie

May would probably have regarded the Facebook hearings as one of those potentially seminal historical moments that was wasted.

Turning to artificial intelligence: in my last year as secretary of defense, the question I would get most often in a wide-open press availability would be about "autonomous weapons." I would remind people that way back in 2013 or 2014, I had promulgated a directive on that subject that governed the conduct of the Defense Department as it developed the technology of artificial intelligence. It stated that for every system capable of executing or assisting the use of lethal force, there must be a human being making the decision. That is, there would be no literal autonomy. So that is how things stand on the books.

I was motivated to do that by imagining myself standing in front of the press the morning after, let us say, an airstrike that had mistakenly taken the lives of women and children. Imagine further that I tried to assign responsibility by saying, "The machine made a mistake." I would be crucified. So also will be the designer of a driverless vehicle that kills a little old man and cannot explain. Judges simply aren't going to accept anything other than an accounting of human responsibility.

I believe that accountability and the transparency to promote it are the key issues for the designers of artificial intelligence systems today. Now there are some who will tell you that the AI system they have developed simply does not enable the tracing of the method of decision that underlies an algorithm's recommendation. In almost four decades of working on technology projects, I've heard that many times from engineers about the difficulty of incorporating some desired feature or another that they haven't bothered to include in their design. My retort to these scientists is: if you want your algorithm to be adopted, you had better make it transparently accountable. If this requires an adjustment in design, which I can well imagine it does, then make that adjustment.

Before I leave the subject of artificial intelligence, I need to say something about the Google employees who resisted working on artificial intelligence for the US Department of Defense. I imagine what I would say to them in a Google town hall or if I were the Google leadership. I'd tell them they should think about and reconsider their decision. First of all, they should understand that the US Defense Department is governed by the memorandum I have described. Our nation takes its values to the battlefield. But second, more fundamentally, and following everything I've said so far in this lecture, who better than they at Google, who are immersed in this technology, to steer the Pentagon in the right direction? Shouldn't they be like the atomic scientists and help find solutions rather than sitting on the sidelines? And last, I'd ask them whether they're comfortable working for the People's Liberation Army. Because they work in and for China. China is a communist dictatorship, and there is no boundary there. There is no getting around that working in China is working indirectly for the People's Liberation Army or that all of their work is available to the PLA.

I've talked a bit about social media curation and about artificial intelligence, but I do not have time in this lecture for the elephant in the room that is China. I would only say this: we have never been in a sustained economic relationship with a communist-controlled economy. The Soviet Union was such an economy, but our approach to it was not to trade with it at all and to hermetically seal it off from the Western tech world. But we are in an intense trade relationship with China. Because it is a communist dictatorship, China is able to bring to bear on US companies and our trading partners a combination of political, military, and economic tools that a government such as ours cannot match. This puts us at an inherent competitive disadvantage. Though it is not a matter for a secretary of defense, I felt that international economists have failed utterly to provide the US government a playbook for dealing with this situation. The approach preferred over the past decades was rules-based free trade, destined to fail with communist China and in any event abandoned by the US itself when it walked away from the Trans-Pacific Partnership. What is left is a spotty trade "war" and some important but partial limits on Chinese investment in "sensitive" technologies. One additional thing I will say, and as a former secretary of defense, is that it is important to play offense and not just defense. Major national investments in areas like artificial intelligence and public-private partnerships (like the National Manufacturers Institutes founded by the Pentagon during my time) are needed.

Let me now turn to the biological sciences. This is not an area in which I have any particular expertise. But I have attempted to learn about it, and my jobs in the Pentagon gave me plenty of opportunity to be acquainted with some parts of it. I've learned a lot also from Eric Lander, John Deutch, George Church, and others.

• • • • •

It seems likely that a biosciences revolution is looming that will be at least as consequential in coming decades as has been the revolution in the information sciences of the past several decades. The resulting "disruptive" change will be enormous, for both good and bad. The first reason is the sheer number of avenues of innovative change that are being paved by quite recent breakthroughs in biological science. The second factor is a new investment climate that will follow. Let me begin with the remarkable variety of avenues of innovation.

One avenue, of course, is clustered regularly interspaced short palindromic repeats (CRISPR) and the possibility of editing even the human genome. If this passes from laboratory to clinical stages and from animal models to human models—and above all from not only contributing to therapeutics for serious illnesses but also to physical and cognitive enhancements of human potential—then the choices in front of us about where to draw the line are very consequential indeed.

In addition to the obvious moral issues associated with binding one's children and their successors with the decisions a parent makes when offspring cannot conceivably provide any sort of consent, and the moral issues involved in tampering with life itself, there's a serious distributive issue as people of means can purchase a new kind of unequal opportunity that makes any previous form of discrimination pale in comparison.

A different innovative avenue is the growing capacity to create new kinds of designer cells. This has gotten a lot of attention, including by the Defense Department, in the matter of novel pathogens with high lethality and flu-like ability to spread. But it extends to organisms and tissues custom-made for a wide range of purposes, which may be more or less benign.

Yet another category consists of biosensors, and another of biomanufacturers. Biosensors may revolutionize the ability to change environmental signals into processable and storable data in a way we have become well accustomed to with the revolution in electro-optical and other electronic and electromechanical transducers. These sensors could potentially reliably detect even quite subtle and seemingly intangible factors like mood and behavior. Biomanufacturers are custom organisms that can synthesize novel proteins or biological materials in very large scale, thereby making compounds previously only available in trace amounts available in bulk.

There is, next, quite a literature on self-defending cells. These are animal or plant cells provided with new or enhanced ability to defend themselves. These self-defenses could, in turn, be part of the long-sought solutions to cancer, viral infection, or antibiotic resistance.

Additionally, there is the avenue of bio-inspired engineering. Those of you familiar with robotics know that many of them are modeled on either human or animal locomotion with either legs, tails, or cilia. The wheel is an interesting invention in that it has no clear biological precursor, but most of the locomotion chosen by robotics engineering is modeled on nature. So also are biologically inspired exoskeletons and other structural features, and cognitive and behavioral models used in artificial intelligence.

Finally, with all this innovation of all these kinds goes another encompassing avenue of disruptive potential: the union of the information revolution and the biological revolution. It is becoming quite possible, for example, to do a "big data" collection of a cell's DNA, RNA, and protein inventory, not just on a sample basis from a single organism, but cell-by-cell within the organism.

The sheer number and profundity of these bioscience avenues of innovation is the first factor in the coming revolution.

The second factor is who will be able to use all these avenues.

The disruptive avenues of biotech I noted have been until now *laboratory techniques* requiring PhD-level talent and institutional-scale investment and instrumentation. They are becoming *platforms* on top of which scientifically minor, but still socially significant, innovation can build. It is already possible to send off a DNA sample and get an entire sequence returned overnight by email. This took Eric Lander and his colleagues a decade and billions of dollars to do just a few years ago. Someone who knows nothing about the underlying science can sit atop this same platform and think only about novel applications. Many digital unicorns were founded by an entrepreneur using the powerful computational platform on their laptop whose underlying digital technology they neither created, appreciably added to, or even understand.

In Cambridge, Massachusetts, where I work, which is probably the leading biosciences hub in the country, there are a number of bio incubators, including old warehouses, where kids with an innovative idea can set up a little shop and at no expense make use of laboratory equipment that costs millions of dollars to buy. This would have been completely out of the reach of even a pretty well-funded start-up a few years ago.

What this second, nonscientific factor shaping the biosciences revolution means is that the scale and the cost of meaningful innovation will go way down, and the speed of socially (while perhaps not scientifically) consequential innovation will go way up. Sound like digital?

For many purposes, the multibillion dollar, decade-long investment cycle of traditional pharma will be supplemented by something much shorter that can be fueled by fast money—venture capital money. There will shortly be innovators and investors sitting atop the platform exploiting those new bioscience avenues I described who will not necessarily have the culture or the values of research scientists or who have not been brought up with the norms and regulations that come with, for example, National Institutes of Health and Food and Drug Administration funding and approvals, with their rules concerning use of human subjects, protection of personal information, and so on.

What I'm describing here is a climate that looks very much like the early digital era. While a lot of good came out of this combination, including by a lot of people who were essentially amateurs at digital technology itself, we cannot say in hindsight that it came out at all the way we might have hoped.

• • • • •

The third tech-driven revolution of our time is in the future of work and training. I only have time to say this about what is a gargantuan challenge: unless our fellow citizens can see that in all this disruptive change there is a path for them and their children to the American dream or its equivalent, we will not have cohesive societies.

There are a lot of smart kids at MIT and around Boston working on the driverless car. LIDAR (light detection and ranging), which along with passive imagery and radar provides inputs to the steering algorithms, was in fact invented for the military at MIT's Lincoln Labs. I always say to these smart kids, "Save a little bit of your innovative energy for the following challenge: How about the carless driver? What is to become of the tens of thousands of truck, taxi, and car drivers whose jobs are disrupted?"

For these drivers, this unstoppable transition will be like the farm-to-factory transition. We owe it to them to create a Progressive Era of supporting conditions so it all comes out well.

Ash Carter is the director of the Belfer Center for Science and International Affairs at Harvard Kennedy School. He is also an innovation fellow at MIT. For over thirty-five years, Secretary Carter has leveraged his experience in national security, technology, and innovation to defend the United States and make a better world. He has done so under presidents of both political parties as well as in the private sector. As secretary of defense from 2015 to 2017, Secretary Carter pushed the Pentagon to "think outside its five-sided box." He changed the trajectory of the military campaign to deliver ISIS a lasting defeat, designed and executed the strategic pivot to the Asia-Pacific, established a new playbook for the US and NATO to confront Russia's aggression, and launched a national cyber strategy. Secretary Carter also spearheaded revolutionary improvements to the Department of Defense, developing new technological capabilities, leading the "Force of the Future" initiative to transform the way the department recruits, trains, and retains quality people, opening all military positions to women, and building bridges to America's technology community. He was also elected a fellow of the American Academy of Arts and Sciences and is a board member of the Council on Foreign Relations and a member of the Aspen Strategy Group. Currently, Secretary Carter is a member of Delta Airline's board and advises the MITRE Corporation and Lincoln Laboratories on technology matters. Secretary Carter earned a BA from Yale University and a PhD in theoretical physics from the University of Oxford as a Rhodes Scholar.

"Unlike most prior technologies of strategic importance, government is not the primary funder of these technologies, and the primary developers are not defense contractors. Large commercial markets exist, driving the bulk of research, development, and operation."

–JASON MATHENY

Four Emerging Technologies and National Security[1]

Jason Matheny
Former Assistant Director of National Intelligence
and Former Director of IARPA

This chapter discusses four emerging technologies: biotechnology, small satellites, quantum computers, and cognitive enhancement. I'll briefly describe the four technologies and their implications for US national security, and close with some general observations.

Biotechnology

Progress in biotechnology can be attributed to four developments spanning sixty years. First, in the 1950s, the mechanisms of biological inheritance were identified as sequences of nucleic acids read and written by organisms as a ticker tape. Second, in the 1990s, practical methods were developed for reading these sequences at scale. Reading the sequence of the first human genome was completed in 2003 at a cost of over $2 billion. Today it costs less than $10,000 to sequence a human genome—a 200,000-fold reduction in fifteen years. The third development was, over the last fifteen years, a 1,000-fold reduction in the cost of writing sequences. These advances made practical the synthesis of individual genes that could be inserted into existing organisms, as well as the synthesis of whole genomes of microorganisms. The fourth development, in the early 2010s, was the development of cost-effective tools for editing genomes using the CRISPR/Cas9 system. With these tools, biotechnology can be used to create microbiological factories for making small things in large numbers. One can instruct the factories to produce beneficial things, like medicines, fuels, food, textiles, and catalysts. Or one can instruct the factories to produce destructive things, like viruses and toxins.

Even without the application of human creativity, biology is among the gravest threats to national security. In 1918, a single influenza virus caused over fifty million deaths during a twelve-month period, ranking as the most severe loss of life in human history, including wars and other disasters. The Plagues of Antonine and Justinian, the Black Death, and the Cocoliztli epidemics were also far more lethal than contemporary wars. Biology is powerful because of two properties. The first is self-replication, by which a small arsenal is efficiently turned into a large one. (Apart from biological weapons, the only other class of weapon that can self-replicate is cyberweapons.) The second is natural selection, by which random traits advantageous to an organism will become more prevalent over time. Biotechnologies allow these two properties to be modified toward human goals. For example, one biotechnology called a "gene drive" allows human engineers to ensure that traits become ubiquitous within a species even if the trait is disadvantageous to the organism. (A benevolent application of gene drives is engineering mosquitoes incapable of transmitting malaria. A malevolent application is engineering honeybees incapable of pollination to devastate key parts of agriculture.) Another category of biotechnology, gain of function, uses artificial selection to amplify the characteristics of an organism, such as lethality or transmissibility, beyond rates that can occur naturally.

Gain of function research is openly published despite security concerns. On the one hand, the results of such research can be used to develop advanced biological weapons. On the other hand, the results can be used to assess risks and develop better defenses. The benefits of publication are limited if biology has substantial offense dominance—if, at the margins, offense is more effective than defense. As one example, in 2016 a close relative of the smallpox virus was synthesized for $100,000 using commercially available equipment and materials. Meanwhile, a new defensive vaccine costs over $1 billion to develop. Offense has a 10,000-fold cost advantage over defense, and the advantage has historically increased over time. Even with effective distribution of medical countermeasures, a poxvirus released in the United States would likely kill millions of Americans. It is unclear what defensive technology could improve the odds. For the moment, we live in a vulnerable world in which a single sophisticated misanthrope is capable of killing millions of people.

Small Satellites

For national security, space plays two key roles: as a transit area for weapons, such as intercontinental ballistic missiles, and as a perch from which to observe and communicate, as with satellites. In 1945, when Arthur C. Clarke first proposed communications satellites, he recognized space as the ultimate high ground—from a position in orbit, satellites could maintain lines of sight to many points on Earth, receiving signals from one location and transmitting them to any other. It was twelve years before Clarke's concept was realized in the first satellite. The subsequent sixty years have seen his concept elaborated, with thousands of satellites now used to send and receive communications, enable timing and navigation, and photograph Earth's surface.

Delivering objects to space is costly. For the last sixty years the cost of launching objects into a low Earth orbit (LEO) has remained steady, on the order of $10,000 per kg. Some new technologies, such as reusable boosters, might reduce that cost by 10 to 50 percent, but moving objects to space will remain costly due to the inconvenience of gravity. A more encouraging trend is the miniaturization of payloads. Over the last sixty years, microelectronic weight per unit performance has decreased by a factor of around one million. Miniaturization has substantially increased the processing capacity of satellites of a given size, with cubesats—small satellites with standard dimensions in multiples of 10 x 10 x 10 cm—as one example of what can be accomplished in small payloads. Still, there are limits to miniaturization, particularly for the parts of satellites that involve apertures and antennas whose dimensions have lower bounds dictated by physical laws. While the sizes of apertures and antennas might remain constant, their costs might be substantially reduced by on-orbit servicing, assembly, and manufacturing (OSAM) methods that transform less fragile, small payloads of parts and raw materials into large, complex structures. Future automation of OSAM, combined with rapid launches of raw materials into orbit, could create a formidable space infrastructure.

The democratization of space by commercial satellites has meant that wealthy states no longer monopolize detailed imagery of the earth's surface. Great powers have lost their advantage in one important dimension of national intelligence, and there are now fewer places on Earth left to hide. The democratization of space has also led to a growing number of active spacecraft and inactive objects (such as rocket bodies) in orbit—currently on the order of 16,000 objects above 10 cm in size in LEO alone. This number is likely to double within the next decade, increasing the probability of collisions and complicating space traffic management. Fortunately, space is big—the volume to geostationary orbit is about 100,000 times larger than the volume of Earth's oceans. But tracking objects in space, particularly foreign military and intelligence spacecraft, will remain a national security priority.

Space infrastructure is highly vulnerable. Satellites are fragile—they're optimized for weight and volume, so they're difficult to armor and vulnerable to damage by collision, missiles, lasers, radio jamming, and space weather. Satellites are also costly to replace because of their innate cost and their cost of delivery. Protecting these systems, finding ways to quickly replace them, and developing redundant services that don't depend on space are all difficult national security problems. The US is more dependent on space than other countries, but the gap is decreasing as more countries launch their own communications and GPS systems, placing their own infrastructure at risk.

Quantum Computers

Quantum computers exploit two quantum phenomena: superposition, the ability for a quantum bit (qubit) to simultaneously maintain more than one state, and entanglement, the ability for two or more qubits to share their states at distance. In the early 1980s, physicists hypothesized that a computer exhibiting both phenomena could be well suited to some types of computation, such as simulations of quantum behavior. It was unclear whether such a machine would have security applications until 1994, when Peter Shor discovered that a large quantum computer should be able to efficiently factor large numbers. Most current encryption methods, including many used to protect military communications and online financial transactions, depend on the difficulty of factoring. If factoring becomes easy, then many existing encryption techniques are no longer secure.

Progress in quantum computing has been difficult. Although substantial gains have been made in theory, materials, and measurement, since 2001 the size of quantum computers has increased from 7 to only 72 physical qubits. A useful quantum computer would require several thousand qubits, and the difficulties of error correction increase with every addition. To date, no single logical qubit that could serve as the building block of such a computer has been created. In contrast, progress has been faster in creating new encryption methods resistant to quantum attacks. Several such methods now exist, though they are computationally less efficient than traditional encryption.

Cryptanalysis is believed to be the quantum computing application of greatest importance to national security, but there are other potential applications. These include simulating molecules to design chemicals or catalysts and solving hard optimization problems, such as those used in logistics, scheduling, and machine learning. There has not yet been a useful computation by a quantum computer that could not otherwise be performed at lower cost by a classical computer. At some point, however, researchers are likely to achieve this milestone, sometimes called "quantum supremacy."

If current trends continue, large-scale quantum computers will probably be introduced well after quantum-resistant encryption has been widely deployed. If so, the computers' main intelligence value would be in decrypting historical data. The countries that most benefit would be those that store intercepted foreign communications for long periods. By policy, the US limits retention of most SIGINT to five years, casting doubt on the relative value of large-scale quantum computing for US cryptanalysis. Still, it is important for the US not to be surprised by developments in quantum computing, in part because the rate of progress will inform our own schedule for deploying quantum-resistant encryption methods. Since these methods come at a computational cost, we will need to decide how long to invest in their refinement and when to simply accept the costs as necessary.

Cognitive Enhancement

In the future, humans might invent technologies to make themselves better inventors. The scientific method, scientific instruments, peer review, computers, and patent systems were enhancements to the infrastructure used for invention. If human intelligence is itself enhanced, then the rate of invention might be more substantially increased.

Potential pathways to cognitive enhancement can be divided into three categories: devices, drugs, and genes. Devices that link brain tissue to external devices, such as prosthetic limbs, have become increasingly sophisticated due to military research following the increase in non-lethal combat injuries. Some start-ups are now exploring whether such neural interfaces can be used to offload processing and memory tasks directly from the brain to peripheral devices.

The increased availability of drugs for neurocognitive conditions such as attention deficit disorder, narcolepsy, and dementia has supported a subculture of healthy adults who consume these drugs to enhance cognition. It's unclear whether any of them succeed—there is scant evidence that these drugs improve cognitive performance in healthy adults any more than caffeine or nicotine. Still, it is possible that future discoveries will lead to cognitive gains.

The last category of cognitive enhancement could be the most effective and the most controversial. The decreasing cost of DNA sequencing has enabled large genomic studies of human intelligence, in which some variance in cognitive ability can be attributed to genes. China's Cognitive Genomics Project is one example of such a study, and surveys suggest that most Chinese parents would choose to use genotyping services to select more intelligent children, if such services were available. Surveys suggest that American and European parents would be far more reluctant to use such services. If these surveys are representative, the rate of technology adoption could be substantially affected by differences in cultural attitudes. Since the rate of invention is one leading determinant of national power, any "cognitive arms race" that affects rates of invention could have important geopolitical and economic consequences.

Observations

Government faces a number of security challenges in addressing these four technologies, along with a fifth technology, artificial intelligence (AI) addressed in this volume with a chapter by Jack Clark. These security challenges result from the following three features shared by the technologies.

Commercial markets. Unlike most prior technologies of strategic importance, government is not the primary funder of these technologies, and the primary developers are not defense contractors. Large commercial markets exist, driving the bulk of research, development, and operation. In some cases, the components of key national security systems are dependent on commercial markets—for example, GPUs (Graphics Processing Unit) developed for video games are the basis of supercomputers used for nuclear weapons simulations. In other cases, cutting-edge technology is unavailable to some parts of government—for example, Google has a policy against designing AI for weapon systems.

Dual-use characteristics. Because of the commercial applications of these technologies, military programs are difficult to distinguish from commercial ones. There are few signatures that distinguish a legitimate biotechnology effort from a biological weapons program. Similarly, it is difficult to detect whether a computing system is being used for benign scientific research or to decrypt military communications.

Rapid change. Unlike most weapons, some of these technologies exhibit exponential improvements in performance over time. Such rapid change makes planning and adaptation difficult, especially on the typical time scales of policy making and government procurement.

Over the next decade, competitions in these technologies (and others) will challenge US national security decision makers. They will be pushed to formulate national strategies for accelerating innovation, including increases in R&D funding, changes in immigration policy to prevent brain drain, expansions of IP protection to limit technology theft, and temptations to compel commercial technology firms to address national security priorities. Defense agencies will be challenged to partner with technology organizations outside their traditional contractors and to operate on procurement time lines closer to commercial expectations. Intelligence agencies will be pressed to monitor foreign commercial technology firms and foreign academic labs, where key strategic technologies are being developed. Lastly, diplomats will be moved to negotiate in unfamiliar territory, such as obscure international standard-setting bodies that influence whose technology becomes the global norm.

Jason Matheny was assistant director of national intelligence, and director of IARPA, the U.S. intelligence community's research organization. Before IARPA, he worked at Oxford University, the World Bank, the Applied Physics Laboratory, the Center for Biosecurity, and Princeton University and was the co-founder of two biotechnology companies. He is a member of the National Academies' Intelligence Community Studies Board, is a recipient of the Intelligence Community's Award for Individual Achievement in Science and Technology and the Presidential Early Career Award for Scientists and Engineers, and was named one of *Foreign Policy*'s "Top 50 Global Thinkers." He has served on various White House committees related to high-performance computing, biosecurity, artificial intelligence, and quantum information science. He co-led the National AI R&D Strategic Plan released by the White House in 2016 and was a member of the White House Select Committee on AI, created in 2018. He holds a Ph.D. in applied economics from Johns Hopkins University, an M.P.H. from Johns Hopkins University, an M.B.A. from Duke University, and a B.A. from the University of Chicago.

[1] This chapter is based on a talk given to the Aspen Strategy Group, largely unedited and thus containing all the mistakes of the original. My thanks to Richard Danzig, Nick Burns, and the Aspen Strategy Group for the kind invitation to participate.

"If governments feel they are in a non-collaborative race with other private or state actors, they will be incentivized to cut corners on techniques to increase the predictability, robustness, and reliability of such systems. This could lead to powerful AI systems being deployed with unpredictable risk profiles—inviting us to wonder what might an AI-Chernobyl look like?"

–JACK CLARK

Artificial Intelligence and Countries: What Might Happen, What We'd Like to Happen, What We Should Avoid Letting Happen

Jack Clark
Policy Director
OpenAI

There's an old story that scientists who have worked adjacent to national security like to tell each other, usually after talking about the difficulties they are having with their current artificial intelligence work.

The story goes like this: once upon a time an intelligence agency had some desire to automatically identify tanks from satellite images and to distinguish between the different tank platforms fielded by different nations. The intelligence agency embarked on an expensive and long-running data collection exercise in which it had teams of human analysts laboriously label the tanks in a variety of satellite imagery. After the tanks had been labeled, the agency trained a simple, supervised learning system to distinguish between the different types of tanks within the testing dataset, thereby (theoretically) automating a large job within the intelligence apparatus and likely improving the awareness and preparedness of the national security establishment.

Eventually, the bureaucracy decided to deploy this solution into the field. With expectations high, the researchers gathered around to see the first prediction made by the system when given a contemporary satellite image: its prediction was false. "No matter. It had an accuracy of 99.99 when we tested it in simulation, so this means the next few thousand classifications will be accurate," said one of the developers of the system (and a lousy gambler). But the results of the next prediction were also wrong. And the next. And the next.

Several hours later, and with a statistical success rate that was actually worse than coin-flipping, the test was concluded, and the developers of the system were sent back to whatever R&D broom closet they were based in. It took many weeks for them to isolate the cause of the problems with the AI system, and, in the end, it had nothing to do with tanks and everything to do with the weather. After further studying the images used to train the system, the scientists discovered that all the tanks from one country had been photographed under some specific environmental conditions—partially cloudy, with frequent snowfalls on the ground—while those from other countries had been photographed under a very different set of circumstances. Therefore, the AI they had built was functionally useless for the purpose of tank detection but potentially very useful for automatically classifying weather conditions via images created by spy satellites.

Few of the problems identified in this story have been solved. Today's AI systems are still vulnerable to problems like over-fitting to a particular distribution of data, like the particular environmental contexts of the tanks in the above story; experimental methodologies for AI are still based around empiricism so frequently that we have to test systems against real-world data to uncover flaws; and there continues to be a wild mismatch between what people expect AI systems to be able to accomplish and what they can actually do.

But some fundamental things have changed: computers have become significantly faster, datasets have become more plentiful, and the range of actors capable of developing and deploying this technology has broadened dramatically. AI has matured to the point that now many entities around the world are developing and deploying it—and the majority of these entities don't work for any government.

Today's AI: A Machine Readable World

So, how did we get here? In the past few years artificial intelligence techniques based around the use of neural networks have matured to the point that today's AI can, in certain highly constrained domains, outperform humans at tasks relating to the classification of data, whether that means automatically labelling images, transcribing text, or dealing with more abstract and hard-to-discern streams of data, like monitoring the flows of traffic over a digital network for signs of anomalous behavior.

Systems based on such techniques are now being deployed to do tasks as varied as identifying the potential signs of tumors in medical imagery, predicting potential locations for earthquakes, reading lips from video footage, tracking individual pedestrians as they move across cities, facial recognition, providing automated disaster response analysis via analysis of drone-based footage, and more. As a rule of thumb, AI systems are now good enough that if you have a large quantity of well-labeled data and access to a large computer, it is relatively easy to develop a system that can classify that data with greater than 90 percent accuracy and, in some cases that have received more commercial development, to accuracies of 99 percent or higher.

Today, the confluence of digitization (increasing amounts of digital communication, a fall in the cost of acquiring satellite imagery, a proliferation of cheap sensors as a consequence of the smartphone boom, and so on) and AI has made the world essentially machine readable, which in turn is making it substantially cheaper and easier for interested parties to use AI to more effectively and efficiently achieve their goals. In the same way that companies such as Google and Amazon use AI to let their businesses run more rapidly and efficiently, intelligence agencies and militaries around the world are now exploring how to use AI to strengthen their signals analysis capabilities and let them passively monitor more of the world's signals in real time.

Implications of Today's AI: Deployment

These AI capabilities, as basic as they are, have already influenced the plans of nations: China has committed to be the global leader in AI by 2030; countries like South Korea, France, Canada, the UK, and others have all announced national AI strategies (backed up by increased funding); and Canadian Prime Minister Justin Trudeau has said AI gives Canada the opportunity to "be relevant in a positive way on the world stage."[1]

Countries are choosing to use AI as a way to signal their desires for technical advancement because of its clear utility and significance. In many ways, such comments are a replay of what nations said about cloud computing and big data during the past decade. The difference is that AI is a technology that makes use of many previous innovations, including big data and cloud computing, and is continuing to improve at a fundamental technological level far in excess of prior revolutions. This rate of improvement means that nations and other actors will soon gain AI-based capabilities that go beyond being able to classify the world. Soon, AI systems will take increasingly sophisticated sets of actions with less and less oversight (for instance, Google's data centers are now managed by AI algorithms that observe a stream of data from sensors within each facility, then automatically adjust machinery to optimize the efficiency with which these data centers use electricity). This shift from pure observation to observations chained into actions that themselves condition the observations will dramatically broaden the utility of the technology and the chance for accidents when using it. We will see deep ramifications on how nations relate to one another, how nations relate to their own private companies, and how bad actors will seek to utilize AI to cause harm to others.

Tomorrow's AI: A World in Which (Unreliable) Machines Act

The next few years of AI will be distinguished by the creation of machines that not only classify the world, but that take actions within it as well. Systems like self-driving cars are perhaps the most visible example of this trend, but other examples abound. Techniques in development today are already yielding software to automatically manage the flow of power across utility grids; systems to teach computers to learn to construct software fundamental to cloud computing such as search algorithms for databases; better auto-navigation systems for delivery robots and drones; manipulation systems that will give us robots that can map, navigate, and open doors within buildings along with specialized robotic hand-based systems that can be taught to manipulate objects with similar capabilities to humans; and new systems based around giving AI greater imaginative and memory capabilities that will let us develop machines that can analyze complex situations and conceive of plans and strategies to implement within them, such as work by OpenAI on the strategy game Dota 2 and by Alphabet-subsidiary DeepMind on the e-sports game StarCraft (and prior to that, Go). These systems are going to become increasingly adept at automatically analyzing and acting within their environments, whether that is autonomously exploring areas of science or mathematics or powering physical things that begin to influence the world.

These action-oriented AI systems are vulnerable to many of the same flaws as the previously discussed tank-detector with a crucial distinction: as we build systems that take increasingly long sequences of actions that are themselves learned without human oversight, the chance of them failing in subtle or hard-to-anticipate ways increases. We've already seen examples of this at OpenAI in our work on creating increasingly autonomous robots. For one recent project, we trained a robot to reach to move a green hockey puck to a particular point on a simulated table.[2] Our robot succeeded at this task, with all of the graphs showing that, after many hours of training, our simulated robot was successfully moving the puck to the specified point on the table. But when we looked at the video footage from inside the simulator to see what the robot was actually doing, we made a disconcerting discovery: the robot had learned to use its arm to move the table the puck was sitting on. By moving the table, it was able to move the puck to the point we had specified.

Why had this happened? It was because we hadn't set up our simulation correctly: the location the puck needed to be moved to wasn't bound to the surface of the table but was instead a point hovering in space. The puck, meanwhile, was simulated to have weight and sit on the table. Therefore, the robot had learned that sometimes it was more efficient for it to move the table and therefore also move the puck toward the specified space. Our solution to "fix" this problem was to make the simulated table weigh 200 tons, making it impossible for the robot to move it; a hacky fix, as we term it in the technology industry, but reasonable given that we're interested in using these simulations to develop more efficient algorithms. We don't need our fix to necessarily correspond to how you might fix the system in the real world.

But we worry that as the technology further improves, governments and industry will be motivated to deploy such systems in increasingly sensitive areas (while putting in place controls to spot problems like our table-moving robot or a miscalibrated tank-detector) before the research community develops sufficiently powerful techniques that guarantee system robustness and predictability so we can confidently rule out the chance of such failures. Why might governments do something so seemingly risky? Because the advantages of deploying AI systems to gain a strategic upper hand over a competitor are significant, and today's incentives for most parties are to try to develop a profound technological lead over other parties, leading to a kind of technological race.

Why Compute, Not Data, Will Drive the Next Phase of AI Development

There's one additional factor that will condition how states approach AI in the future: hardware-based computational abilities (compute). Today's AI systems rely on a combination of data, algorithms, and access to large amounts of computers. It's worth noting that for many systems, we can substitute compute for data. So in our tank example, we could generate a broad range of synthetic images to use to train our machines. Or in the robot example, we can use computers to simulate the robot and its environment, thereby paying for computers to create synthetic data for us. Because the cost of computation is consistently falling over time due to competitive market pressures, the arrival of new computer chips, and algorithmic improvements, it is becoming cheaper and cheaper to use computers to generate or augment datasets.

Behind all these advancements lies one unifying trend: the industrialization of AI. This trend stems from the confluence of increased computational capabilities, better engineering practices, a rise in the amount of digital data, growth in the research sector, and the arrival of multiple end-users of the technology ranging from governments to large companies to individuals.

One key factor within this industrialization trend is the growth in compute. Today, we have a line of sight to computers that are as much as ten or a hundred times faster than those we use today, and these bits of hardware will likely arrive this year or next. Alongside that, the AI sector is getting better at developing systems that utilize many machines in parallel—for instance, OpenAI's work on training an AI to beat humans at the competitive e-sport Dota required us to use over 100,000 computers in parallel. This means that AI is actually progressing faster than many people's intuitions would suggest. An analysis recently performed by OpenAI shows that the amount of compute utilized by significant AI applications has grown by over 300,000 times in around six years because of the combination of faster computers and our ability to use more of them in parallel. If this trend holds—and with the arrival of new computer hardware, that looks to be guaranteed for the next few years—then increasingly powerful AI systems are going to become easy to develop. The general utility of these systems will be far greater than those we have today, which will broaden the capabilities available to actors deploying them and increase the pressure to deploy them.

Why AI Progression Requires a 'Collaboration Moonshot'

Because AI is embodied in software rather than hardware, the technology will proliferate more rapidly than prior technologies with significant relevance to national security. It will also be innately difficult to track and apply controls to—again, because of its embodiment in software. This means that as our algorithms become more capable, and as the costs of developing and deploying the technology become ever cheaper, an ever larger group of actors are going to be motivated (and have the means) to develop this technology. Much of this development is going to be positive, leading to more secure nations, greater economic productivity, and astonishing advances in health care and utility infrastructure. But because AI is a dual-use technology—pretty much any new AI algorithm can be used for both military and nonmilitary purposes—in parallel, increasingly sophisticated AI-based tools will be developed for attacking others, whether through cyberwarfare, the weaponization of robotic platforms or drones, and more.

If governments feel they are in a non-collaborative race with other private or state actors, they will be incentivized to cut corners on techniques to increase the predictability, robustness, and reliability of such systems. This could lead to powerful AI systems being deployed with unpredictable risk profiles—inviting us to wonder what might an AI-Chernobyl look like? What would happen to the world if the fallout of such an event were not contained and instead spread over global digital networks, leading to the software equivalent of radioactive material being distributed across the world?

We've already had indications that non-AI-based government-developed technologies can have tremendous fallout effects. The Stuxnet virus was designed to target Iranian enrichment facilities, but ended up "breaking out" of the Iranian facility and corrupting computers in other countries. The more recent NotPetya malware was thought to be developed by the Russian government to target computers in Ukraine but subsequently spread worldwide, causing billions of dollars in damage and taking down hospitals and utility grids along the way. We do not want to know what the AI version of such situations might be.

What if, instead, a government—such as the United States—invested in techniques relating to AI safety, predictability, and robustness and distributed many of these techniques around the world? This would have two significant effects: it would increase the links between that government and the academic and corporate technical communities, letting the three different constituencies work together on a matter of common concern, accelerating development here. The second effect would be more subtle but potentially more transformative for geopolitics: by increasing the robustness and predictability of such systems, the United States would have a logical point of collaboration with other governments, even ones with ideologies or positions in opposition to the United States.

By successfully pioneering the invention of these techniques, and the diffusion of them across the world, governments could increase the rate of their own AI development and create the infrastructure for ongoing (selective) collaboration with others around the world.

By diffusing such safety techniques into the world, all governments would benefit. It would be easier to establish the "rules of the road" for AI-mediated conflict with other states, and it would lay the groundwork for subsequent work on forming "safety-first" agreements with other states that would potentially incorporate verification regimes, as has been done with nuclear power and nuclear weaponry.

The good news is that AI is not yet at the point where the transformatively *destructive* applications are feasible, giving the international community time to work on the problems of AI safety (and the associated policy challenges). This means we have time to think about national AI strategies that incorporate AI safety as a route to potential future collaborations that will create global infrastructure to minimize accidents and misuse. We should begin to move deliberately—imagine how we might feel in six years after another 300,000 times increase. Will we be able to look at each other and say we did enough?

Jack Clark is the policy director for OpenAI, an AI research and development organization. At OpenAI, he leads the organization's efforts around AI policy and regularly meets with politicians and VIPs across the world. He also sits on the steering committee of the AI Index, an initiative from the Stanford One Hundred Year Study on AI to track and analyze AI progress. Additionally, Mr. Clark is a member of the Center for a New American Security (CNAS) Task Force on AI and National Security. This year he participated in the 2018 Assembly program on ethics and governance in AI at the MIT Media Lab and Berkman Klein Center at Harvard. Mr. Clark spends a lot of time thinking about the potentially destabilizing effects of rapid technological progression and is one of the co-authors of *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*. He writes a weekly newsletter about cutting-edge AI research and applications called *Import AI* (www.importai.net), which is read by more than 15,000 experts worldwide.

[1] Darrell Etherington, "Justin Trudeau Explains Why Canada Really 'Gets' AI and Smart Cities," TechCrunch, November 2, 2017, techcrunch.com/2017/11/02/justin-trudeau-explains-why-canada-really-gets-ai-and-smart-cities/.

[2] Surprising as it may be, this is a difficult task for today's AI systems because we're trying to teach robots to learn how to move their arm to move *any* puck to *any* position on the table. Traditional approaches would simply program the robot to reach a particular puck at a particular point on a particular table.

“Society needs to adopt a view of continuous learning and reskilling throughout a person’s life, with emphasis on more demand-driven and vocational training and a greater focus on team-based skills and cross-functional and disciplinary integration.”

–DIANA FARRELL

Technology and the Future of Work

Diana Farrell
President and CEO
JPMorgan Chase Institute

The technological changes that drove the first, second, and third industrial revolutions in developed countries in the eighteenth and nineteenth centuries saw huge portions of the population moving off the farm and into factories. The resulting mass migration and social dislocation produced years of disruption, requiring time and major policy interventions to navigate.

Today, the world economy and its participating workforce is engaged in a shift of equal or larger proportion, brought about by the technological advances of our age. **Digitization** is translating huge amounts of economic and social activity into machine readable formats that can be analyzed and re-imagined. Companies are using **artificial intelligence**, from machine learning to predictive analytics, in a growing range of contexts. **Robotics** and other semi-autonomous machines are taking on tasks that were long believed to be executable only by humans. Finally, advances in **materials and biological** sciences, such as nanotechnologies and bio-engineering, are rendering robots more human-like and able to change how humans work.

Many of these technologies have been used in some form for decades, but they are now entering a new phase. As in the legend of the rice and the chess board, in which each square on the board produces a doubling of the wager from one grain of rice to two, then from two to four, the improvements and use cases for these technologies are now scaling "on the second half" of the technological chess board.

Much of this change is full of promise. Yet there are real risks that the gains from technological advancement will only benefit some industries and some societies and will not translate broadly to improve prosperity and well-being. The future of work might see huge proportions of the workforce shifting into lower-productivity, lower-paid occupations and jobs.

We have historical precedent for how rapid technological advances produce radical changes in society. We don't know whether the "fourth industrial revolution" in which we are currently engaged is going to look the same as the technological advances that preceded it, but the possibility requires us to learn the lessons of the past. It is not too soon to mobilize—now—to imagine the future of work and design policies and approaches that can optimize its benefits. That process begins by understanding the dynamics already underway.

Simultaneous Impacts from Technology

The technological advances of today are affecting the future of work in three key ways: by scaling and speeding up human capabilities, by substituting labor with machines, and by enabling new ways to access and supply labor.

Scaling and speeding up human capability. Technology has been enhancing human capability since the invention of the Gutenberg press. Modern applications of technology to improve productivity are legion, though two examples offer representative scope of today's possibilities to change the work landscape.

One is *MITx*, a pilot project launched by MIT to create a digital replica of the classroom. MIT offered the prototype course "Circuits and Electronics" in March 2012 to over 154,000 students from more than 160 countries. As is typical with massively open online courses (MOOCs), less than 5 percent of registered students passed the course, but that percentage clouds the course's impact. In absolute terms, 7,157 students passed in one semester—as many as MIT could accommodate in person in forty years.[1]

A second example comes from the field of robotics, in which recent developments in top artificial intelligence labs are resulting in the development of modern robotic hands that can spin, grip, pick, make beds, and place objects of any size.[2] Efforts to enhance and support aging or injured human bodies, and to improve memory and access to information, show equal promise.

Substituting labor with machines. Manufacturing provides less than 9 percent of US employment, down from over 30 percent in the 1950s.[3] Advancements in automation account for much of this decrease, allowing today's factories to produce far more product with far fewer people. Jobs across a range of service industries will soon see a similar decline.

In fact, in their assessment of the susceptibility of US labor to computerization, Frey and Osborne examined 702 occupations and concluded that roughly two-thirds of US jobs are at high risk or medium risk of being computerized. More specifically, the analysis showed that 47 percent of US employment is in high-risk jobs that have a 70 percent to 100 percent probability of becoming computerized in the next ten to twenty years; 19 percent of employment is medium risk, with a 30 percent to 69 percent probability of becoming computerized; and 33 percent is low risk, meaning not likely (0 percent to 29 percent probability) of being computerized.[4] The most vulnerable to least vulnerable jobs move along a continuum from simple, repetitive, and routine to optimizing, complex, and creative.

High-risk jobs include equipment operators (notably transportation drivers), jobs that require basic cognitive skills such as inputting data, and a range of administrative jobs. On the low-risk end of the continuum are jobs that require higher cognitive ability, such as creativity, cross-disciplinary faculties, technical abilities, and social or emotional connections.

These trends are not necessarily negative from a labor, economic, or social perspective. Automation can be welcome when it solves a labor shortage or other source of economic friction. For instance, Japan is experimenting with robots in assisted living facilities to address the mismatch between a rising elderly population and a shortage of specialized care workers. Japan's Ministry of Economy, Trade, and Industry has provided 4.7 billion yen ($45 million) in development subsidies since 2015, and the labor ministry spent 5.2 billion yen ($50 million). These investments have resulted in the introduction of more than 5,000 robots in assisted care facilities from March 2016 to March 2017.[5]

In another context, language translation is reducing friction in online trading, with significant economic impact: a recent study led by AI expert Erik Brynjolfsson at MIT found that the introduction of a machine language translation application on an online platform led to a 17.5 percent increase in international trade between participants.[6] The authors conclude that language barriers between humans limit trade, and machine-based translation reduces the need for those humans to engage in translation-motivated searches before executing a transaction.

Notwithstanding these positive applications of automation technology, the speed and scale with which human tasks are becoming automated is uncertain and is very likely to result in significant displacements for workers in the most vulnerable occupations.

New ways to access and supply labor. As workers are automated out of jobs in the traditional economy, many are embracing independent work accessed through online platforms as an alternative or supplement to traditional jobs. Growth in the Online Platform Economy has enabled more employers and employees, suppliers and consumers, and even suppliers and businesses to find each other and exchange labor and goods.

The number of significant online exchange platforms has grown from over forty to more than 120 in the last two years. The JPMorgan Chase Institute segments the Online Platform Economy into four sectors: the transportation sector, which involves drivers transporting human riders and/or goods; the non-transportation labor sector, which deals in services such as dog walking, home repair, etc.; the selling sector, in which independent sellers find buyers through online marketplaces; and the leasing sector, in which people rent their homes, parking spaces, and other assets. Transportation dominates in terms of both the number of participants and total transaction volume, accounting for over 60 percent of total platform economy participants and generating as much revenue as the other three sectors combined.

The JPMorgan Chase Institute's most recent report shows 1.6 percent of the 2.3 million families in its sample earned platform income in the first quarter of 2018, up from 0.3 percent in the first quarter of 2013; 4.5 percent earned platform income at some point during the prior year. Generalizing these numbers to the 126 million US households suggests that 5.5 million households earned income from the Online Platform Economy at some point during the year.[7] As a fraction of all jobs, there were almost as many people working in the platform economy in the first quarter of 2018 as there are workers in the information sector (1.8 percent of all jobs in 2016); the 4.5 percent that earned platform income for all of 2017 is comparable to the percentage of workers in public administration (4.6 percent of jobs in 2017).[8]

Despite the scale of employment in the Online Platform Economy, participant engagement and income from online platforms is more sporadic than in most traditional jobs, and the work typically lacks both benefits and opportunities for advancement. Monthly platform earnings represent an average of 20 percent or less of total take-home income for individuals who participated in the Online Platform Economy in the twelve months captured in the JPMorgan Chase Institute's data set.[9] Of the transportation platform participants who drove at any point in the twelve-month period, 58 percent had earnings in just three or fewer months. In the other sectors, engagement was more sporadic. Moreover, as platforms have grown in number and size, earnings have declined or remained stagnant. Specifically, average monthly earnings in the transportation sector fell by 53 percent between 2014 and 2018, and average earnings in the other three sectors were flat.

In sum, online platform work does not yet seem to be replacing more traditional sources of income, but it is changing the composition and wages of work for growing numbers of people.

Unintended Consequences and Emerging Risks

Many will view some of the issues described above, particularly the push toward automation, as trends developed society has seen and recovered from as new jobs in heretofore unimagined new areas emerge. That has been true to a considerable degree, but there are five troubling and related signs that deserve close attention.

Shrinking share of total income to labor. Over the past few decades, economies around the world have seen a fall in the share of total income allocated to labor. In 1947, 66 percent of total income accrued to labor in the US. In 2016, that number had fallen to 58 percent. The pattern for OECD countries is similar, showing an income-to-labor decrease from the 1990s to the late 2000s of 66 percent to 61 percent. The IMF estimates that in advanced economies about 50 percent of the decline is attributable to technological advancement, while another 25 percent is attributable to global integration, which is itself facilitated by technology.[10] More technological advances will likely reinforce this long-standing trend.

Income inequality. OECD research suggests that countries where labor income share has decreased likewise show an increase in income inequality.[11] Indeed, the long-standing rise in both income and wealth inequality in the US and other countries is well documented. The average income of the top quintile earners in the US, after transfers and taxes, was over twenty times greater than for the bottom quintile earners.[12] Differential income growth rates by wage level has caused that gap to widen. Between 1979 and 2016, real hourly wages increased 51.7 percent for workers at the ninety-fifth percentile of wages and 4.4 percent for workers at the tenth percentile.[13]

Rising income inequality is increasingly linked to direct and indirect economic growth limitation, as noted in a recent IMF study, and is a challenge to political stability and trust in the legitimacy of the system.[14]

Living wages. When the majority of income-generating resources are controlled by a small percentage of society, and labor holds less power in the economy, members of the working class are less able to negotiate wages to allow them to adequately support a family through income from full-time work. Indeed, the cost of living in many places in the US has long been rising while incomes remain stagnant. The extent of the problem was laid bare in 2003 when the MIT Living Wage Lab created a tool to estimate the cost of living and the minimum living wage necessary to cover basic living expenses across many US cities, including housing, groceries, basic health care, etc. Calculations from the tool show that the median living wage in the US in 2017 was $16.07 an hour for a family of four with two working parents, more than twice the federal minimum wage of $7.25 an hour. A typical family of four needs to work nearly four full-time minimum-wage jobs to earn a living wage.[15]

Today, 42 percent of Americans make less than $15 an hour and 30 percent make less than $10 an hour.[16] Two out of every five people struggle to access consistent housing, food, and health care. A full-time minimum-wage worker can afford a one-bedroom rental home at fair-market value in only 22 out of 3,000 US counties.[17] Even before the technological disruption we anticipate on the horizon has fully revealed its impact, a significant share of jobs in today's economy barely meet the "sufficient" threshold.

The downward pressure on wages is further exacerbating cost-of-living challenges felt by so many—even those in developed countries experiencing overall economic growth. According to an McKinsey Global Institute report, the share of households that experienced flat or falling incomes, after taxes and transfers, increased from less than 2 percent between 1993 and 2005 to between 20 percent and 25 percent from 2005 to 2014.[18] The most recent US Bureau of Labor Statistics data show a 0.2 percent decrease in real average hourly earnings from July 2017 to July 2018.[19]

Widespread income and expense volatility. Low wages are not the only source of pressure on low- and moderate-income households. Many also struggle with irregular incomes. Extensive research from the JPMorgan Chase Institute provides robust evidence that Americans at all income levels experience high degrees of income and spending volatility. In fact, family income and spending typically fluctuate by about 30 percent on a month-to-month basis. These fluctuations are only slightly positively correlated, meaning families could simultaneously experience an expense spike and an income dip.[20]

To withstand typical volatility, families need emergency savings and a cash buffer, yet few Americans have either, which makes them vulnerable. A 2016 Federal Reserve survey found 46 percent of families could not readily pay for a $400 emergency expense, or would have to borrow or sell something to do so.[21] When such an expense is associated with medical services or auto repairs, common sources of high-value expenses, the lack of a cash buffer can affect a family's financial and physical health and the ability to maintain steady employment.

Insufficient transition to higher-value-added, higher-income jobs. If the economic environment makes it difficult for people to earn higher wages for the same work, or to smooth earnings and expenses, workers can still seek better, higher-paying jobs as a way to increase income and spending power. Yet that too is becoming out of reach for many, as new job creation falls in the sectors where new technologies are advancing most rapidly. In the past, creative destruction has led to new and higher-value-added job creation and may do so in the future, at least to some extent. But the trend suggests fewer new jobs created in industries affected by technological disruption: in the 1980s, 8.2 percent of US workers shifted into jobs that emerged for the first time during the decade. In the 1990s, that number was 4.4 percent; in the 2000s, it was only 0.5 percent.[22] Such dramatic decline in the share of new jobs suggests a real barrier to workers moving into newly created, higher-value-added jobs.

It is important to point out that the long-term impact on the job market remains something of a puzzle. There are still many "jobs" in the economy. Indeed, the current unemployment rate in the US is at historical lows. But many of the jobs available to people with lower skill levels offer inconsistent hours, are contract-based, and are in lower-value-added sectors that do not represent a shift toward the higher-value, higher-paying jobs associated with new technologies. MIT economist David Autor suggests the job market is polarized, with continued availability of low-paying, low-skilled, and unstable jobs on the low end and high-skilled jobs on the high end, but a disappearance of the administrative, clerical, labor, and craft jobs that once allowed huge portions of the middle class to thrive.[23] The focus on jobs alone, therefore, may be missing the point. The more relevant issue may be the ability to create new jobs at scale in high-income sectors.

Call to Action

Absent direct policies and concerted action to address these risks, the economy of the future could fail to create a broad base of well-paying jobs with stable wage growth and benefits. As important as these issues are for the US and other developed nations to examine and address, they also have important ramifications for developing countries. Many nations have been on a deliberate path to mimic the economic growth path taken by developed countries. These countries have embraced the process of advancing urbanization and shifting the workforce from agriculture to manufacturing to services, with a view toward preparing more of the labor force to participate in higher-value, higher-income activities. Yet technological developments may have created a break in the development chain, rendering that path unpassable by emerging economies. For a continent like Africa, where the population is expected to grow four-fold by 2100, the risk of a burgeoning labor force with no jobs to do is real.

A future of work that is broad-based and broadly accessible by the majority of people will require a range of purposeful interventions by government, the private sector, and the civic sector. Society must renew its focus on workforce preparation so that more people can participate in new, higher-value activities. This is important not only for future generations, but also for existing workers likely to be displaced and affected by changes. Closing the existing skills gap will require both investments and changes in the traditional PK-12 basic education system, as well as changes to our ideas about when education "stops." Society needs to adopt a view of continuous learning and reskilling throughout a person's life, with emphasis on more demand-driven and vocational training and a greater focus on team-based skills and cross-functional and disciplinary integration.

As important and necessary as these education investments are, their impact will be limited if not undertaken in concert with a holistic range of economic and social policies. This may well include a new look at a range of economic levers, including changes to the tax code, increases to the minimum wage, reform of unemployment and wage insurance, effective savings programs, and modified compensation and tax structures. Beyond income, more may need to be done to address the cost of and access to decent housing and health care at the community level.

The future of work is already here. Now is the time to mobilize.

Diana Farrell is the founding President and Chief Executive Officer of the JPMorgan Chase Institute. Previously, she was a Senior Partner at McKinsey & Company where she was the Global Head of the McKinsey Center for Government and the McKinsey Global Institute. Ms. Farrell served in the White House as Deputy Director of the National Economic Council and Deputy Assistant to the President on Economic Policy for 2009-2010. During her tenure, she led interagency processes and stakeholder management on a broad portfolio of economic and legislative initiatives. She coordinated policy development and stakeholder engagement for innovation and competition strategies broadly, and led financial policy initiatives including the passage of major legislation. She also served as a member of the President's Auto Recovery Task Force. Ms. Farrell currently serves on the Board of Directors for eBay, The National Bureau for Economic Research (NBER), The Urban Institute, and Washington International School. In addition, she is a Trustee Emeritus of Wesleyan University, a Trustee of the Trilateral Commission, and served as a Co-Chair of the World Economic Forum's Council on Economic Progress. Ms. Farrell is a member of the Council on Foreign Relations, Economic Club of New York, Bretton Woods Committee, and Aspen Strategy Group. Ms. Farrell holds a M.B.A. from Harvard Business School, and has a B.A. from Wesleyan University from where she was awarded a Distinguished Alumna award.

[1] Larry Hardesty, "Lessons learned from MITx's prototype course," July 16, 2012, news.mit.edu/2012/mitx-edx-first-course-recap-0716.

[2] Mae Ryan, Cade Metz, and Rumsey Taylor, "How Robot Hands Are Evolving to Do What Ours Can," July 30, 2018, *New York Times*, www.nytimes.com/interactive/2018/07/30/technology/robot-hands.html.

[3] YiLi Chien and Paul Morris, "Is U.S. Manufacturing Really Declining?" On the Economy Blog, Federal Reserve Bank of St. Louis, April 11, 2017, www.stlouisfed.org/on-the-economy/2017/april/us-manufacturing-really-declining.

[4] Carl Benedikt Frey and Michael A. Osborne, *The Future of Employment: How Susceptible Are Jobs to Computerisation?* Oxford Martin, 2013, 114.

[5] Malcolm Foster, "Aging Japan: Robots May Have Role in Future of Elder Care," *Reuters*. March 27, 2018, www.reuters.com/article/us-japan-ageing-robots-widerimage/aging-japan-robots-may-have-role-in-future-of-elder-care-idUSKBN1H33AB.

[6] Erik Brynjolfsson, Xiang Hui, and Meng Liu, *Does Machine Translation Affect International Trade? Evidence from a Large Digital Platform,* NBER Working Paper No. 24917, 2018.

[7] Diana Farrell, Fiona Greig, and Amar Hamoudi, *The Online Platform Economy in 2018: Drivers, Workers, Sellers and Lessors*, JPMorgan Chase Institute, 2018.

[8] Diana Farrell, Fiona Greig, and Amar Hamoudi. *The Online Platform Economy in 2018: Drivers, Workers, Sellers and Lessors*, JPMorgan Chase Institute, 2018.

[9] Diana Farrell, Fiona Greig, and Amar Hamoudi, *The Online Platform Economy in 2018: Drivers, Workers, Sellers and Lessors*, JPMorgan Chase Institute, 2018.

[10] Dao Mai Chi, Mitali Das, Zsoka Koczan, and Lian Weicheng. "Drivers of Declining Labor Share of Income," IMFBlog, April 12, 2017, blogs.imf.org/2017/04/12/drivers-of-declining-labor-share-of-income/.

[11] International Labour Organization, International Monetary Fund, Organisation for Economic Co-operation and Development, and World Bank Group, *Income Inequality And Labour Income Share in G20 Countries: Trends, Impacts and Causes,* September 2015, www.oecd.org/g20/topics/employment-and-social-policy/Income-inequality-labour-income-share.pdf.

[12] Congressional Budget Office, "The Distribution of Household Income, 2014," 2014.

[13] National Low Income Housing Coalition, *Out of Reach: The High Cost of Housin*g, 2018, nlihc.org/sites/default/files/oor/OOR_2018.pdf.

[14] Jonathan D. Ostry, Andrew Berg, and Charalambos G. Tsangarides, *Redistribution, Inequality, and Growth*, International Monetary Fund, April 2014.

[15] Amy K. Glasmeier, "New Data Up: Calculation of the Living Wage," MIT Living Wage Calculator, January 26, 2018, livingwage.mit.edu/articles/27-new-data-up-calculation-of-the-living-wage.

[16] Irene Tung, Yannel Lathrop, and Paul Sonn, *The Growing Movement for $15*, National Employment Law Project, November 2015, www.nelp.org/wp-content/uploads/Growing-Movement-for-15-Dollars.pdf.

[17] National Low Income Housing Coalition, *Out of Reach: The High Cost of Housing*, 2018, nlihc.org/sites/default/files/oor/OOR_2018.pdf.

[18] McKinsey Global Institute, *Poorer Than Their Parents? Flat Or Falling Incomes in Advanced Economies*, McKinsey & Company, July 2016.

[19] Bureau of Labor Statistics, "Real Earning – July 2018," 2018, www.bls.gov/news.release/pdf/realer.pdf.

[20] Diana Farrell and Fiona Greig, *The Monthly Stress-Test on Family Finances*, JPMorgan Chase Institute, 2017.

[21] Board of Governors of the Federal Reserve System, *Report on the Economic Well-Being of U.S. Households in 2015,* May 2016, www.federalreserve.gov/2015-report-economic-well-being-us-households-201605.pdf.

[22] Carl Benedikt Frey and Michael A. Osborne, *The Future of Employment: How Susceptible Are Jobs to Computerisation?* Oxford Martin, 2013, 114.

[23] David Autor, *The Polarization of Job Opportunities in the U.S. Labor Market: Implications for Employment and Earnings,* Center for American Progress, April 2010.

"As technological change is accelerating, the United States needs to show the same level of public and private commitment to meeting this challenge as it showed when the country transitioned from an agrarian to an industrial economy just over 100 years ago."

–PENNY PRITZKER AND EDWARD ALDEN

The National and Economic Security Imperative of Helping More Americans Adapt and Thrive

Penny Pritzker
Chairman and Founder
PSP Partners

Edward Alden
Bernard L. Schwartz Senior Fellow
Council on Foreign Relations

Introduction

The United States today faces twin challenges—building its global leadership in the next generation of transformative technologies and rebuilding economic opportunities for more of its citizens. The first cannot be done successfully without also doing the second. Innovation and competition are the great drivers of prosperity, but they have also created a growing gap between the economic winners and those struggling to get by. Unemployment in the United States has fallen below 4 percent, and the well-being of Americans has been improving as the economy continues to grow at a strong pace. Yet four in ten US households still report that they are unable to cover an unexpected $400 expense without borrowing money or selling something they own.[1] More than a decade after the last recession, economic insecurity remains widespread.

This continued economic insecurity poses a growing and fundamental threat to America's economic competitiveness and national security. While technology and global competition have helped raise incomes and living standards around the world, they have also created huge new challenges in the labor markets of many of the advanced economies, from the disappearance of once well-paying manufacturing jobs to the growth of the gig economy and other contingent work that comes without traditional employment benefits. Americans need far better access to the education and retraining opportunities required to prosper in this rapidly changing economy, and government support systems must be updated so that working Americans can again have greater confidence about their futures. The reality is that for more than thirty years we have failed as a nation in this regard.

In the United States, where the social safety net is especially porous and support for job retraining is weaker than in any other wealthy country, labor market disruption has already contributed to social and political upheaval. Donald Trump was elected president in 2016 on a platform that promised greater restrictions on both international trade and immigration to the United States, blaming both for the economic challenges facing many Americans. Since taking office, the president has approved the largest increase in tariffs on imports since the 1930s, has slashed refugee admissions to their lowest levels since the refugee program was created in 1980, and has taken a series of steps to reduce the entry of highly skilled immigrants to the United States.

Such restrictions on trade and immigration will erode America's technological and economic leadership. Immigrants today—many of them initially attracted by the high quality of American universities—are more than twice as likely to start a business as native-born citizens; from 1996 to 2011, the business start-up rate for immigrants increased by more than half, while the native-born start-up rate fell by 10 percent, to a three-decade low.[2] Of the eighty-seven start-up companies that had reached a value of more than $1 billion by 2016, immigrants founded more than half, and over 70 percent had immigrants as part of the top management and product development teams.[3] On trade, internationally engaged American companies—those that both export and invest abroad—are America's most innovative companies, accounting for nearly three-quarters of private sector research and development.[4] The success of these firms depends on markets that are open to both trade and investment. And while the United States has imposed few restrictions on the deployment of new technologies, some 75 percent of Americans today are worried about a world in which computers and robots do more of the work, fearing for their job prospects, their family's future, and that inequality will worsen.[5]

Polls indicate that the public does not favor tariffs on imports, sharp restrictions on immigration, or regulations that curb technological innovation. But the public is wary about what technology and global competition mean for their jobs and their future. Public support for economic openness can no longer be assumed; it must be rebuilt. That requires rebuilding the connection between economic openness, innovation, and better work and life opportunities for Americans. The US education system must do a better job of preparing Americans for the world of work by expanding career-related offerings; better support is needed to allow mid-career workers, or those displaced by technology or trade competition, to return to school and retrain for new careers; and the benefits that are now available to most full-time workers—health care, sick leave, vacation pay—need to be available to everyone with a job. Improving and rebuilding the links among education and workforce training, good jobs, and greater economic security is vital to our future security and economic competitiveness. As technological change is accelerating, the United States needs to show the same level of public and private commitment to meeting this challenge as it showed when the country transitioned from an agrarian to an industrial economy just over 100 years ago.

Meeting the twin challenges of technological leadership and rebuilding opportunity must be the primary goals for US economic policy. Given the seismic forces of innovation, automation, and globalization, the nature of work is fundamentally changing; we must help more Americans adapt, adjust, and thrive. America needs a more forward-looking, comprehensive economic competitiveness strategy that includes an innovation leadership agenda, modernization of our workforce training and education systems, immigration reform, and expanded multilateral trade. If the United States fails to meet these challenges, it will have neither the resources nor the political support needed to play a large global role.

The Technology Challenge

First, the United States must develop a strategy for leading in the next generation technologies. This is critical for economic leadership and national security. China is moving forward with its *Made in China 2025* program, investing heavily in research and development aimed at giving the country a leading position in electric vehicles, robotics, high-speed rail, semiconductors, and satellites. In July 2017, China rolled out plans to gain global leadership in Artificial Intelligence (AI) by 2030. Many other countries, including Canada, Finland, Singapore, and Sweden have also developed national strategies to support the development of AI.

Artificial intelligence—which is broadly defined as machines that can learn and improve their ability to carry out tasks—will be critical to driving growth and enhancing productivity in the future. It holds out the promise of reducing use of scarce natural resources, lowering costs for a wide range of businesses, and augmenting human capabilities. It will also be vital for future US defense needs. The United States currently leads the world in development of AI-driven technologies, but future leadership is far from assured. China and the United States, in particular, will vie for global leadership. US companies such as Facebook, Google, and Amazon are still out front, but Chinese companies like Alibaba and Tencent are catching up rapidly. China has some significant advantages in developing AI; its huge population generates enormous quantities of data that have the potential to speed up improvements in

AI applications. Most Chinese are already using mobile pay as well; in the US, the adoption has been slower.[6] China's investments in AI-related R&D have been growing much faster than in the United States. The Obama administration took steps toward developing a national strategy for AI in 2016, from which China has borrowed liberally. The Trump administration has not moved forward with developing a similar comprehensive strategy, though it has highlighted the importance of AI in a series of approach documents, and the Department of Defense has pledged $2 billion in research support for AI over the next five years.[7]

The stakes are extremely high. AI and robotics will be the commanding heights of the economy of the future, as important as steel, cars, and aircraft were in the twentieth century. AI will be a driver of future economic growth and competitiveness, reducing costs and improving efficiencies across a range of industries. Much as its early leadership in information technology was a key ingredient of US military superiority, AI has a wide variety of military applications, including logistics, cybersecurity, and cyber-warfare as well as autonomous and semi-autonomous weapons systems. This is not a technology race that the United States can afford to lose.

If the United States is to retain its edge in the new technologies, it must also find a way to adapt to the labor market impacts. As the United States should have learned from the disruptions brought about by growing trade competition, the failure to adjust to the labor market impacts of a rapidly changing economy will undermine public confidence in political leadership.[8] While no one fully knows what the scale of disruption will be, one thing is clear: as robotics, AI, and other new technologies begin to meet their potential, many of the jobs that human beings have long taken for granted will no longer be needed. Estimates of the job disruption vary widely, but even the more modest ones predict significant changes in the labor market. McKinsey estimates that about 15 percent of jobs will be fully automatable by 2030, while some 60 percent of jobs are likely to see some of their tasks automated.[9] Those vulnerable include retail workers, truck drivers and other transportation workers, and a range of white collar jobs from legal research to financial services. As the technology develops, few occupations are likely to be untouched. AI and robotics will have significant applications in health care, education, agriculture, warehousing, and food and cleaning services. The United States is far from alone in facing this challenge; McKinsey has estimated that half of all current jobs in China could be automated, making it the most susceptible country to technology-induced job disruption.[10] Technology will challenge the unity and stability of countries everywhere, developed and developing alike.

The challenge for each country is to find ways to ensure that such job disruption is managed successfully, and that the gains that AI and other forms of automation will bring in increased productivity and efficiency are broadly shared. The United States, unfortunately, faces many obstacles in meeting that challenge.

The Workforce Challenge

The US economy is amid its second longest expansion in 60 years. For the first time since 2000, when data on job openings began to be collected nationwide, the number of unemployed Americans is slightly fewer than the number of open jobs.[11] The strong job market has been good news for many Americans; while overall wage growth remains surprisingly weak given the tight labor market, wages have risen most strongly among the lowest-paid workers.[12] But the weakness in wage growth suggests that employers continue to have many options other than raising wages, such as investing in labor-saving technologies and outsourcing work. Nearly two decades into the twenty-first century, the United States must do more to demonstrate that the new economy can bring broad benefits for its population.

America's rise as a manufacturing power in the nineteenth and twentieth centuries was accompanied by a broad improvement in living standards. Much of that was driven by mass education; by World War II, the United States was graduating more than half its citizens from high school, the first country to reach that benchmark. It led again after the war by offering free college education to every one of the nation's 16 million veterans. In the mid-twentieth century, the United States also developed the broad social programs that brought greater economic security to most Americans, including Social Security, unemployment insurance, Medicare, and Medicaid. Continued US leadership as a technology and innovation powerhouse will require that the benefits be spread similarly to the vast majority of Americans.

The United States has continued to make significant strides in raising American education levels; nearly two-thirds of Americans today get at least some education beyond high school. But educational offerings have not been nimble enough to keep up with the rapid changes in the economy. Americans not only need more education, but better-targeted education that leads to better work opportunities, even as the target will continue to move as new technologies are adopted in the workplace. There are currently, for example, more than 300,000 job openings in cybersecurity, a number that is expected to rise to more than 500,000 by 2021.[13] There are also many openings for jobs that require training beyond high school—a two-year associate's degree, specialized certificates, or apprentice training—but not a four-year college degree.

The cost of higher education is a huge and growing barrier for many young Americans. Tuition increases have far outstripped the overall cost of living, and students are going deeper and deeper into debt to finance their education. Total student debt has grown over the past decade from $600 million to more than $1.5 trillion, and the average student graduates with $37,000 in debt.[14] Even for those who acquire the credentials to succeed in the economy, many barriers remain. Roughly 25 percent of the workforce today requires some sort of state-level occupational license, up from just 5 percent in the 1950s. Most of the country's three million teachers, for example, hold state credentials that are not automatically recognized in other states, limiting their ability to move from one job to the next. Finally, the United States does far too little to offer new education and retraining options for mid-career workers who lose their jobs or want to move to better ones. US support for mid-career retraining is the weakest of all the advanced economies.[15]

While the United States must educate its people for the jobs of the future and create a culture of lifelong learning, it also needs to update its social safety net to provide greater economic security for the way that more and more Americans are working. Workers in alternative arrangements, including temporary employees, independent contractors, and so-called "gig economy" workers, now make up about 16 percent of the workforce.[16]

The US safety net was built for the twentieth century compact, one in which benefits were tied closely to employers able to offer stable, full-time employment, often throughout an employee's working life. More than 80 percent of full-time workers have access to a full range of employment benefits, including health insurance, retirement benefits, and paid vacation and sick leave. Among part-time workers, however, fewer than 40 percent enjoy the same benefits.[17] Very few independent contractors or gig economy workers enjoy any employment benefits at all. This is a perverse situation that creates significant incentives for companies not to hire full-time employees, while denying a level of economic security to millions of working Americans.

The Affordable Care Act helped fill a hole in the benefits system by providing more options for the self-employed and for others who lack employer-based health care. Roughly one in five Obamacare recipients who signed up in the first year of the program in 2014 was a small business owner or self-employed.[18] But the full range of employment benefits needs to be available to Americans regardless of the nature of their employment relationship. Many models have been proposed to create a system of "portable" benefits that would cover all working Americans.

Immigration and Trade: Moving in the Wrong Direction

The failure to deal with these labor market challenges has already produced significant public pressure for retrenchment on immigration and trade. Donald Trump spoke to many of these fears during his 2016 presidential campaign, and since capturing the White House, he has been implementing a series of measures to protect American companies and workers from trade competition and competition from new immigrants.

On immigration, while most of the public attention has focused on Trump's so-far unfilled promise to expand the wall along the US border with Mexico, his administration has been using its regulatory powers to discourage and reduce legal immigration to the United States. The administration has taken specific aim at the H-1B program, which is the temporary work visa for immigrants with at least a four-year college education. Since its creation in 1990, the program has become the main vehicle for skilled immigrants seeking green cards to remain in the United States—including many scientists and engineers. The administration has also tightened the requirements that allow international graduates in science, technology, engineering, or mathematics from US universities to continue working while they're trying to apply for a work visa. Under the April 2017 executive order "Buy American, Hire American," the president ordered officials to "rigorously enforce and administer the laws governing entry into the United States of workers from abroad."[19] Since the order, denials of company applications for H-1B workers have increased by 40 percent, while requests by the government for additional information—which slows application processing times—have doubled.[20]

The administration is taking other measures as well. It is moving to rescind a 2015 rule that permitted the spouses of H-1B workers to hold jobs. Due to quota restrictions, many H-1B workers can face waits of ten years or longer before they can receive green cards permitting them to remain permanently. The Trump administration move will force nearly 150,000 spouses of H-1B workers to leave their jobs, harming not just those individuals, but the companies that employ them. The administration is also rolling back the international entrepreneur rule, another Obama-era innovation that permitted immigrants who started their own companies to remain in the United States for five years. The new administration is also sharply cutting admissions of refugees and working to eliminate the Deferred Action for Childhood Arrivals (DACA) program, which has allowed nearly two million young people who were brought to the country illegally by their parents to enter the workforce or further their education. The effect of these initiatives has been to make the United States a far less attractive destination for migrants and has led some companies to move research and other facilities outside the United States in order to attract the world's best scientists and engineers.

Other countries are moving to take advantage of the increasingly harsh immigration climate in the United States; in 2017, Canada created a "global talent stream" program that allows fast-growing Canadian-based companies to hire qualified foreign workers with hard-to-find qualifications in as little as two weeks. While the program is still a small pilot, it underscores Canada's commitment to attracting the best immigrants in high-demand fields such as advanced manufacturing, clean technology, animation and on-line gaming. A recent survey of high-tech firms in Toronto showed a spike in international applicants, with an average 45 percent increase in international hires; the city is now North America's fastest growing tech market. The survey found that one of the key reasons was due to US immigration policies. Kate Mitchell, the founder of Scale Venture Partners in Silicon Valley and former chair of the National Venture Capital Association, says that "the fight over tech talent is not something that is coming in the future. It's happening right now. And we are losing."

The administration's trade policies have moved in the same troubling direction. President Trump pulled the United States out of the Trans-Pacific Partnership (TPP) trade agreement on his third day in office, cutting the US out of a deal that would have encompassed nearly 40 percent of the global economy. The new rules in the TPP covering digital trade and intellectual property protection would have been enormously helpful for US advanced technology companies, opening new markets abroad, helping to protect their data and trade secrets, and fostering the growth of a global digital economy.[21] The administration instead has put in place a range of new import-restricting tariffs on steel, aluminum, solar panels, and nearly half of all Chinese exports to the United States. Those tariffs are raising costs not just for US consumers, but for manufacturers as well. Ford Motor Company reported in September that the new tariffs have cost the company $1 billion in lost profits already.[22] A new round of tariffs against China, which the administration has threatened, would hit a significant portion of the high-tech supply chain, including assembly of smart phones.

The United States certainly faces problems around the world, and especially in China, with ensuring that foreign markets are open to US goods, services, and investment, and that US intellectual property is protected. And the administration has claimed some progress in renegotiating trade agreements, including the US-Korea trade agreement and the North American Free Trade Agreement (NAFTA) with Canada and Mexico. But the administration's approach to trade has been narrowing rather than expanding opportunities for US-based companies and their American workers. That increases the risk that the companies will look to other countries for the opportunities to develop and commercialize the next generation of technologies. An America that is protecting its companies behind tariff walls rather than competing with the world will not be the economic leader of the future.

Meeting the Twin Challenges

A comprehensive economic strategy begins first with the United States ensuring it leads in the next generation of transformative technologies. It has all the right ingredients—the world's top scientific talent, its best universities, its most innovative companies. But those advantages cannot be taken for granted. The United States also needs supportive government policies that will build on those advantages. And government needs to play a big role, working closely with the private sector where appropriate, in developing far better education and workforce strategies that spread the benefits of technological advances more broadly to Americans.

Here is what the US must do to create a strong and comprehensive economic strategy for the twenty-first century:

America must lead in artificial intelligence. While the United States has a strong position in AI and has attracted or developed many of the world's best AI scientists, it is facing a rising challenge from other countries, especially China. AI, for all its potential benefits, will be disruptive to many jobs. The United States needs an AI competitiveness strategy that will maintain US technological leadership while mitigating the negative impacts on the job market. The Trump administration has taken small steps toward solidifying its own AI strategy, such as convening officials, business leaders, and academics for a DC summit last May. But the consensus from researchers and companies is that the White House is not doing nearly enough. The US government needs to provide far greater support, for example, for R&D on artificial intelligence. Kai-Fu Lee, the CEO of Sinovation Ventures, which is investing in AI development in both the US and China, says that doubling that funding would be "a good start."

America must do more to support and help its workers succeed. Research funding alone is not sufficient. The United States also must develop and implement more robust education and workforce solutions that will both develop the next generation of scientific and engineering leaders and prepare all Americans for the workplace disruptions that will take place as the new technologies mature. The country has a long way to go; the new World Bank Human Capital Index—which looks at how well countries are doing in advancing the education and health of their people—ranked the United States twenty-fifth in the world, among the lowest of the high-income countries.[23] The US education system too often falls short in preparing Americans for the faster-growing, better-paying jobs that require some mixture of soft skills, specific technical skills, some practical on-the-job experience, and a capacity for lifelong learning. The growing student debt burden must also be addressed to open the door to higher education for more Americans.

The Trump administration has taken some encouraging steps, including efforts to expand apprenticeship training.[24] But the United States still lags far behind countries like Germany, Switzerland, Singapore, and many others in addressing its workforce challenges. The tax bill passed by Congress in 2017 was a huge opportunity missed. While the bill cut corporate taxes and increased incentives for companies to invest in research and new technologies, it did nothing to encourage companies to upgrade their workforce training. Companies should be rewarded for investing in their workforce in the same way they are rewarded for investing in research and technology. Congress should also create lifelong training accounts, which are employee-owned accounts with matching funds from employers and government-matched savings plans, which could be used at any time during an employee's career to pay for the education needed to upgrade skills and move to better jobs.[25]

America needs robust immigration reform to attract and keep talent. Sensible immigration policies are vital to winning the race for next-generation technologies. The United States continues to attract an outsized share of the world's best-educated and more entrepreneurial immigrants. The size of the college-educated immigrant population in the United States has more than tripled since 1990; some 45 percent of new immigrants since 2010 have had a four-year bachelor's degree or higher.[26] But the Trump administration is moving the country in the wrong direction, embracing a misguided theory that immigrants are stealing jobs from Americans and holding down wages. Instead, the United States should be doubling down on welcoming the world's best and brightest, including an easier path for foreign students who graduate with science and engineering degrees from US universities to remain in the country. To win the race for the future, the United States needs to be developing the education and skills of Americans to the fullest and continuing to attract the world's smartest immigrants.

America needs a smarter approach to trade. The United States must aggressively pursue larger multilateral trade agreements rather than just bilateral deals. And it needs to move away from unilateral tariffs that weaken respect for trade rules and work with other countries to meet global trade challenges, such as ensuring that China and other rising economies adhere to the rules-based trading order. The other eleven TPP countries are moving forward without the United States, while the European Union has entered new trade agreements with Canada and Japan. Unless the United States has access to the world's markets on the most favorable terms, companies will look to invest in other places. US companies need to know they can grow and expand in this country and do business anywhere in the world. The United States must embrace multilateral trade, and a good first step would be to rejoin the TPP. The United States should also recommit to strengthening the World Trade Organization (WTO) rather than pursuing unilateral policies that are weakening long-established rules and norms for international trade.

America must create more inclusive economic growth. Finally, and most importantly, the United States has failed for more than three decades to create more inclusive prosperity for all its citizens. We need a broader generational cultural change in our approach to education, workforce training, and the basic link between learning and a job. Americans will need to reimagine their careers, embracing the need for lifelong learning, and government and employers will need to help make more opportunities available and affordable. The United States needs to restructure the relationship between jobs and benefits. With so much of the current and projected job growth in part-time, contingent, or gig employment, it no longer makes sense to tie employment benefits such as retirement and sick leave to particular jobs. Rather, portable systems of employment benefits should be introduced that follow the individual from job to job. Good models have been suggested by the Aspen Institute, and several states, including Washington, California, New York, and New Jersey, are experimenting with new schemes.[27] We should embrace this type of workplace innovation.

Conclusion

The United States won the twentieth century because it finally got the big challenges right—education, scientific excellence, innovation, immigration, and trade. Yet, in recent decades we have not done all that we can as a nation to adapt government policies and approaches to the rapid pace of economic and technological change. Too many Americans have been left behind by the rapid changes in the economy, without the necessary tools and resources to prosper. The reality is we can do better.

With diminishing opportunities, it is not surprising that Americans have been susceptible to populist promises. The United States has been here before and risen to such challenges in the past. We must do so again as our national and economic security depend on it.

Penny Pritzker is the founder and chairman of PSP Partners and its affiliates, Pritzker Realty Group, PSP Capital, and PSP Growth. From June 2013 through January 2017, she served as US secretary of commerce. Secretary Pritzker is an entrepreneur, civic leader, and philanthropist, with more than 25 years of experience in numerous industries. She founded Vi Senior Living and co-founded The Parking Spot and Artemis Real Estate Partners. She is the former chairman of the board of TransUnion and is a past board member of Hyatt Hotels Corporation, Wm. Wrigley Jr. Company, Marmon Group, and LaSalle Bank Corporation. Secretary Pritzker is currently a member of the board of Microsoft, a member of the Harvard Corporation, chairman of the board of trustees of the Carnegie Endowment for International Peace, a member of the Aspen Economic Strategy Group and Aspen Strategy Group, and a co-chair of the Cyber Readiness Institute. Secretary Pritzker was formerly a member of the board of the Council on Foreign Relations, the board of trustees of Stanford University, the Harvard University Board of Overseers, and she founded Skills for America's Future. She also served on President Obama's Council on Jobs and Competitiveness and his Economic Recovery Advisory Board. Secretary Pritzker and her husband, Dr. Bryan Traubert, co-founded The Pritzker Traubert Foundation, a private foundation that works to foster increased economic opportunity for Chicago's families. Secretary Pritzker earned a Bachelor of Arts degree in economics from Harvard University and a Juris Doctor and Masters of Business Administration from Stanford University.

Edward Alden is the Bernard L. Schwartz senior fellow at the Council on Foreign Relations (CFR), specializing in US economic competitiveness, trade, and immigration policy. He is the author of the book *Failure to Adjust: How Americans Got Left Behind in the Global Economy*, which focuses on the federal government's failure to respond effectively to competitive challenges on issues such as trade, currency, worker retraining, education, and infrastructure. Mr. Alden's previous book, *The Closing of the American Border: Terrorism, Immigration, and Security Since 9/11*, was a finalist for the Lukas Book Prize, for narrative nonfiction, in 2009. Previously, he was the Washington bureau chief for the *Financial Times* and prior to that was the newspaper's Canada bureau chief, based in Toronto. He worked as a reporter at the *Vancouver Sun* and was the managing editor of the newsletter *Inside U.S. Trade*, widely recognized as a leading source of reporting on US trade policies. He has won several national and international awards for his reporting and has been published in *Foreign Affairs, Foreign Policy, Fortune, Los Angeles Times, New York Times, Toronto Globe and Mail, Wall Street Journal,* and *Washington Post*. Mr. Alden has a bachelor's degree in political science from the University of British Columbia. He earned a master's degree in international relations from the University of California, Berkeley, and pursued doctoral studies before returning to a journalism career. Mr. Alden is the winner of numerous academic awards, including a Mellon fellowship in the humanities and a MacArthur Foundation graduate fellowship.

Ms. Pritzker was co-chair and Mr. Alden was project director for the Council on Foreign Relations Task Force Report *The Work Ahead: Machines, Skills and US Leadership in the Twenty-First Century.*

[1] Board of Governors of the Federal Reserve System, "Report on the Economic Well-Being of US Households in 2017," May 2018, www.federalreserve.gov/publications/files/2017-report-economic-well-being-us-households-201805.pdf.

[2] Adam Bluestein, "The Most Entrepreneurial Group in America Wasn't Born in America," *Inc. Magazine*, February 2015, www.inc.com/magazine/201502/adam-bluestein/the-most-entrepreneurial-group-in-america-wasnt-born-in-america.html.

[3] Stuart Anderson, "Immigrants and Billion Dollar Startups," *National Foundation for American Policy,* March 2016, nfap.com/wp-content/uploads/2016/03/Immigrants-and-Billion-Dollar-Startups.NFAP-Policy-Brief.March-2016.pdf.

[4] Matthew J. Slaughter, "The 'Exporting Jobs' Canard," *Wall Street Journal*, June 14, 2017, www.wsj.com/articles/the-exporting-jobs-canard-1497482039.

[5] Erica Smith and Monica Anderson, "Americans Attitudes Towards a Future in Which Robots and Computers Can Do Many Human Jobs," Pew Research Center, October 4, 2017, www.pewinternet.org/2017/10/04/americans-attitudes-toward-a-future-in-which-robots-and-computers-can-do-many-human-jobs/.

[6] Eric Johnson, "If They Don't Want to Lose Their Jobs to a Machine, Doctors Will Need to Become Compassionate 'Human Connectors,'" Recode, September 17, 2018, www.recode.net/2018/9/17/17867990/kai-fu-lee-ai-superpowers-book-artificial-intelligence-jobs-doctors-kara-swisher-decode-podcast.

[7] Drew Harwell, "Defense Department Pledges Billions Toward Artificial Intelligence Research," *Washington Post,* September 7, 2018, www.washingtonpost.com/technology/2018/09/07/defense-department-pledges-billions-toward-artificial-intelligence-research/?noredirect=on&utm_term=.92395977c54c.

[8] Edward Alden, *Failure to Adjust: How Americans Got Left Behind in the Global Economy,* Lanham, MD: Rowman & Littlefield, 2017.

[9] McKinsey Global Institute, *Jobs Lost, Jobs Gained: Workforce Transitions in a Time of Automation*, December 2017, www.mckinsey.com/~/media/mckinsey/featured%20insights/future%20of%20organizations/what%20the%20future%20of%20work%20will%20mean%20for%20jobs%20skills%20and%20wages/mgi-jobs-lost-jobs-gained-report-december-6-2017.ashx.

[10] McKinsey Global Institute, *A Future That Works: Automation, Employment, and Productivity,* January 2017, www.mckinsey.com/~/media/McKinsey/Featured%20Insights/Digital%20Disruption/Harnessing%20automation%20for%20a%20future%20that%20works/MGI-A-future-that-works_Full-report.ashx.

[11] Sho Chandra, "US Job Openings Rise to Record, Exceeding Number of Unemployed," *Bloomberg*, June 5, 2018, www.bloomberg.com/news/articles/2018-06-05/job-openings-in-u-s-increased-in-april-to-record-6-7-million.

[12] Economic Policy Institute, "Average Wage Growth Continues to Flatline in 2018, While Low-Wage Workers and Those with Relatively Lower Levels of Educational Attainment See Stronger Gains," July 18, 2018, www.epi.org/blog/average-wage-growth-continues-to-flatline-in-2018-while-low-wage-workers-and-those-with-relatively-lower-levels-of-educational-attainment-see-stronger-gains/.

[13] CompTIA, "US Cybersecurity Worker Shortage Expanding, New Cyberseek Data Reveals," June 6, 2018, www.comptia.org/about-us/newsroom/press-releases/2018/06/06/us-cybersecurity-worker-shortage-

expanding-new-cyberseek-data-reveals; CSO, "Cybersecurity Labor Crunch to Hit 3.5 Million Unfilled Jobs by 2021," June 6, 2017, www.csoonline.com/article/3200024/security/cybersecurity-labor-crunch-to-hit-35-million-unfilled-jobs-by-2021.html.

[14] Jillian Berman, "Student Debt Just Hit $1.5 Trillion," *MarketWatch*, May 12, 2018, www.marketwatch.com/story/student-debt-just-hit-15-trillion-2018-05-08.

[15] Organization for Economic Cooperation and Development, "Active Labour Market Policies: Connecting People with Jobs," www.oecd.org/employment/activation.htm.

[16] Lawrence F. Katz and Alan B. Krueger, "The Rise and Nature of Alternative Work Arrangements in the United States, 1995-2015," National Bureau of Economic Research, September 2016, www.nber.org/papers/w22667.

[17] US Bureau of Labor Statistics, March 2017, cited in John Engler et al., *The Work Ahead: Machines, Skills, and US Leadership in the Twenty-First Century*, Council on Foreign Relations Independent Task Force No. 76, April 2018.

[18] Dan Mangan, "Small-Business Owners and Self-Employed More Likely to Buy Obamacare Plans," *CNBC*, January 12, 2017, www.cnbc.com/2017/01/12/small-business-owners-self-employed-bought-many-obamacare-plans.html.

[19] Exec. Order. No. 13788, April 18, 2017, www.whitehouse.gov/presidential-actions/presidential-executive-order-buy-american-hire-american/.

[20] Nelson D. Schwartz and Steve Lohr, "Companies Say Trump Is Hurting Business by Limiting Legal Immigration," *New York Times*, September 2, 2018, www.nytimes.com/2018/09/02/business/trump-legal-immigration-h1b-visas.html.

[21] IBM Government and Regulatory Affairs, "Viewpoint from IBM CEO Ginni Rometty: Data, IBM and the Future of Global Trade," www.ibm.com/ibm/ibmgra/data-ibm-future-trade-09162016.html.

[22] Nick Carey and David Shepardson, "Trump Metals Tariffs Will Cost Ford $1 Billion in Profits, CEO Says," *Reuters*, September 26, 2018, www.reuters.com/article/us-ford-motor-tariffs/trump-metals-tariffs-will-cost-ford-1-billion-in-profits-ceo-says-idUSKCN1M61ZN.

[23] World Bank Group, *The Human Capital Project,* October 2019, www.worldbank.org/en/publication/human-capital?mod=article_inline.

[24] Exec. Order. No. 13845, July 19, 2018, www.whitehouse.gov/presidential-actions/executive-order-establishing-presidents-national-council-american-worker/.

[25] Edward Alden and Robert E. Litan, *A New Deal for the Twenty-First Century,* Council on Foreign Relations, May 2017, cfrd8-files.cfr.org/sites/default/files/report_pdf/Renewing_America_Twenty_First_Century_Deal_OR.pdf.

[26] Migration Policy Institute, "MPI Reveals Striking Finding: Nearly Half of Immigrant Adults Arriving in US since 2011 Have College Degree, a Sharp Increase over Earlier Groups," June 1, 2017, www.migrationpolicy.org/news/mpi-reveals-striking-finding-nearly-half-immigrant-adults-arriving-us-2011-have-college-degree.

[27] Robert Maxim and Mark Muro, "Rethinking Worker Benefits for an Economy in Flux," Brookings Institution, March 30, 2018, www.brookings.edu/blog/the-avenue/2018/03/29/rethinking-worker-benefits-for-an-economy-in-flux/.

“While we still can, we should articulate, encode, and embed the best of humanity, while protecting against the worst, into our AI companions.”

–JOHN DOWDY & CHANDRU KRISHNAMURTHY

Defense in the 21st Century: How Artificial Intelligence Might Change the Character of Conflict

John Dowdy
Senior Partner
McKinsey & Company

Chandru Krishnamurthy
Senior Partner
McKinsey & Company

As artificial intelligence (AI) has grown in sophistication over the past decade, governments around the world are asking how AI-enabled autonomy will affect national security. Technological innovation has played a major role in the evolution of warfare, from the advent of gunpowder to the invention of nuclear weapons, and more recently the development of stealth and precision-guided weapons. But AI has the potential for even more profound consequences, with US Secretary of Defense James Mattis wondering whether its adoption will change the fundamental nature of conflict.[1]

In the past several years, the United States, China, and Russia have each articulated plans to gain advantage through the application of AI to national security. The US Third Offset strategy, announced by Secretary of Defense Chuck Hagel in 2014, seeks to renew America's narrowing battlefield advantage by boosting innovation in several technologies, including AI-enabled robotics and system autonomy, in closer collaboration with innovative private sector enterprises. China released its AI strategic plan last summer, declaring its intentions to become the world's primary AI innovation center by 2030 through a strategy of military-civil fusion. And Russian President Vladimir Putin has gone so far as to suggest that the nation that leads in AI "will become the ruler of the world."[2]

AI will disrupt both the physical and virtual domains (and blur the boundary between them). In this paper, our focus is on the physical world. The past decade has produced startling advances in AI software, algorithms, data, and computing power. As these are fused with related commercial technologies, such as small, low-cost sensors, better communications, and continued digitization, new possibilities will arise for the configuration, connection, and control of current and new weapons and delivery platforms. With that, the glare of the spotlight on AI will grow very bright.

Several critical questions have already come up. What threats are posed by fully autonomous weapons, and can those threats still be mitigated by humans "in the loop"? Could AI level the playing field between Leviathan nation-states and smaller states and non-state actors? How will militaries operating in the context of different political, economic, and social systems adapt? What moral, ethical, and legal frameworks and governance mechanisms need to be updated or created for an age of fused man-machine conflict?

In the first part of this chapter, we highlight three implications for the business of AI-enabled warfare. First, the increasing physical separation of the human operator from weapon delivery can deliver near-term transformational economics. Second, the rise of "swarm" warfare has the potential to upend the effectiveness and strategic control points of warfare and further alter their economics. Third, we highlight one little discussed but substantial benefit of AI in military organizations: we estimate that over the next two decades, AI and its companion technologies can *cut in half* what MIT's Andrew McAfee and Erik Brynjolfsson call the 4Ds (dangerous, dirty, dull, and dear)[3] of today's defense workloads. It is less difficult, and more palatable, for a military to rally around reducing the 4Ds than to resolve debates about ceding control to Skynet and the Terminators.

Like all game-changing technologies, there is no free lunch with AI. In the second part of the chapter, we touch on two aspects of paying for the AI lunch: the profound and potentially wrenching changes required to the "operating system" for R&D, military-industrial-political relations, procurement, and talent, and how responsible actors can embed and enforce accountability for "*humanity* in the loop," even as the rapidly increasing speed, mass, diversity, and detachedness of conflict potentially obsolete "*humans* in the loop."

How AI Might Disrupt the Business of Conflict

Distancing the Human

AI-enabled autonomy is already finding its way into many tactical applications ranging from surveillance image processing to suppression of enemy air defenses. Greg Allen and Taniel Chan from the Belfer Center for Science and International Affairs at the Harvard Kennedy School suggest that "in the short term, advances in AI will likely allow more autonomous robotic support to warfighters, and accelerate the shift from manned to unmanned combat missions."[4] Deputy Defense Secretary Bob Work captured the sentiment well: "I'm telling you right now, ten years from now if the first person through a breach isn't a . . . robot, shame on us."[5]

"Narrow AI" (capable of performing a specific task in a particular domain) is likely to find its way into defense missions, taking over many "dangerous, dirty, dull, and dear" tasks currently performed by uniformed personnel. Autonomy has the potential to transform many of the military platforms we rely on today, from tactical aircraft to armored fighting vehicles to submarines. Many attributes of these platforms are determined by their need to carry, and protect, human occupants.

Take tactical aircraft as one example. Several cost and weight penalties are associated with a human pilot, including constrained forebodies, large canopies, displays, and environmental control systems. The aircraft's maneuver capabilities are limited by the pilot's physiological limits, such as G tolerance. Separating the pilot from the vehicle imparts benefits in terms of speed, maneuverability, and endurance and eliminates man-rating requirements, pilot systems, and interfaces. These design freedoms can be exploited to produce a smaller, simpler, faster, and nimbler unmanned combat aerial vehicle (UCAV) about half the size, a third to a fourth of the weight, and a third of the cost of a manned aircraft. The difference for surface ships is even more profound. The US Navy's new Sea Hunter, an autonomous unmanned surface vehicle (USV), has operating costs that are a small fraction of the cost of operating a destroyer, $15,000-$20,000 per day compared to $700,000. That is transformational.

The Rise of Swarm Warfare

The introduction of AI-enabled autonomy on the battlefield has the potential to bring profound changes to both the future character of conflict and to the economics of defense. Those economics are already in trouble. Since World War II, the US has procured an ever-smaller number of increasingly more powerful, accurate, and lethal weapon systems. But this has come at great cost, prompting General James Cartwright, former vice chairman of the Joint Chiefs, to bemoan in 2008 that the military must end its quest for "exquisite" weapon systems that are too costly, take years to design and build, and don't reach troops fast enough, or in quantities large enough, to address ever-changing threats.[6]

Strategists and policy makers would do well to pay heed to General Cartwright, as the convergence of AI and automation is accelerating the decrease in weapons' cost while increasing their capabilities. This shifts the cost/effectiveness calculus and control points away from fewer, individually complex lethal end-points in favor of multitudes of individually simple but collectively complex swarms of lethal end-points.

Consequently, as weapons become smaller and cheaper, they inevitably become more numerous as well, returning mass to the battlefield. This will in turn facilitate new swarm tactics, where large groups of small, autonomous systems coordinate to attack an adversary. "Exquisite" platforms, and for that matter large formations, could become more vulnerable to swarm attacks. Consider that a new Arleigh Burke-class guided missile destroyer costs $1.8 billion. The Sea Hunter costs only $20 million and clearly has nowhere near the capability of an Arleigh Burke, but it could be developed to act in an anti-surface ship capacity. And you can buy ninety Sea Hunters for the cost of one Arleigh Burke. Even with the additional cost of arming it, the comparison is stark—one needs to ask, could a destroyer defend itself from nearly 100 Sea Hunters? Even the Sea Hunter begins to look expensive once commercial and consumer technologies come into military use. Recall the 2000 attack on the USS Cole (itself an Arleigh Burke-class destroyer). She was successfully attacked when a small fiberglass boat carrying C-4 explosives and two suicide bombers approached the destroyer and exploded. A swarm of such vessels, guided autonomously, costing far less than $20 million each, would be difficult to stop.

In another, clearly extreme case, a US ally recently used a $3 million missile to shoot down a $200 commercial drone. This example illustrates in stark terms the challenge that militaries face in attempting to counter cheap, readily available technology with extremely expensive, high-end custom hardware. As much more affordable consumer products adapted to perform military tasks become available, expensive and high-capability platforms could become vulnerable to swarms, and swarm tactics could become the norm.

Vice Chairman of the Joint Chiefs of Staff Paul Selva succinctly articulates this economic challenge: "the US is on the absolute wrong end of the cost-imposition curve. We are doing a $10 solution for a 10-cent problem. We need a 10-cent solution for a $10 problem."[7]

Cutting the Cost of 4Ds in Half

Swarms are not here yet, and humans are still needed at the strategic, operational, and even tactical levels of warfare, yet the impact of automation on the defense workforce is likely to be profound. Although it might seem like defense is inherently different from the commercial sector, over two-thirds of the cost and operational structure of the Department of Defense (DoD) is driven by critical but relatively mundane processes that help forces get and stay ready for conflict. For example, a 2014 Defense Business Board report suggested that DoD spends about $125 billion annually on several "back-office" functions analogous to those in industry, such as HR, IT, finance, legal, procurement and supply chain, health benefits, and so on.

A recent study by the McKinsey Global Institute found that roughly 45 percent of tasks in the US economy could be automated in the next decade using existing technology,[8] and our analysis shows something similar for DoD. Of the approximately 1.3 million uniformed personnel in the US military, only 14 percent have a combat specialty. Although other deployed forces also operate in circumstances that require real-time judgement, and hence wouldn't be candidates for automation, that still leaves a great many "4D" uniformed jobs, particularly those performed at home, as candidates for automation. The same (with even less implementation risk) is true of many of the approximately 788,000 DoD civilians and is probably the case for most of the approximately 600,000 contractors as well. Our analysis of more than fifty DoD job codes, and their civilian equivalents, suggests that in aggregate nearly 40 percent of tasks performed by DoD personnel could be automated, and within some job codes the AI-enabled automation potential is nearly 75 percent (Exhibit 1). This automation will likely occur over decades, but given that the US military recruits uniformed staff for at least twenty-year careers, planning needs to begin today for the jobs the military hopes to automate via AI and for how training programs need to change.

While all the angst about Skynet/Terminator/HAL-like dystopian AI futures is completely appropriate, it is striking that the clearly beneficial, remarkably mundane potential of AI-enabled automation to transform the military back-office is barely mentioned. With DoD being "eaten alive" by military pay and health care costs, such automation could double the fiscal and organizational flexibility available to redeploy on next-gen systems over the next fifteen to twenty years.

Exhibit 1. AI-enabled automation might reshape the structure of the DoD's enlisted forces

Enlisted occupation group	Number of active duty personnel	Automation potential, %	Potential impact
Machine operator and production	23,293	74	17,305
Vehicle and machinery mechanic	160,690	62	99,140
Support service	30,171	61	18,542
Transportation and material handling	187,544	59	111,551
Other services	39,264	49	19,227
Construction	30,322	47	14,289
Protective service	82,178	41	33,494
Administrative	51,283	39	31,046
Media and public affairs	20,424	36	7,264
Electronic and electrical equipment repair	129,209	35	44,947
Engineering, science, and technical	214,668	35	74,676
Healthcare	95,332	33	31,488
Human resource development	38,023	25	9,661
Combat specialty	187,244	N/A	-
Total	1,318,875		512,630

SOURCE: BLS.org, Office of Personnel Management, McKinsey Global Institute

No Free AI Lunch

Wrenching Change for Military Operating Systems

Harnessing the full potential of the three areas highlighted above, and guarding against the risks, will require profound changes to current DoD operating models. To keep pace with the rapid development of new technologies, the defense procurement system needs to become drastically more agile than it is today. It currently takes seven years to complete the development phase of major defense acquisition programs. With new technologies and applications emerging regularly from commercial sources—already, McKinsey Global Institute has found more than 400 commercial use cases[9]—this is no longer fast enough to stay ahead of our adversaries. Moore's law may be waning, but it has shown us that even eighteen months is a lifetime in the commercial world. According to Andrew Ilachinski from CNA, there is a growing mismatch—even dissonance—between the accelerating pace of technology innovation in commercial and academic research communities and the timescales and assumptions underlying DoD's existing acquisition process.[10]

The center of gravity for research and development for the most relevant new technologies is squarely in the private sector versus government or traditional defense contractors. Commercial tech giants in the US and China are spending 10 to 20 percent of their revenue developing new technology versus just 1 to 3 percent true "at-risk" R&D by traditional defense primes. AI leaders Google, Facebook, Intel, Microsoft, and IBMs' aggregate spending on internal R&D as a percentage of revenue is more than five times that of the top five US defense contractors (Exhibit 2).

Exhibit 2. The US defense industry invests five times less in R&D than the AI leaders

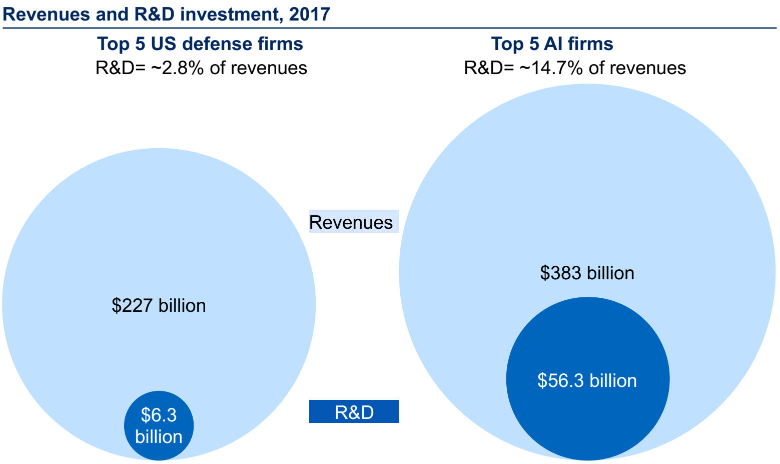

SOURCE: CapIQ, McKinsey analysis

And as with many early-stage technologies, there are numerous start-ups. Our analysis shows that the US has led the way, though lately China has gained ground. The US share of start-up investment declined from 71 percent in 2016 to 50 percent in 2017, and China's share rose from 13 to 30 percent (Exhibit 3).

Unlike the First Offset strategy, in which government-funded development of nuclear technology found its way into civilian applications, or the Second Offset, in which advanced technologies such as stealth and precision-guided weapons were primarily developed by DoD-sponsored research and development programs, success now critically depends on how well the government can embrace and integrate

Exhibit 3. Most investment in AI goes to companies headquartered in the US, followed by China

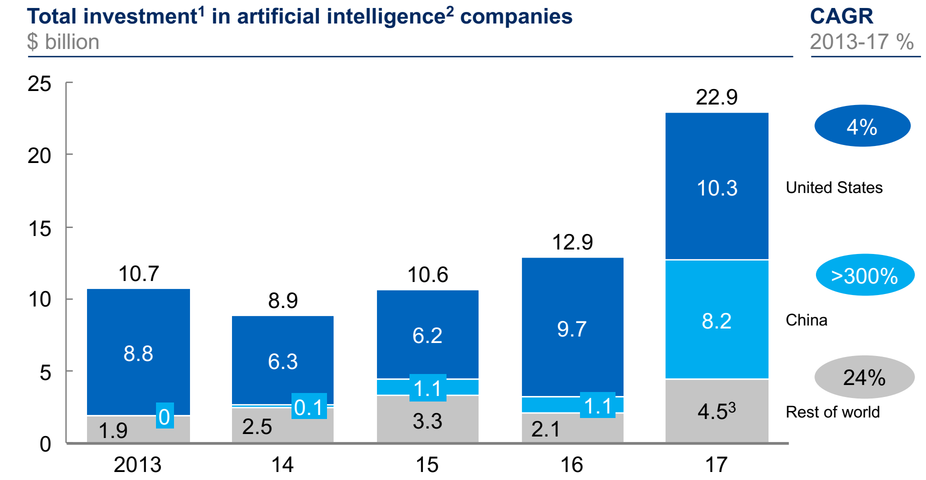

[1] Includes Venture Capital, Private Equity, Corp/Strategic M&A, IPO/Liquidity, Acquisition Financing, Asset Acquisition, Asset Divestiture, Corporate Divestiture, Leveraged recapitalization and Secondary Transactions (Open market and Private)
[2] Companies developing technologies that enable computers to autonomously learn, deduce and act, through utilization of large data sets. Applications for Artificial Intelligence & Machine Learning include speech recognition, computer vision, robotic control and accelerating processes, gene-sequencing in life sciences, etc
[3] Excludes Israel's increase in investment caused by Mobileye's M&A for $15.3 Billion

SOURCE: Pitchbook, SILA, McKinsey analysis

commercial technologies. The Pentagon will need to interface with the private sector to (a) adapt private-sector AI innovations to defense applications and (b) turn commercial AI powerhouses into partners for developing AI tools for national security.

In this context, the Third Offset is a direct challenge to the business model the Pentagon has practiced since World War II. Unlike the First and Second Offset, where the Pentagon's business model consisted primarily of cost-plus research and development contracts with companies that derived a substantial part of their enterprise value from government work, this will require dealing with primarily commercial enterprises. That presents problems for DoD, which must persuade unwilling partners to collaborate. Writing in the *New York Times*, Chris Kirchhoff, a former member of the Pentagon's Defense Innovation Unit (DIU), the West Coast outpost established in 2015, points out that critics in the Valley have two objections: (1) that collaboration risks compromising Silicon Valley values and (2) that contact with the military's inefficient processes would contaminate the fast-moving culture of Silicon Valley.[11]

The Pentagon is starting from a difficult position. According to a Center for New American Security study, 80 percent of top Silicon Valley executives surveyed rated the relationship between the Valley and DoD as "poor" or "very poor."[12] The most visible manifestation of this cultural dissonance came earlier this year when Google, in response to employee objections over the company's support of Project Maven (a relatively small $9 million contract with DoD to use AI to analyze drone footage), first declared its intention to end its support to that particular contract, then formally announced that it will not allow its algorithms to be used in weapons systems.

The fact is, all these technologies—search, storage, machine vision, natural language processing, autonomy—like the internal combustion engine and integrated circuit before them, have both civilian and military applications. Controlling their spread, and military use, will be challenging. The US will be at a real disadvantage if DoD can't access the best and latest technologies. This will require the department to transparently address both moral and practical concerns, as we discuss below.

Even with these concerns, there is still a massive opportunity for DoD to collaborate with the commercial innovators. As we indicated, reducing the 4Ds at scale and with intensity should be a palatable and even compelling message to innovators. Who wouldn't want to make our soldiers safer, basic processes faster and cheaper, and military work more meaningful and less routine? Working on common, less controversial applications of AI together first will at least deepen and broaden engagement and create pathways to solve much harder problems jointly.

Practically, it will be necessary to fundamentally reform DoD's systems, processes, talent strategy, mindset, and culture to create the agility[13] needed to take full advantage of the AI opportunity. Mike Griffin, undersecretary of defense for research and engineering, has highlighted the need to reform the Pentagon's slow-moving bureaucracy, saying "we can either retain our national [military] preeminence, or we can retain our processes, but we cannot have both."[14]

Ensuring "Humanity in the Loop"

The prospect that AI might be used to create fully autonomous weapons has justifiably raised alarms. More than 100 of the world's leading robotics and AI pioneers have argued that we must prevent AI's use in military applications, fearing that once this Pandora's box is opened, it will be hard to close. Elon Musk, who is no technophobe, has dramatically argued that in AI we are "summoning the demon."[15] The tension between the military utility of autonomous weapons and the need to manage the legal, moral, and ethical challenges is referred to inside the Pentagon as the Terminator Conundrum. Specifically, the central question is "how much autonomy is too much?"

Many acknowledge the risks of autonomy but see a panacea in human control—a "human in (or on) the loop." Current US policy states: "Autonomous … weapons systems shall be designed to allow commanders and operators to exercise appropriate levels of human judgment over the use of force."[16] We laud and support this intent. It is also likely this will come under intense pressure. First, autonomy is already a feature of *defensive* systems, where speed of response is critical. Defensive systems that autonomously identify and attack incoming missiles, rockets, ships, artillery, and aircraft are becoming entrenched. One of the first, the Phalanx CIWS (pronounced "sea-whiz"), began testing in 1973 and entered service in 1980. During the Cold War, both the United States and the Soviet Union developed autonomous launch systems for their nuclear weapons in case of a successful nuclear "decapitation" strike on the country's leadership. Fully autonomous systems can now be found for tank defense, missile defense, and even sentry guns protecting national borders such as those in Korea's Demilitarized Zone (DMZ). Autonomous offensive systems, while less prevalent, include the Harpy drone, which detects enemy radar emitters and destroys them.

If notions such as swarms and hypersonic weapons become prevalent, the pace and complexity of battle will overwhelm human cognition and response times, putting militaries that require real-time human intervention at a fundamental disadvantage. Further, ongoing human control of (semi) autonomous systems requires robust, secure communications. In denied, contested, or degraded communications environments, systems will require some degree of autonomy to maintain their effectiveness.

Attempts to ban or restrict new military technologies are almost as numerous as the technologies themselves, but mostly unblemished by success. AI is likely to impart real battlefield advantage, and with few exceptions, technologies that impart advantage have proven hard to contain. History is instructive here. The ancient Greeks established some of the earliest, short-lived restrictions on the use of poisons. The crossbow in the Middle Ages enabled commoners with relatively little training to defeat even heavily armored knights, causing Pope Innocent II to ineffectually ban their use against Christians in 1139. Poison gas was prohibited in war by The Hague Conventions of 1899 and 1907 but saw widespread use in the First World War. The US and the UK agreed at the 1907 Hague Convention to ban the installation of weapons on aircraft, but both went on to research and test aircraft-borne weaponry and used aircraft in lethal roles in World War I. Similarly, British attempts to ban submarine attacks on merchant shipping were ineffectual in World War II. Nuclear weapons escalated into a perpetual two-horse race, and post-Cold War nuclear nonproliferation treaties have a mixed record.

AI is even harder to regulate than these technologies, because it is not one tangible and discrete "thing," but a way of making decisions across a continuous spectrum of increasingly non-human control, enabled by deeply embedded and mind-bogglingly complex software interacting with devices across multiple networks. Consequently, autonomy to varying degrees is here to stay, and its use will accelerate.

Graduated Autonomy: First Steps to Humanity in the Loop?

Given the practical and historical reasons why it may be difficult to keep autonomous weapons off the battlefield, three legal, moral, and ethical concerns must be addressed. The first is that these systems will fail to adhere morally to just war theory and its concept of "right conduct in war" (*jus in bello*) or legally to international humanitarian law (IHL) with its requirement for distinction and proportionality. Specifically, critics fear that systems will be unable to distinguish between combatants and non-combatants, leading to indiscriminate killing of civilians. The second is that we will somehow lose control of these systems, posing novel challenges for attribution and accountability, and in extremis, they will somehow turn on their creators. Third, there is a moral concern that the decision to take a human life is sacrosanct, and we should not outsource life and death decisions to machines. Paul Scharre from the Center for a New American Security captures the sentiment nicely: "Use of force decisions are particularly significant and must remain under human control."[17]

The question as to whether a machine can consistently and accurately distinguish between enemy, civilian, and friendly troops is an essential one. Will humans or machines be better at following the laws of armed conflict, effectively discriminating between combatants and non-combatants? Autonomous systems do, and probably always will, make mistakes (even with a human in the loop). They do precisely what they are programmed to do, and it is impossible for programmers to envision all possible scenarios. Those against autonomy would reference several incidents where autonomous systems erroneously targeted allied military and even commercial aircraft. The errors occurred both with a human in the loop and in fully autonomous mode. Officers overseeing these operations tended to defer to the machine logic programmed by the manufacturer, blurring accountability for the accidents.

One might even invert the argument about AI-controlled warfare and ask whether AI could reduce the dangers of humans in the loop? Alarmists sometimes conveniently gloss over three facts: (1) *human* decision-making is highly variable in the "fog of war"—for every hero with thoughtful, timely judgment, there is an ordinary soul who panics; (2) throughout history, humans (in the loop but devoid of humanity) have caused more misery than any AI algorithm; and (3) more automation, intelligence, and precision have begun to *reduce* unintended civilian and military deaths from conflict and have acted as stabilizing forces. Definitive figures are hard to come by, but even the most expansive estimates of civilian casualties from drone strikes show only one civilian killed per seventy-two combatant deaths in 2016, far lower than that observed for other forms of strikes.[18] Yet, for two identical bad outcomes, why is it somehow still more palatable to accept human error than machine error?

While we still can, we should articulate, encode, and embed the *best* of humanity, while protecting against the worst, into our AI companions. The US Navy has been attempting to do just that for years in using graduated autonomy in its surface warfare combat systems. Some of the more advanced systems are controlled by a set of algorithms with four settings: manual; two semi-autonomous settings (that automate parts of the firing solution but still require human authorization to fire); and autonomous (automatically firing against threats that meet defined parameters), with different fine-tunings possible against different threats (e.g., cruise missiles versus aircraft). Systems can be configured to handle non-saturation (when there is time to thoughtfully respond to each potential threat) and saturation (when there isn't), thus providing the captain a way to pre-delegate his or her decision making against certain threats.

Graduated autonomy goes a significant way toward a balanced approach addressing the legal, moral, and ethical issues with lethal autonomous weapons. Graduated settings keep a human in the loop as much as possible, and most retain a form of human control. The captain remains in command, satisfying the need for attribution and accountability. Fully automated settings are only used in "saturation" settings, when they are arguably warranted under just war theory.

• • •

Just as Isaac Asimov created his Three Laws of Robotics to dictate robot behavior in his fictional works, we can algorithmically embed our values into the systems we design and field to create "ethical autonomy."

> "A robot may not injure a human being or, through inaction, allow a human being to come to harm. A robot must obey orders given it by human beings except where such orders would conflict with the First Law. A robot must protect its own existence as long as such protection does not conflict with the First or Second Law."

The authors wish to thank Edward Barriball and Jonathan Leydon for their contributions to this paper.

John Dowdy is a Senior Partner in the London office of McKinsey & Company. He led the firm's global Defense and Security practice from 2010-2016. Mr. Dowdy has conducted more than 100 projects on defense and security issues in eight different countries over his 27 year McKinsey career. He has worked extensively in countries including the US, the UK, Denmark, Australia, Japan and Canada. Over the past decade, he has been involved in projects improving the efficiency and effectiveness of air, land and maritime forces, headquarters organization, defense supply chain and logistics processes, and counter terrorism, among others. Most recently, he led McKinsey's support to the Equipment Support Plan (ESP) Review in the UK, which delivered £2.5 billion in audited savings. Mr. Dowdy leads McKinsey's research on best practices in defense. He led McKinsey's benchmarking on the efficiency and effectiveness of 33 defense forces around the world. He is a fellow at the Royal United Services Institute (RUSI), where he serves as a member of the Board of Trustees. Mr. Dowdy holds an MBA with high distinction from Harvard Business School, where he graduated as a George F. Baker scholar, and a B.S. in Electrical Engineering and Computer Science with honors from the University of California at Berkeley. He is a private pilot.

Chandru Krishnamurthy is a Senior Partner at McKinsey & Company. He leads McKinsey's US Defense Sector. In his 25 years at McKinsey, Mr. Krishnamurthy has led McKinsey's client service to a broad range of private and public institutions in the telecommunications, technology, private equity, defense, and national security sectors. He founded and led McKinsey's Digital Practice in the Southern US for over eight years, and has led its North American Cybersecurity Practice. In the most recent 5 years, Mr. Krishnamurthy has focused most heavily on defense and national security, on topics such as technology-driven productivity, innovation and organizational transformation. Mr. Krishnamurthy has been following (and in his distant past, developing) advances in artificial intelligence for the last 20-25 years. His master's thesis was on the testability of complex learning systems. He programmed very primitive and slow neural networks in the late 1980s, when he worked for FMC Corp in Silicon Valley as a software engineer. Mr. Krishnamurthy has an MBA from The Wharton School, University of Pennsylvania (1993), an MS in Electrical Engineering from Vanderbilt University (1988), and a BTech in Electrical Engineering from the Indian Institute of Technology, Delhi (1985).

[1] Press gaggle by Secretary of Defense James Mattis en route to Washington, DC, February 17, 2018, www.defense.gov.

[2] Russian President Vladimir Putin speaking to a group of Russian students on Knowledge Day, September 1, 2017, www.rt.com.

[3] Erik Brynjolfsson and Andrew McAfee, *Machine, Platform, Crowd: Harnessing Our Digital Future* (W.W. Norton & Company, 2017).

[4] Gregory C. Allen and Taniel Chan, *Artificial Intelligence and National Security*, Belfer Center for Science and International Affairs, July 2017, belfercenter.org/sites/default/files/files/publication/AI%20NatSec%20-%20final.pdf.

[5] Deputy Secretary of Defense Bob Work, "Reagan Defense Forum: The Third Offset Strategy," Reagan Presidential Library, Simi Valley, CA, November 7, 2015, dod.defense.gov/News/Speeches/Speech-View/Article/628246/reagan-defense-forum-the-third-offset-strategy.

[6] General James Cartwright, vice chairman of the Joint Chiefs, to the Military Officers Association of America in Arlington, VA, November 17, 2017, www. stripes.com.

[7] Air Force General Paul J. Selva, at the Brookings Institution event hosted by the Center for 21st Century Security and Intelligence, January 21, 2016, www.brookings.edu.

[8] James Manyika, Michael Chui, Mehdi Miremadi, Jacques Bughin, Katy George, Paul Willmott, and Martin Dewhurst, *A Future that Works: Automation, Employment and Productivity*, McKinsey Global Institute, January 2017, www.mckinsey.com/mgi/overview/2017-in-review/automation-and-the-future-of-work/a-future-that-works-automation-employment-and-productivity.

[9] Michael Chui, James Manyika, Mehdi Miremadi, Nicolaus Henke, Rita Chung, Pieter Nel, and Sankalp Malhotra, *Notes from the AI Frontier: Insights from Hundreds of Use Cases*, McKinsey Global Institute, April 2018, www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-applications-and-value-of-deep-learning.

[10] Andrew Ilachinski, *AI, Robots, and Swarms: Issues, Questions and Recommended Studies*, CNA, January 2017, www.cna.org/cna_files/pdf/DRM-2017-U-014796-Final.pdf.

[11] Christopher M. Kirchhoff, "Why Silicon Valley Must Go to War," *New York Times*, May 2, 2018, www.nytimes.com.

[12] Loren DeJonge Schulman, Alexandra Sander, and Madeline Christian, "The Rocky Relationship Between Washington and Silicon Valley," CNAS, July 19, 2017, www.cnas.org/publications/commentary/the-rocky-relationship-between-washington-silicon-valley.

[13] John Dowdy and Kirk Rieckhoff, "Organizational Agility in Defense and National Security: Elephants Learning to Dance," in *America's National Security Architecture: Rebuilding The Rocky Relationship Between Washington and Silicon Valley the Foundation,* eds. Nicholas Burns and Jonathon Price (Washington, DC: Aspen Strategy Group, 2016).

[14] "Regaining the Strategic Advantage in an Age of Great Power Competition: A Conversation with Michael Griffin," Hudson Institute, April 13, 2018, www.hudson.org.

[15] Greg Kumparak, "Elon Musk Compares Building Artificial Intelligence to Summoning the Demon," *Techcrunch*, October 27, 2014, techcrunch.com/2014/10/26/elon-musk-compares-building-artificial-intelligence-to-summoning-the-demon.

[16] US Delegation Statement on "Appropriate Levels of Human Judgment," Geneva, Switzerland, April 12, 2016, geneva.usmission.gov.

[17] Paul Scharre, "Autonomy, 'Killer Robots,' and Human Control in the Use of Force—Part I," Just Security, July 9, 2014, www.justsecurity.org/12708/autonomy-killer-robots-human-control-force-part.

[18] Jessica Purkiss and Abigail Fielding-Smith, "White House Releases Annual Counterterrorism CIVCAS Figures," The Bureau of Investigative Journalism, January 20, 2017, www.thebureauinvestigates.com/stories/2017-01-20/white-house-releases-annual-counterterrorism-civcas-figures.

"As the Information Age advances, the United States needs to recognize that data is a precious commodity that warrants a much higher national security priority."

–ERIC ROSENBACH & KATHERINE MANSTED

Can Democracy Survive the Information Age?

Eric Rosenbach
Co-Director
Belfer Center for Science and International Affairs
Harvard Kennedy School

Katherine Mansted
Fellow
Belfer Center for Science and International Affairs
Harvard Kennedy School

States have always used a combination of diplomatic, military, economic, and informational measures to advance their national interests, and technological change has altered each of these levers of power. The Information Revolution, however, has most radically reinvented the way in which states wield information power, ushering in changes to the nature of state competition, conflict, and international relations in the twenty-first century. Today's digitally enabled information operations bear no resemblance to the mass letter drops and radio broadcasts of the Cold War. Even the counter-messaging campaigns focused on al-Qaeda during the decade following 9/11 now seem anachronistic: the United States had little control over who tuned in to its messages, struggled to segment the audience or create customized content, and relied on expensive, yet imprecise, infrastructure.

Democracy is built on the crucial compact that citizens will have access to reliable information and can use that information to participate in government, civic, and corporate decision making. The technologies of the Information Age were largely built on the assumption that they would strengthen this compact. However, as typified by Russia's ongoing use of information operations against the United States and Europe, key information technologies have evolved quickly over the past five years and have been weaponized against democracies. The trajectory of data-driven technologies, including machine learning and other aspects of artificial intelligence, will increase the scale, complexity, and effectiveness of adversary information operations. As technology advances, and as geopolitical and ideological tensions between democratic and authoritarian states rise, information operations are likely to become more numerous, insidious, and difficult to detect. Democracy is resilient: few, if any, democracies will crumble under the coming wave of information warfare. But absent a new national security paradigm and real action, the weaponization of information technologies threatens to jeopardize democracies' ability to govern and protect their national security and to undermine people's trust in democracy as a system of government.

This paper explores both the politics and technologies that are changing the face of information power in the twenty-first century. In Part 1, we explain why states, especially those with authoritarian forms of government, are increasingly seeking to "game" democracy's strengths by using information operations.[1] We also highlight that authoritarian governments—and China in particular—have largely succeeded in controlling their domestic information environments. This affords them a degree of impunity to engage in information operations but in the long run may make them more brittle in the face of conflict and dissent in the Information Age. As an antidote to the national security community's tendency to "fight the last war" (in this case Russian information operations), we then explain why China is also capable and increasingly likely to engage in information operations to advance its national goals. In Part 2, we assert that advances in artificial intelligence, coupled with the growing abundance and importance of data, will turbocharge the scale and effectiveness of adversary information operations. Information operations to date are only prototypes of the sophisticated platforms and messages that non-democratic states will weaponize in coming years. This points to the urgency of reorienting America's national security strategy to focus on the emerging information operations threat.

In Part 3, we set out steps that the United States can take to build a whole-of-nation strategy to defend itself in the Information Age. By mobilizing action across civil society, the private sector, and government agencies, a whole-of-nation approach will play to the strengths of democracy. The linchpin of any strategy must be a clear national deterrence posture that explicitly includes the option of counterattacks that would threaten adversaries' control of their own information environments.

I. Democracy and Security in the Information Age

In the 1990s, as the internet started to be commercialized, it was widely assumed that technology would accelerate the global spread of democracy. The design of the internet itself—a decentralized network that empowers individuals to freely associate and share ideas and information—reflected liberal principles. As an American innovation, the internet, and the Information Revolution that followed, attested to the merits of democracy and appeared to cement the United States' position as a world superpower. During the 1990s and early 2000s, governments that sought to co-opt the internet to serve the state were seen to be resisting an immutable, democratizing force. In March 2000, President Bill Clinton expressed the dominant view in the United States when he derided China's nascent attempts to censor the internet as "like trying to nail Jell-O to the wall."[2] However, contrary to predictions of the futility of control in the Information Age, China ultimately developed the most sophisticated national-level system of censorship in history: the Great Firewall. Other non-democratic states have also invested significant resources in protecting and controlling their domestic "information environments." About eight years ago, authoritarian governments began to reap rewards from their investment in information control, and in mastering the tools of the Information Age. Actions by Russia, China, and terrorist networks like the Islamic State upended the conventional wisdom of the preceding two decades and demonstrated that information technologies can be used to exploit the vulnerabilities of democracy to advance nefarious interests.

The Vulnerabilities of Democracy

In the same way that island building in the South China Sea or "little green men" in Ukraine enable our adversaries to achieve foreign policy objectives without triggering thresholds that would invite a significant US response, information operations permit less powerful states to challenge core US national interests. Our adversaries are emboldened because they see technologically advanced democracies like the United States as digital "glass houses" with four specific vulnerabilities: (i) weak mechanisms for distinguishing facts from fiction; (ii) the long, media-driven nature of elections; (iii) the tech sector's profit-oriented culture; and (iv) the inability of the government to oversee and coordinate issues related to the information environment.

First, free speech is a core value of democracy. However, with the advent of social media platforms, the internet is no longer just a static bulletin board, but a place where any individual (or bot) can participate in the public debate, in real time. The nexus between the internet and social media means that, without resorting to diplomacy or conflict, adversaries can change a democracy's behavior by influencing its citizens at scale and in real time. The institutions democracies previously relied upon to provide objective facts have not adapted to the reality of the Information Age. Now, the information that citizens use to inform how they vote, protest, and debate in the public square is distributed via a largely unmediated social media environment and frenetic 24/7 news cycle. By contrast, authoritarian states often control the media, censor the internet, and in many cases shield their nations from outside information through national firewalls.

Second, elections, the heart and soul of a democracy, are vulnerable to both information operations and cyberattacks. The internet, social media, and data analytics will be the center of gravity for future political campaigns. All three played a key role in the Obama campaigns' ability to mobilize grassroots support,[3] while the Trump campaign established a new paradigm for presidential campaigns, using social media and big data analysis to both drive media coverage and to mobilize probable voters. The length of campaigns in the United States provides adversaries with a luxurious amount of time to plan, execute, and adapt information operations. Conversely, authoritarian states restrict the ways in which the public square can influence national decision making. Most obviously, their leadership does not hinge on genuine elections that can be disrupted by manipulating public opinion.

Third, healthy democracies rely on the private sector to drive economic growth and prosperity in a way that is compatible with the overall public good. Over the past decade, profit-focused technology firms, especially Facebook and Twitter, have amassed an enormous amount of valuable data and honed the capability to drive citizens' decisions and opinions. This power has not been matched by a sense of public purpose, much less a sense of responsibility to contribute to national security. Although there are some early indications of change, these companies' indifference remains a vulnerability. Conversely, authoritarian states closely align their industrial policy for the tech sector and national security priorities.[4] A state like China has unfettered access to domestic communications, has rules against online anonymity, and can instantly order shutdowns of websites or designated accounts.

Fourth, in democracies, the executive and legislative branches perform the "inherently governmental functions" of national security and regulation. However, in the United States, government has not kept pace with adversaries' strategies of exploiting information technologies. Moreover, information operations (by design) fall into the seams between the public and private sectors. For example, ahead of the 2016 presidential election, there was ample evidence of Russia's intent and capability to meddle in US politics. The intelligence community knew that Russia had been publicly signaling an increasingly aggressive political warfare posture from at least 2013[5] and had tested many of the information tools it used in 2016 in Ukraine's presidential election in 2014. As early as 2012, researchers had also reported the Syrian government's use of similar tools during the Syrian Civil War.[6] Many campaign staffers and technologists knew that in 2012 and 2015, respectively, Google and Facebook had promoted their platforms' abilities to influence voter behavior.[7] We also know now that Russian operatives stationed in the United States began laying the groundwork for the 2016 campaign as early as 2014 and purchased advertisements on Facebook and Google.[8] However, researchers, government officials, and tech platforms failed to work together in a way that let them connect the dots and anticipate—or even detect the full extent of—Russia's actions.

The Authoritarian Information Paradox

Authoritarian states have recently deployed information operations to advance their foreign policy, but propaganda and censorship have always been essential tools for maintaining control and power at home. For authoritarian governments in the Information Age, however, the internet and related technologies are also a major vector for instability—since they make news and ideas accessible and allow people to mobilize in ways that can threaten the ruling party.[9] Consider this stark observation in China's 2017 Cybersecurity Strategy: "If our party cannot traverse the hurdle represented by the Internet, it cannot traverse the hurdle of remaining in power for the long term."[10] Authoritarian governments do not just fear that their citizens will use the internet to organize or rebel; they also believe that democracies use the internet to advance pro-democratic narratives to undermine their regimes. Russia's president has derided the internet as a "CIA project,"[11] while China's president characterized the competition between major powers as "rivalry for ideology, for the power of discourse."[12] It is easy to dismiss these statements as a diversionary tactic. However, while democratic governments, as a general rule, do not engage in information operations to undermine their competitors, commercial and civil society actors do actively promote democratic and liberal principles—indeed, they are the primary agents for much of the "soft power" appeal of the US system of government.[13] This dynamic means that authoritarian states do not just view control of their information environments as a domestic matter; they increasingly believe that offensive action might be required to counter what they perceive as foreign information incursions.

Centralized control of the internet does, however, make authoritarian states brittle. In democracies, a plethora of decentralized nongovernment actors play a role in disseminating trusted information and debunking propaganda. Even a small chink in the armor of authoritarian states' information control systems may have existential ramifications for those in power. This helps explain why authoritarian governments are prepared to engage in hostile information operations to defend their information environment, including to suppress liberal ideas and to discredit alternative systems of government. It also explains why most efforts for creating international "norms" against state information operations are bound to fail. Thus, in the long-run, authoritarian states, not democracies, may prove to be the real glass houses.

The China Example

Russia's ongoing interference in US democracy is the most politically salient instance of an information operation, but we should not assume that Russia is the only actor in this space. We should also not assume that all information operations will employ similar tools and tactics or share similar objectives. For example, researchers have observed that China is increasingly embracing "sharp power," which centers on using information for the purposes of distraction, manipulation, and intimidation.[14] China's domestic propaganda and censorship capabilities provide it with a powerful vehicle for information operations, particularly as the popularity of Chinese platforms like WeChat and internet penetration in countries with large Mandarin-speaking populations continue to rise. China's use of government operatives (colloquially known as "50-centers") to flood social media with an estimated 488 million posts annually to advance pro-government narratives and drown out negative stories is well-documented.[15] Additionally, China has been increasingly implicated in spreading "fake news" stories designed to foment civil unrest and distrust in democracy in Taiwan.[16] Even if China confined propagandistic content to its domestic platforms, the borderless nature of the internet means that this content will inevitably spill over to other countries.

China's cyber agencies are also increasingly using offensive operations to advance China's regional foreign policy interests—particularly in connection with China's territorial disputes. Chinese hackers blocked or vandalized Japanese websites in response to tensions over the disputed Senkaku (Diaoyu) Islands,[17] and they took down Philippines government websites[18] and hacked Vietnam Airlines systems[19] following the 2016 UN Permanent Court of Arbitration ruling on the South China Sea. Perhaps most concerning of all, China's cyber agencies also have a history of hacking government, political, and media networks in the lead-up to democratic elections in its region, most recently in Taiwan[20] and Cambodia.[21] At this stage, it is not clear whether this activity is limited to intelligence-gathering or could be in preparation for information operations.

Chinese information operations within the Asia-Pacific region threaten US allies and interests. It is also increasingly conceivable that China might use information operations to directly target the United States or other Western democracies. As China's power grows, it is increasingly using influence and intimidation tactics, for example, against media companies, civil society, and academia in Europe, the United States, Australia, and New Zealand, to suppress information contrary to the interests of the Chinese Communist Party.[22] To this point, these tactics are executed by human operatives rather than as part of information operations exploiting information technologies.

Currently, China predominantly uses its extensive cyber capabilities for government and commercial espionage; however, it could easily repurpose these capabilities, or information it has collected, for use in information operations against Western targets. Instead of using its access to a penetrated network to exfiltrate data, China could easily use that position to poison government datasets, seed false content into the information environment, or obtain and leak sensitive documents. Recalling that China was responsible for the 2015 data breach of the US Office of Personnel Management—resulting in the theft of sensitive information on four million people who had undergone US government security checks—it is apparent that China already controls large troves of sensitive information that could be strategically leaked or used for highly targeted information operations. China also has the technical tools to launch "blunt force" attacks. The Great Firewall can be repurposed into an offensive weapon (which researchers have dubbed the "Great Cannon") that can launch massive cyberattacks to shut down websites or to inject false or misleading content into targeted systems.[23] The potential of the Great Cannon, and China's willingness to use it against Western democracies, was demonstrated in 2015, when Chinese hackers launched a massive distributed denial-of-service attack against US-headquartered website GitHub.[24] GitHub, the world's biggest repository of open source code, had hosted content that provided technology to subvert Chinese online censorship. As the "authoritarian information paradox" discussed above predicts, China is willing to use offensive measures to suppress information that challenges its domestic control of information.

II. The Coming AI Wave

To this point, information operations have relied on human operatives to generate content and have used a combination of human "trolls" and basic automated algorithms to disseminate that content. Social media companies currently have the capacity (if not always the will) to defend against most of the tools that today's information adversaries employ. Most automated algorithms can be identified because they exhibit predictable bot-like patterns. Similarly, after public and political pressure, Facebook, Twitter, and Google have announced efforts to update their algorithms to deemphasize fake news. Facebook has also introduced labeling requirements for electoral advertisements and advertisements on topical political issues,[25] while Twitter has taken action to shut down bot networks and delete inauthentic accounts.[26] These are all positive steps, but advances in artificial intelligence (AI) technologies in coming years may allow adversaries to outpace our abilities to defend against them using technology alone.

The Data Explosion

The primary driver of advances in AI is the growing abundance of data. The amount of data in the world is growing at an exponential rate (around 90 percent of the data in the world today was created in the last two years). By 2020 there will be some 20 billion Internet of Things sensors embedded around the world—collecting data from wearable devices, home appliances, and city infrastructure. Parallel developments, including rising use of geotagging by phone apps, smart cars, and financial services firms; improvements in facial recognition technology; and the rise of affective computing (whereby machines can discern human emotions from text, facial expressions, and voice patterns[27]) will add to increasingly rich sets of data. The data explosion is overwhelmingly driven by economics. Data enables firms—whether they are online platforms, traditional brick and mortar retail chains, financial companies, or insurance brokers—to better target consumers.

Access to data also has a multiplying effect, due to advances in machine learning (a subset of AI). Firms with more data can serve their customers more effectively: for example, Netflix's personalization algorithms have made it more valuable than Disney.[28] Firms with more data are also able to create better AI software: for example, Facebook trained its algorithms to recognize faces by learning from billions of users tagging pictures of their friends. Today, nearly every industry is either using or exploring machine-learning applications. Erstwhile titans of the industrial age, like GE, market themselves based on their ability to aggregate and process customer data.[29] The rapid pace of data collection is likely to accelerate. By some estimations, up to 80 percent of the current share price of Facebook and up to 60 percent of that of Google is attributable to future growth—indicating that the market is confident in their ability to collect and monetize increasingly vast tranches of data.[30]

AI Will Give Adversaries a Technological Edge

The rise of data is underwriting a new wave in the development of AI technologies. These advances will benefit those who defend against information operations (for example, by helping to train algorithms to detect and filter propaganda). In the short term, however, they are likely to magnify the scale and effectiveness of adversary information operations. With China articulating a national plan to be the world leader in AI investment and research,[31] we can no longer assume that the benefits of Information Age advancements will accrue to American interests. Likewise, while the research and development costs of cutting-edge AI are high, breakthroughs are likely to spread quickly and widely, equipping both state and non-state adversaries with a technological edge.

Hyper-personalized content

Data is valuable commercial information, but in the hands of adversaries, it can be extraordinarily dangerous. Russia has a long history of exacerbating divides on fractious social issues by targeting susceptible identity groups: for example, its operatives posted divisive content to Facebook groups affiliated with Black Lives Matter supporters and detractors and to individuals who "liked" content related to race.[32] With advances in machine learning, adversaries will be able to build, buy, or steal detailed profiles on nearly every citizen. These profiles will not just be based off known data inputs (like whether a person is a member of a racial justice group); machine learning tools will also be able to predict with increasing accuracy people's personality and preferences, political and religious beliefs,[33] real-time emotions, and identity characteristics like sexual preference.[34] Social media micro-targeting is already one of the more difficult information operation tactics to counter—since messages are only seen by select individuals or groups and for a short time. As machines begin to know us better than we know ourselves, adversaries will increasingly be able to identify and target those who are most susceptible to influence. They will then be able to deliver highly personalized content that achieves maximum effectiveness by exploiting individuals' unique characteristics, beliefs, needs, and vulnerabilities.[35]

Taking humans out of the loop

Humans, such as China's 50-centers and the hundreds of Russian "troll farm" operators, are currently needed to develop content for information operations. Bots are often used to amplify the content but do not create it. However, the next wave of AI research—which is focused on creating tools that are better able to understand human language and to process it in the right context[36]—may put bots in the driver's seat. Today's AI tools can only interact with humans in highly circumscribed contexts, but they are constantly improving their ability to generate original and dynamic content, to persuade, and to tailor their content to appeal to their interlocutor's mood. As AI becomes better at mimicking human behavior, fake, automated online personas will become harder to detect. Additionally, once a certain piece of AI software exists "in the wild," the marginal cost of using it to create one additional bot or to influence one additional citizen will be almost zero.[37] Finally, AI tools are by nature learning systems. AI-enabled bots will be able to conduct experiments and learn from successes and failures in real time to recalibrate their methodologies for maximum impact. As a result, once an adversary sets an objective for an online information operation, it may become most effective and economical to take humans out of the loop at the tactical level.

Deepfakes

Advances in AI are also making digital manipulation of audio and video cheaper and harder to detect. Currently, high-quality audio or visual material is considered highly reliable evidence in factual and legal disputes and is largely taken at face value by the media and individuals to be legitimate.[38] However, according to expert predictions, researchers are only several years or even months away from being able to produce realistic video forgeries that will fool the human eye. Already, user-friendly software exists online to produce realistic fake audio, provided there is a sufficiently large training dataset of the particular voice.[39] Adversaries will soon be able to create entirely false audiovisual content or, even more insidiously, modify existing content to create highly effective information operations.[40] Some analysts have also pointed to the risk that increased use of digital forgeries by both legitimate and illegitimate actors could pose a more systemic risk to democracy by eroding people's trust in even completely truthful information.[41]

We are more sanguine about this risk. Just as the internet has evolved to require security certificates for trusted websites, it is likely that, over time, audiovisual certificate systems will become increasingly sophisticated. Blockchain technologies could also be used to ensure these certificates are authentic. In the short term, however, adversaries will be able to take advantage of the gap between the emergence of better forgery technology and new authentication norms. Even once that gap closes, forged audiovisual content is likely to remain a significant concern on particularly contentious issues such as international crises or fraught political issues—where news will spread fast, and quick decisions will need to be made.

III. A Whole-of-Nation Security Strategy for the Information Age

Our response to Vladimir Putin's ongoing attempts to undermine the strength of American democracy will be a defining issue of the Information Age. The most important lesson the United States should internalize from Russia's interference campaign in 2016 is the price of complacency: we ignore continued cyberattacks and information operations at our own peril. Adversaries are able to turn our democratic system against itself because we lack a coherent national security strategy for the Information Age. A whole-of-nation approach that recognizes and harnesses the expertise of academia, civil society, and companies, and across government agencies, will play to the strengths of democracy. At a minimum, this approach should contain the four actions we outline below.

Establish a clear deterrence posture against information operations and cyberattacks that begins with explicit declaratory policy and includes the threat of offensive information and cyber counterattacks.

Strategic signaling matters as much in the Information Age as it did to the nuclear threat during the Cold War. Failures by both the Obama and Trump administrations to confront Russia will lead our adversaries to continue to believe that the United States will hope to weather the blows of information operations and cyberattacks. This puts us in an invidious position: the scale, potency, and likely impact of future information operations will not be determined by how prepared we are to combat them; we are instead leaving it up to our adversaries to determine when their ends justify the means of deploying an information operation against us. As highlighted earlier in this paper, authoritarian governments recognize that they exist in a brittle information environment. This fact makes them particularly susceptible to the threat of information and cyber operations against key elements of their media, technology, and government ecosystems. With that in mind, the United States, led by the president and other senior political leaders in the executive and legislative branches, must explicitly state that information (and cyber) attacks against the United States and our allies will result in counteraction. The counterpunch should use all available levers in the foreign policy toolkit but should always include some aspect of an incisive information operation.[42] Looking over the horizon, it is particularly important that the United States signal to China that information operations against democracies will result in significant strategic risk to China's core national interests.

Recalibrate the intelligence community to provide the United States and its allies with the intelligence necessary to detect and expose sources and content related to information operations.

During the run-up to the 2016 presidential election, the US intelligence community provided impressive, clear, and unique intelligence about Putin's intent to interfere in the presidential election. The intelligence community provided relatively little intelligence, however, on some of the specific operations that the Russian intelligence services used to influence American voters. Understandably, the intelligence community simply was not postured to provide indications and warnings about manipulation of social media via botnets and false personas. The intelligence community cannot shift its full attention away from more existential threats to the country, but it must invest and innovate to bolster the early warning system of attribution for information operations that target US democratic institutions. This will require better collaboration and information sharing with Silicon Valley firms, which in turn need to shed a post-Snowden reluctance to cooperate with government on pressing national security issues. Additionally, the government should enlist the support of private sector threat intelligence firms, which have excellent capabilities and will be key to any information-sharing arrangement.

Attribution is not only a requirement for an effective deterrence strategy, but also a core aspect of the best antidote to information operations: public shaming and fact-based counter-messaging. The United States can learn from and follow the examples set by both the French and German governments, which used intelligence about probable Russian information operations to warn their citizens and publicly call out the Russians. Just hours after a massive online leak of emails, the 2017 Macron campaign issued a statement blaming the leaks on hackers intent on "sow[ing] doubt and disinformation"[43] about the French elections. After receiving intelligence from her national intelligence organization, Chancellor Merkel warned the German public of possible Russian interference in Germany's elections, a course of action that many experts believe influenced Russia not to leak data stolen in a 2015 hack of the Bundestag.[44]

Enact national legislation that requires social media platforms to increase transparency about their algorithms and political, bot-driven content.

Social media companies are an essential aspect of American economic power in the Information Age, but they have also created tools and systems that can be used to subvert democracy. Ensuring that social media is not gamed by our adversaries cannot be left to self-regulation but needs to be incentivized and in some cases mandated by government. In particular, citizens have a right to know when they are seeing paid political advertisements and, in some cases, why they are being targeted by certain political or social campaigns. While Facebook has recently introduced internal requirements for displaying disclosures on political advertisements, this issue is too important to be left to the discretion of individual companies. Congress should enact legislation that mandates at a minimum the same disclosure for political advertisements on social media as for traditional media.

Additionally, Congress should pass laws that require platforms to identify and label foreign actor-driven bots and to provide users with the option to block them. Researchers have already developed tools that can screen for and identify bots; social media companies should now provide their users with more protection from them.[45] As we noted earlier in this paper, bot-spotting is likely to become more difficult with advances in AI. Therefore, requiring social media firms to label accounts that they should be reasonably able to identify as bots will force these firms to keep ahead of the technological curve and to adopt the latest best practices on bot identification from researchers and the security community. Private sector–led bot-spotting also supports the whole-of-nation ethos. Bot labelling does not silence speech or censor content but equips citizens with tools to better understand the content they are engaging with and to judge its veracity for themselves.

Finally, social media companies must adjust their algorithms to reflect their role in democracy. Because of their market dominance, platforms like Twitter, Facebook, and YouTube do not just house public discourse; they shape it. Currently, social media algorithms are optimized for user engagement because clicks, views, and likes maximize profit. As a result, social media platforms often promote and prioritize controversial information—something that Russian trolls and bots have exploited to great effect. It is technically uncomplicated, and necessary for the overall public good, to adjust these algorithms.

Enact national legislation that establishes a high level of protection for citizens' private data.

As the Information Age advances, the United States needs to recognize that data is a precious commodity that warrants a much higher national security priority. Data protection has historically been viewed as a niche issue relevant only to consumer privacy, but recent incidents, such as the hacks of the Office of Personnel Management and credit reporting agency Equifax, illustrate that data is a high-priority target.

Looking over the horizon, adversaries will greatly increase operations to steal sensitive and valuable information from the private sector. Given the richness of data held by companies like Google, Amazon, Facebook, and the financial and health care sectors, it is highly likely that adversary intelligence services will expand their traditional targets to include corporate datasets that could be used to train AI systems and to hone information operations. We should not be surprised if Facebook eventually confirms that Russian intelligence services accessed and operationalized the same private user data acquired by Cambridge Analytica.[46]

National data protection legislation is necessary because even the serious repercussions that arose from the Equifax data breach, including costly litigation and reputational damage, have not sufficiently changed most corporate perspectives on data protection. While Europe's General Data Protection Regulation is by no means a perfect model and in some respects is inconsistent with other US values,[47] it has been effective at driving corporate investment in data protection. Data protection legislation passed in California in June 2018 will need fine-tuning before taking effect in 2020, but it establishes important principles that could serve as the foundation for national legislation.[48]

IV. Conclusion

Information technologies have not just revolutionized lives, societies, and economies; they are reshaping the nature of twenty-first century politics and conflict. In many respects, our adversaries have learned and adapted to this reality more quickly. As technology continues to advance, and as our adversaries learn from the successes of Russian information operations, democracies should brace themselves for increasingly sophisticated and aggressive information attacks. Leaders in democracies must also realize that they can no longer advance and defend their national interests through conventional military, economic, and diplomatic means. For America, there remains a narrow window in which a coherent whole-of-nation strategy can be devised to combat the threats of the Information Age.

The integrity, and legitimacy, of our system of government may depend on it.

Eric Rosenbach is Co-Director of the Harvard Kennedy School's Belfer Center for Science and International Affairs and a Kennedy School Public Policy lecturer. Mr. Rosenbach previously served as the Pentagon's Chief of Staff from 2015-17 and Assistant Secretary of Defense for Global Security, responsible for leading all aspects of the department's cyber activities and other key areas of defense policy. On Capitol Hill, he previously served as national security advisor for then-Senator Chuck Hagel. In the private sector, Mr. Rosenbach worked as the Chief Security Officer for a large European telecommunications firm. Mr. Rosenbach is a former Army intelligence officer and commander of a telecommunications intelligence unit. He has co-authored several books on national security. He was a Fulbright Scholar. He has a JD from Georgetown, an MPP from Harvard and is a proud graduate of Davidson College.

Katherine Mansted is a nonresident fellow at Harvard Kennedy School's Belfer Center for Science and International Affairs, with a focus on emerging technologies, cybersecurity, and the Asia-Pacific. Her research explores the impact of the information revolution on traditional systems—including democracy, national defense, and international relations. Her publications include work on cyber-enabled foreign interference in Australia and internet privacy. Previously, Ms. Mansted practiced law as a commercial solicitor, served as a policy adviser to an Australian Cabinet Minister, and served as an associate to a justice of the High Court of Australia. She holds a Master of Public Policy from the Harvard Kennedy School, and a Bachelor of Laws with First Class Honors and a Bachelor of International Relations from Bond University.

[1] We use the term "information operations" to refer broadly to the use of information to "influence, disrupt, corrupt, or usurp" decision making within a state, especially via the internet and related information technologies. While the phrase in quotation marks comes from US Joint Publication 3-13 (Information Operations), we are not using the current military definition of the term—which focuses on exploiting adversary information flows, and protecting our own, during extant military operations.

[2] William J. Clinton, "Remarks at the Paul H. Nitze School of Advanced International Studies," Washington, DC, March 8, 2000, www.presidency.ucsb.edu/ws/index.php?pid=87714.

[3] In 2013, Google published a case study, *Obama for America uses Google Analytics to Democratize Rapid, Data-driven Decision Making*, explaining how Google Analytics had been critical to Obama's 2012 "data-driven re-election campaign." Claiming that digital marketing and analytics were responsible for "providing much of the winning margin" for the campaign, the case study outlines how Google Analytics was used to understand voter motivations and to shape the information voters were served when they searched to verify claims made during debates or as they considered how to vote. See: analytics.googleblog.com/2013/08/obama-for-america-uses-google-analytics.html.

[4] This alignment between industrial and national security policy is most obviously reflected in the Chinese government's use of cyberattacks to steal US commercial secrets, an effort that continues despite a 2015 US-China bilateral agreement to cease industrial espionage.

[5] For a summary of these public signals, see: Linda Robinson et al., *Modern Political Warfare: Current Practices and Possible Responses* (Santa Monica, CA: RAND Corporation, 2018), 42–48, www.rand.org/pubs/research_reports/RR1772.html.

[6] Norah Abokhodair, Daisy Yoo, and David W. McDonald, "Dissecting a Social Botnet: Growth, Content and Influence in Twitter," in *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing* (New York, 2015), 849, doi.acm.org/10.1145/2675133.2675208.

[7] Ashley Parker, "Facebook Expands in Politics, and Campaigns Find Much to Like," *New York Times*, July 29, 2015, www.nytimes.com/2015/07/30/us/politics/facebook-expands-in-politics-and-campaigns-find-much-to-like.html.

[8] Robert S. Mueller, "United States of America v. Internet Research Agency & Ors. Indictment by the Grand Jury for the District Court of Columbia." (Case 1:18-cr-00032-DLF, February 16, 2018).

[9] For example, when President Xi Jinping came to power, a memo referred to as "Document No. 9" was allegedly distributed to senior party leaders that listed seven "perils" to the Chinese Communist Party's (CCP) leadership. These included "Western constitutional democracy"; promotion of "universal values" like human rights, media independence, and civic participation; and "nihilist" criticisms of the CCP's past. See: Chris Buckley, "China Takes Aim at Western Ideas," *New York Times*, August 19, 2013, www.nytimes.com/2013/08/20/world/asia/chinas-new-leadership-takes-hard-line-in-secret-memo.html.

[10] Translated by Elsa Kania, et al., "China's Strategic Thinking on Building Power in Cyberspace," New America, September 25, 2017, www.newamerica.org/cybersecurity-initiative/blog/chinas-strategic-thinking-building-power-cyberspace/.

[11] Ewen MacAskill, "Putin Calls Internet a 'CIA Project' Renewing Fears of Web Breakup," *The Guardian,* April 24, 2014, www.theguardian.com/world/2014/apr/24/vladimir-putin-web-breakup-internet-cia.

[12] Quoted in Fergus Ryan, "Money Talks in China's Cloistered Internet," *The Strategist* (blog), December 15, 2017, www.aspistrategist.org.au/money-talks-in-chinas-cloistered-internet/.

[13] Weatherhead Center for International Affairs, "Hard Times for Soft Power: A Q&A with Joseph Nye," Harvard University, May 30, 2017, epicenter.wcfia.harvard.edu/blog/joseph-nye-qa.

[14] Christopher Walker and Jessica Ludwig, "The Meaning of Sharp Power: How Authoritarian States Project Influence," *Foreign Affairs,* November 16, 2017, www.foreignaffairs.com/articles/china/2017-11-16/meaning-sharp-power.

[15] See, for example, David Wertime, "Meet the Chinese Trolls Pumping out 488 Million Fake Social Media Posts," *Foreign Policy,* May 19, 2016, foreignpolicy.com/2016/05/19/meet-the-chinese-internet-trolls-pumping-488-million-posts-harvard-stanford-ucsd-research/.

[16] See, for example, David Spencer, "Why the Risk of Chinese Cyber Attacks Could Affect Everyone in Taiwan," *Taiwan News*, July 13, 2018, www.taiwannews.com.tw/en/news/3481423.

[17] Bill Gertz, "US Officials Say China behind Cyber Attacks on Japan," *Washington Free Beacon* (blog), September 25, 2012, freebeacon.com/politics/cyber-blitz/.

[18] Janvic Mateo, "68 Gov't Websites Attacked," *The Philippine Star,* July 16, 2016, www.philstar.com/headlines/2016/07/16/1603250/68-govt-websites-attacked.

[19] Charlie Osborne, "Chinese Hackers Take down Vietnam Airport Systems," ZDNet, August 1, 2016, www.zdnet.com/article/chinese-hackers-take-down-vietnam-airport-systems/.

[20] Scott Morgan, "Taiwan Prepares for Spike in Chinese Cyber-attacks in Lead-up to Elections," *Taiwan News*, July 9, 2018, www.taiwannews.com.tw/en/news/3477568.

[21] FireEye, "Chinese Espionage Group TEMP.Periscope Targets Cambodia Ahead of July 20148 Elections and Reveals Broad Operations Globally," July 10, 2018, www.fireeye.com/blog/threat-research/2018/07/chinese-espionage-group-targets-cambodia-ahead-of-elections.html.

[22] Thorsten Benner, et al., *Authoritarian Advance: Responding to China's Growing Political Influence in Europe*, Global Public Policy Institute, February 5, 2018, www.gppi.net/publications/rising-powers/article/authoritarian-advance-responding-to-chinas-growing-political-influence-in-europe/.

[23] Bill Marczak, et al., "China's Great Cannon," The Citizen Lab, April 10, 2015, citizenlab.ca/2015/04/chinas-great-cannon/.

[24] Lorenzo Franceschi-Bicchierai, "China Is Behind DDoS Attack on GitHub, Activists Say," Motherboard, March 30, 2015, motherboard.vice.com/en_us/article/8qx7wz/china-is-behind-ddos-attack-on-github-activists-say.

[25] Facebook, "Making Ads and Pages More Transparent," April 6, 2018, newsroom.fb.com/news/2018/04/transparent-ads-and-pages/.

[26] See, for example, Twitter, "Confidence in Follower Counts," July 11, 2018, blog.twitter.com/official/en_us/topics/company/2018/Confidence-in-Follower-Counts.html.

[27] See, for example, the Affective Computing research group at the MIT Media Lab: www.media.mit.edu/groups/affective-computing/overview/.

[28] "Netflix is Moving Television Beyond Time-slots and National Markets," *The Economist,* June 28, 2018.

[29] For example, GE's "Predix Platform" promises to extract value from the "massive amounts of data" generated by customers' industrial operations. See: www.ge.com/digital/predix-platform-foundation-digital-industrial-applications.

[30] Rod Sims, "Don't Rely on Amateur Journalists," *The Mandarin*, July 4, 2018, www.themandarin.com.au/95208-rod-sims-dont-rely-on-amateur-journalists/.

[31] Cade Metz, "As China Marches Forward on AI, the White House is Silent," *New York Times,* February 12, 2018, www.nytimes.com/2018/02/12/technology/china-trump-artificial-intelligencre.html.

[32] Russian operatives also attempted to stoke divides on hot-button issues, including the Black Lives Matter movement, gun control, Islamophobia, immigration, and police violence.

[33] For example, Google offered political interest targeting to advertisers in the 2016 US election cycle, based on whether users had been identified as "left-leaning" or "right-leaning." See: Google, "Security and Disinformation in the US 2016: What We Found," October 30, 2017, storage.googleapis.com/gweb-uniblog-publish-prod/documents/google_US2016election_findings_1_zm64A1G.pdf.

[34] For example, researchers at Stanford have created an AI system able to predict a person's sexual orientation from a photograph with up to 91 percent accuracy.

[35] Matt Chessen, *The MADCOM Future*, Atlantic Council, September 26, 2017, 10, www.atlanticcouncil.org/publications/reports/the-madcom-future.

[36] Venkat Srinivasan, "Context, Language, and Reasoning in AI: Three Key Challenges," *MIT Technology Review,* October 14, 2016, www.technologyreview.com/s/602658/context-language-and-reasoning-in-ai-three-key-challenges/.

[37] Matt Chessen, The MADCOM Future, Atlantic Council, September 26, 2017, 4, www.atlanticcouncil.org/publications/reports/the-madcom-future.

[38] Future of Humanity Institute, et al., *The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation*, February 2018, 46, maliciousaireport.com/.

[39] See, for example, the Canadian company Lyrebird: lyrebird.ai/.

[40] Matt Chessen, *The MADCOM Future*, Atlantic Council, September 26, 2017, 10, www.atlanticcouncil.org/publications/reports/the-madcom-future.

[41] Robert Chesney and Danielle Citron, "Deep Fakes: A Looming Crisis for National Security, Democracy and Privacy?" *Lawfare*, February 21, 2018, www.lawfareblog.com/deep-fakes-looming-crisis-national-security-democracy-and-privacy.

[42] The United States is currently not organized and lacks important capability for this mission. Thus, implementation of this recommendation would likely require the establishment of a national joint task force that combines the unique skills, capabilities, and authorities of the CIA, NSA, SOCOM, CYBERCOM, and DHS.

[43] Adrian Croft and Geert De Clercq, "France Fights to Keep Macron Email Hack from Distorting Election," *Reuters*, May 6, 2017, www.reuters.com/article/us-france-election/france-fights-to-keep-macron-email-hack-from-distorting-election-idUSKBN1820BO.

[44] See, for example, Alina Polyakova and Spencer P. Boyer, *The Future of Political Warfare: Russia, The West, and the Coming Age of Global Digital Competition*, Brookings, March 2018, 18, www.brookings.edu/research/the-future-of-political-warfare-russia-the-west-and-the-coming-age-of-global-digital-competition/.

[45] Zi Chu, et al., "Who Is Tweeting on Twitter: Human, Bot, or Cyborg?" Proceedings of the 26th Annual Computer Security Applications Conference, 2010, dl.acm.org/citation.cfm?id=1920265.

[46] For a relatively good summary, see: en.wikipedia.org/wiki/Facebook%E2%80%93Cambridge_Analytica_data_scandal.

[47] Perhaps most notably in its codification of a "right to be forgotten."

[48] See Daisuke Wakabayashi, "California Passes Sweeping Law to Protect Online Privacy," *New York Times,* June 28, 2018, www.nytimes.com/2018/06/28/technology/california-online-privacy-law.html.

"While regulating technology should be done in a limited manner—technology is always going to move more quickly than policy, and relying too heavily on regulation risks fighting the last battle while missing the next war—we should close off the pathways we know have been exploited in the past while leaving room to address the vulnerabilities of the future."

–LAURA ROSENBERGER

Countering Technologically Driven Information Manipulation

Laura Rosenberger
Director and Fellow, Alliance for Securing Democracy
German Marshall Fund

Just a few years ago, social media was widely hailed as a democratizing force—a tool for activists and democratic figures in authoritarian states to organize and share information with a breadth and speed never seen before. At the outset of the Arab Spring, President Obama commented, "We've seen the incredible potential for technology to empower citizens and the dignity of those who stand up for a better future."[1] Policy experts, such as Sascha Meinrath, then-director of New America's Open Technology Initiative, agreed, observing that "Social media have become the pamphlets of the 21st century, a way that people who are frustrated with the status quo can organize themselves and coordinate protest, and in the case of Egypt, revolution."[2]

Fast forward a few years, and the narrative has been turned on its head, with conversations focused on the threat technology poses to democracy. Congressional and Parliamentary hearings in the US, UK, and EU in the last few months have examined the ways social media has been manipulated to undermine democracy. Senators Mark Warner and Marco Rubio recently wrote that "Twenty-first century social media tools have the potential to further erode public confidence in western institutions and undermine the shared sense of facts that is supposed to be the foundation of honest political debate,"[3] while *The Economist* directly asked: "Do social media threaten democracy?"[4]

How did our view change so drastically in such a short amount of time? As with all innovations, online information platforms are not inherently positive or negative. These "disruptive technologies," as they have often been called, hold the potential to disrupt both authoritarian regimes AND democracies. And in the case of disrupting democracies, our adversaries discovered and exploited these tools before we realized the potential threat existed.

There was no bigger wake-up call to this vulnerability than the Russian government's weaponization of digital and information technology to interfere in the US 2016 presidential election, and its continued efforts to undermine US and European democracies. While propaganda and information warfare are not new tactics, social media and online information platforms allowed the Russian government to supercharge these efforts to a new scope and scale.

Through its campaign to undermine democracy in the United States, the Russian government exhibited a range of methods and tactics that have come to characterize the new age of information manipulation. As technology evolves, so will these threats, and other authoritarian actors are also learning from the Kremlin's success. This paper seeks to provide insight about these threats, their composition, and their future development. This paper will also highlight lessons learned from European countries' responses to such tactics and identify steps, both short and long term, that should be taken by the US government and the tech community to defend our democracy from online information manipulation.

The Age of Information Manipulation

While much discussion of the Russian government's exploitation of social media has focused on the injection and spread of false information, the tactics the Kremlin and its proxies have used to attack democracies are broader and more complex. Using social media's anonymity, the ability to transcend borders, and the power to connect both instantly and at scale, these operations have utilized fake personas and pages that masquerade as Americans or American political organizations[5] and develop followings among targeted audiences. These operations included political advertising aimed at social issues and specific candidates, taking advantage of both legal loopholes and the micro-targeting tools that power all online advertising. They even created fake web pages[6] designed to appear as local news outlets in order to gain credibility. Automated tools (including bots) and machine learning, along with dedicated study, allowed for the manipulation of algorithms to distort the information space—swaying search results and making certain content (especially divisive and extreme messages) seem more prevalent.[7]

The Russian government employs these varied tactics across a range of mediums to shape American perceptions and discussion. Every online platform is part of a broader information ecosystem, and each plays a different role in the Kremlin's manipulation of the information space. At times, information is laundered from one platform to another—and sometimes to traditional media—in order to mask its origin, gain credibility, and achieve broader reach.[8] What starts online does not always remain there, with fake personas organizing Americans to take actions in real life, including attending protests and taking other political actions.[9] The goal of such operations is not necessarily to promote an alternative view, but rather to undermine the very idea of truth, as well as to exploit societal cleavages and promote extreme positions to widen divisions and increase polarization.[10] The Kremlin's disinformation campaigns often do not pick sides, instead promoting contradictory and often inflammatory arguments on both sides of a debate to sow chaos and conflict. The effect creates a distorted image of reality for consumers of information.

As the Russian government's operations have shown, cyberattacks can also be used as part of information operations. During the 2016 US election, Kremlin cyber operatives hacked political organizations and campaigns to steal information that was later weaponized through public releases—including through fake social media personas and front websites operated by the same network that conducted the hacks.[11] The Department of Justice's October 4 indictment of Russian military intelligence officers for hacking and information operations related to the Olympics doping scandal also underscored the connection between these tools, and their use against a broad range of targets to manipulate debate and public opinion.[12] The public release of forged or altered documents has also been used to manipulate opinions; for example, in one case in Sweden,[13] forged documents released via social media reached citizens across the country and were featured in town hall meetings before being debunked. In the United States, election systems in numerous states have also been directly targeted by hackers; most recently, a DDoS attack emanating from an apparently foreign computer targeted election systems in Knox, Tennessee, during the primary in May 2018.[14] For systems and institutions that depend on trust, such as the press and elections, adversaries need not even alter votes or information; simply raising doubts about whether they have been compromised can undermine faith in their integrity.

As bad as this sounds, what we have seen to date will likely seem quaint in the future. Artificial intelligence (AI) is already being used to create manipulated video and audio content indistinguishable to the human eye and ear—so called "deepfakes." Not only will the future bring new kinds of malicious content, but with more people obtaining news and other information primarily online, the algorithms that order this information will hold increasing sway over what we see—and manipulation of them therefore will have a larger impact. As advertising technology increasingly powers the internet, and more of our personal information is stored electronically, data will become an even more valuable resource, subject to misuse and manipulation. The growing pervasiveness of the Internet of Things will also enlarge the surface area for hackers, allowing devices in our homes or our pockets to become weapons against critical infrastructure and other targets. Meanwhile, AI will reduce the cost of conducting cyberattacks—with an army of computers able to replace an army of human hackers. This will lead to a higher frequency and volume of attacks, and potentially increased anonymity and therefore more difficult attribution of blame.[15]

Moreover, the use of these tactics is not limited to Russia; in fact, the Chinese Communist Party (CCP) is developing the means to use sophisticated technology to manipulate the information space in ways that will surpass what we have seen from Russia. With a government that freely collects data on its 1.4 billion citizens and shares it with businesses,[16] China is a gold mine for AI development, which relies on large amounts of data for training those systems. At the same time, the CCP has been investing in methods to harness AI for social control in ways that pose a direct challenge to democracy. The *New York Times* recently described how China is combining AI and data to construct a "high-tech authoritarian future."[17] Using facial recognition technology and a system of 200 million cameras nationwide, China has begun to implement a social credit score system that will be fully operational by 2020. The system assigns every citizen a score that goes up or down based on everyday activities, and that score determines what kind of loans, housing, or transport an individual can get.[18] The Chinese government is using facial recognition software built on AI to recognize individuals through cameras when they are caught breaking laws,[19] allowing the government to store vast amounts of data on all its citizens. Nowhere is China honing these technologies more clearly than in Xinjiang. According to Human Rights Watch,[20] the CCP is collecting DNA, voice samples, and other biometric data from members of the Uyghur Muslim community and is now installing QR code panels on the homes of community members to allow authorities instant access to the personal details of those living there. By scanning the code on a house, authorities can see the identities and personal information of residents as well as how many potentially dangerous tools they own.

And while the CCP is developing these tools at home, it may later export them beyond its borders, much as the Russian government did. Indeed, the CCP is already using technology to engage in political interference outside its borders. The chat app Line has been used to spread disinformation in Taiwan,[21] and a recent report indicated that Chinese hackers infiltrated the opposition party in Cambodia[22] ahead of that country's elections. Senator Marco Rubio recently warned that "the [Chinese Communist] Party is increasingly exporting its authoritarianism abroad, trying to suppress speech, stifle free inquiry, and seek to control narratives around the world."[23] China has even begun to use technology owned by its companies to censor discussion outside its borders—for instance, on the chat app WeChat.[24] China has pressured foreign tech companies to censor content on their platforms; in one case, Chinese authorities pressured Facebook to take down the account of a Chinese business tycoon living abroad because of content he posted critical of Beijing.[25] In a world of information control where China pushes the Great Firewall beyond its borders, Beijing could simply create its own information reality.

Meanwhile, China is using access to its market as leverage to compel US tech companies to take steps that could give the country access to both data and technology, such as cooperative deals with Huawei[26] and Apple's agreement to locate some of its cloud servers in China—a move that was forced by Chinese cybersecurity legislation.[27] Google is reportedly working on a secretive project with the Chinese government to launch a censored, CCP-approved version of its search engine in the country, which could give the Chinese government unprecedented access to Google's data.[28] Chinese companies have also begun to acquire access to Americans' data, including through purchase of data-rich companies, like the LGBT-focused dating app Grindr.[29] These measures could provide the CCP with the sophisticated ability to target and manipulate Americans, should it choose to do so.

Lessons from Europe in Protecting Democracy

While the use of technology by foreign powers to interfere in democracies is new to many Americans, many of the techniques that have now been used against the United States were perfected over the past decade in the Baltics, Georgia, and particularly Ukraine. While there are differences in the European context—in many places, fewer people get their news online, and some online platforms popular in the United States, like Twitter,[30] have lower penetration—lessons from Europe's experience with information operations broadly, and its approach to technology-enabled information manipulation in particular, can inform how the United States tackles this problem. Key lessons from Europe include the importance of exposing malign activities, building resiliency in populations to withstand them, addressing the underlying issues around data protection, mustering the political will and capacity to coordinate a whole-of-government response, and articulating a clear deterrent message.

On a transnational scale, the EU has taken several steps to increase the ability of its member states, and the EU itself, to combat hostile information operations. In April 2018, the EU published a communication, based on the findings of a High-Level Expert Group, on combatting disinformation.[31] It outlined four key principles: improving transparency in the way information is produced or sponsored, promoting a diversity of information, fostering credible information, and fashioning inclusive solutions to disinformation. The communication focuses on cooperating with the private sector (including online platforms and advertisers), governments, and civil society (including "trusted flaggers," journalists, and media groups). Finally, it embraces the development of technological solutions to these challenges—including blockchain and AI—to verify information, as well as long-term solutions like improved media literacy. While the communication presents a holistic, forward-looking approach, it is yet unclear how it will be interpreted and implemented by member states and what impact it will have on future information operations against the EU.

In September 2018, European Commission President Jean-Claude Juncker also announced several new measures and recommendations to help secure "free and fair elections" within the EU.[32] Juncker encouraged member states to set up "national election cooperation network[s]" of relevant authorities to coordinate counter-interference efforts, called for greater transparency in political advertising, and recommended improved cybersecurity for national authorities, political parties, and media. The Commission also announced legislation to tighten rules on European political party funding and proposed the creation of a "Network of cybersecurity competence centres" to share resources and best practices on cybersecurity.

In practice, the EU's approach to countering information operations, along with NATO's, has largely applied a strategic communications frame and focused largely on elections. This has resulted in a focus on fact-checking and debunking false content targeted at those institutions or their member states in addition to ensuring a factual, positive message about them. While necessary, this frame is not sufficient to address either the scale or scope of the problem—particularly given how much of online information operations are about manipulating the information space to widen divisions, increase polarization, and promote extreme narratives.

Outside of explicit steps to combat information operations, the EU has also been active in managing online platforms. European countries tend to be less romantic about technology than Americans, and European governments have been more forward-leaning in regulating technology and data. In particular, Europe has taken steps with respect to protection of individuals' data. The EU's General Data Protection

Regulation (GDPR) went into effect earlier this year, governing how companies can collect data on individuals and how that data should be protected and providing EU residents with both "the right to be forgotten" and requirements for informed consent to data collection.[33] A number of companies, including Facebook, Google, Amazon, and Twitter, announced they would begin enforcing GDPR worldwide. The Commission also worked with representatives of online platforms, leading social networks, advertisers, and advertising industry to develop a self-regulatory Code of Practice on disinformation, released in September 2018.[34] In September 2018, the EU Commission also proposed new rules that would focus on content regulation, requiring online platforms to remove "illegal terror content" within one hour of it being flagged by national authorities.[35] Repeatedly failing to remove content quickly enough could result in a fine of 4 percent of a platform's global annual revenue.

Outside of the EU structure, European countries have also taken individualized steps to combat information operations. Some member states have been particularly inclined to focus on content regulation. In Germany, a hate speech law that went into effect in 2017 requires companies to remove illegal content within 24 hours of being flagged by users or face up to 50 million euro fines.[36] A draft law recently passed by the French National Assembly would give judges emergency power to block "fake" content during sensitive election periods, while also requiring social networks to flag false reports, pass data on such articles to authorities, and make public their efforts against "fake news."[37] Although these measures indicate a recognition of the serious threat posed by information operations, they also come with inherent risks. Focusing on content results in a narrow, reactionary, and temporary response to information operations, and also creates challenges with respect to free speech—an important element of democracy.

The experiences of other European countries reveal two positive lessons for US policy makers: the importance of exposing information operations and of building resiliency in populations to withstand them. The former facilitates deterrence when combined with cost-raising measures, and the latter diminishes the impact of the operations. Raising awareness of malign activities, promoting media literacy, and teaching critical thinking can build resiliency within a population. For example, Finland has already begun a program to train public officials in recognizing and combatting disinformation.[38] It has also invested in education and media literacy training; some of these steps appear to have forced Sputnik's Finnish branch to close due to low readership.[39]

Other European government efforts present a model for creating cross-cutting structures to ensure that analysts and policy makers see the full threat picture and are able to effectively coordinate responses to foreign interference. The Swedish Civil Contingencies Agency (MSB) is an example of such an effort; it works across other relevant agencies to coordinate monitoring of and response to issues of civil protection, public safety, emergency management, and civil defense—and includes foreign interference and hybrid activities in those areas. The effects of this approach on facilitating a unified approach to combatting interference are already evident. Ahead of its September 2018 elections,[40] Sweden invested in media literacy education and media monitoring, including training local media outlets on the threat and working proactively with social media platforms to warn them of disinformation campaigns;[41] trained election workers across the country on identifying disinformation and interference efforts; and educated the public on these issues.[42] While this approach is not replicable in a larger government like in the US, it is a model of the kind of principles that can guide a whole-of-government approach.

It is important to note that while these approaches focus almost exclusively on the defensive side of the equation, deterrence is also critical. Several European cases indicate that clear, high-level deterrent messaging can have an effect. There were significant concerns that elections in France and Germany last year could have been marred by the same kinds of attacks that plagued the US 2016 election. But both governments sent clear signals to Moscow warning against interference and the consequences if there were to be such activity. The Swedish government employed similar messaging surrounding its recent elections, with Prime Minister Stefan Löfven promising to publicly expose interference efforts "without mercy."[43] In the future, these deterrent messages will need to be backed by the credible threat of responses in order to assure adversaries of the high cost for interference efforts.

Building a Response: Necessary Steps for the US Government

Building on the lessons our European allies have learned, and taking into account the specific challenges and domestic context in the United States, the US government should take several immediate steps to help bridge the technology and national security divide, secure itself against future interference efforts, and prepare a coherent, deterrent message.

This needs to start with developing meaningful channels for dialogue with the tech sector. As Mark Zuckerberg's hearings before Congress this spring painfully illustrated, policy makers often do not understand the very technology at the heart of these challenges, nor do they have visibility into new technology being developed. Though recent hearings by the Senate Select Committee on Intelligence exhibited much more constructive and informed engagement from policy makers with the tech community,[44] knowledge and understanding of technology and its implications remain uneven across Capitol Hill and the policy-making community. Similarly, developers working on new technical innovations have little understanding of national security threats and therefore how their creations could be manipulated or misused. As Alex Stamos, the former CSO for Facebook, warned in a letter to his colleagues: "[tech companies] need to think adversarially in every process, product, and engineering decision [they] make."[45] The gap in mutual understanding demands mechanisms to facilitate ongoing dialogue between the tech and national security communities. A consistent channel of communication will help identify potential vulnerabilities in new and existing technologies that can be exploited by adversaries and will allow for the mitigation of those risks.

Meaningful information exchange will also be necessary to identify malicious activity online, particularly on social media. The intelligence community and social media companies hold different information about such activity—the former about the actions and intentions by malicious actors, the latter about the activity actually happening on their platforms. Until these two data sets can be put together, neither the government nor tech companies will be able to see the whole threat picture. The two sectors will need to establish a mechanism to share data and to identify nefarious actors on social media platforms that are linked to foreign nation states, while ensuring protection of Americans' privacy and free speech.[46] Models of such mechanisms exist from counter-terrorism, cybersecurity, and financial integrity efforts.[47] Additionally, the Department of Homeland Security (DHS) and the FBI have established task forces for this purpose, but their discussions with technology companies remain nascent and face barriers.[48] These attempts are further hindered by the trust gap that persists between government—particularly the intelligence community—and Silicon Valley following Edward Snowden's revelations. Restoring this trust is imperative to address the challenges we face.

The US government also needs to be restructured to better address foreign interference. As former CIA and National Security Agency Director Michael Hayden has said, Russia's interference in the 2016 election "hit a seam between law enforcement and intelligence, between 'sigint' [electronic spying] and 'humint' [human spying], between state and federal agencies, between politics and policy."[49] Chris Krebs, under secretary of homeland security for the National Protection and Programs Directorate, recently spoke to this issue at a Congressional hearing, noting that there is no single set of jurisdictions around the issue of foreign interference, rather there are law enforcement authorities and there are counterintelligence authorities.[50] This challenge is therefore one that requires a whole-of-government effort—and that will require a meaningful National Security Council–led process, including the establishment of a senior counter–foreign interference coordinator, a measure included in the currently pending National Defense Authorization Act (NDAA).[51] A National Hybrid Threat Center within the Director of National Intelligence Office would also ensure cross-cutting analysis and integration of reporting across the government. The Department of Justice announced in mid-July a new policy of exposing foreign interference—a positive step that builds on lessons from our European partners that will serve both as a deterrent to our adversaries and make Americans more resilient to their activities.

More broadly, as Christopher Kirchhoff wrote for this forum two years ago, our national security institutions (including those focused on the homeland) need to be reshaped to deal with technological shifts.[52] He recommends, among other things, "positioning the White House to lead on issues of technology by altering the National Security Council [to include a Tech Policy Directorate] and deepening its integration with the Office of Science and Technology Policy."

The Russian government's operations have also taken advantage of several loopholes and vulnerabilities in US law and infrastructure. While regulating technology should be done in a limited manner—technology is always going to move more quickly than policy, and relying too heavily on regulation risks fighting the last battle while missing the next war—we should close off the pathways we know have been exploited in the past while leaving room to address the vulnerabilities of the future. Congress has the authority to address many of these gaps, and bipartisan legislation has been proposed for several of them.

First, the same rules should apply to online political advertising that already apply to such ads on TV, radio, and print outlets. The bipartisan Honest Ads Act[53] would do just this, but absent its passage, companies have begun to undertake their own patchwork of approaches to transparency and regulation of political ads. These varied approaches confuse consumers and have resulted in problematic steps like Facebook's decision to consider publishers political advertisers when their reporting discusses politics—contributing to the degradation of faith in journalism.

Second, we need to harden our election infrastructure against cyber threats. DHS has designated our election systems as critical infrastructure, and Congress should pass legislation to codify that designation. The Secure Elections Act,[54] a bipartisan piece of legislation in the Senate, would articulate important standards for election security and provide additional support to state election officials, including through funding.

Another promising set of proposed regulations would enhance the transparency of automated accounts, including the Bot Disclosure and Accountability Act,[55] which Senator Dianne Feinstein has introduced. This bill would mandate that the Federal Communications Commission create a rule requiring social media companies to disclose any bots on their platform; prohibit political candidates and parties from using bots; and limit political action committees, corporations, and labor unions from using bots in certain political advertising. At the same time, any such disclosure requirements should ensure that anonymity online—which remains an important and empowering force for activists in authoritarian countries—remains protected.

Drawing on the lessons from Europe, we also need to establish a credible deterrent to this malign activity. The Trump administration's recent executive order to impose sanctions on foreign countries interfering in US elections is an important step in this direction, but it will need to be backed by consistent messaging, exposure of interference efforts, and robust implementation.[56] The DETER Act,[57] which enjoys bipartisan support and is being revised to address some concerns, would provide a similar deterrent by requiring the director of national intelligence to inform Congress if a foreign country has interfered in a federal election and would mandate retaliatory sanctions. In principle, this would provide a credible deterrent mechanism to help the United States signal its resolve, willingness, and capability to respond to adversaries' interference efforts.

More broadly, the US needs to continue to update legal frameworks to reflect the evolving challenges posed by technology. This includes following Europe's lead in considering a data protection framework. While the General Data Protection Regulation (GDPR) goes too far for the United States, in the absence of meaningful action by US policy makers, Europe is effectively setting the rules—as tech companies will apply their responses to European regulations to their platforms globally. The 2019 NDAA included several updates and measures to address the national security implications of technology investments and exports.[58] The Financial Investment Risk Review Modernization Act of 2018 (FIRMA) and the Export Control Reform Act of 2018, both of which were included in the NDAA, include important improvements to the Committee on Foreign Investment in the United States (CFIUS) and to US export controls to help hinder harmful foreign investment in certain technologies and prevent the export of technological tools key to national security. Implementation of these measures, including robust export controls around technology, could help close loopholes around technology that have national security implications. Still, more reform will be necessary to address the growing challenges and risks inherent in future technological developments.

Finally, the US Government should lead the way in establishing effective and robust information sharing mechanisms with our partners and allies. Online foreign interference is a transnational problem and often seeks to undermine alliances between governments as well as the governments themselves. Additionally, as indicated by the lessons from European efforts, there is much to be said for cooperation and information sharing in developing strategies to defend democracies from online attacks. The G7 recently announced a mechanism for information sharing both among governments and between governments and the private sector. In July 2018, NATO leaders also committed to establishing "counter-hybrid support teams"[59] to provide assistance to alliance members in preparing and responding to hybrid threats. While these measures are promising, it is yet unclear how they will be implemented and operationalized. As such, the United States should take a leadership role in building an open and consistent dialogue between its partners and allies—through organizations like NATO—for collaboration and sharing best practices.

The Role of the Private Sector

Just as government needs to step up and rethink its approach to technological threats to democracy, so too does the tech community.[60] Silicon Valley was also too slow to recognize the exploitation of its platforms and has too often taken a defensive approach in response. This needs to change. While progress has been made over the past year, with all major tech companies now admitting a problem exists, their approaches often remain reactive and retrospective as opposed to proactive efforts to detect and halt information operations going forward. Although senior executives from Facebook and Twitter acknowledged in recent testimony that they were too slow to recognize the problem and still need to do more to address it, they remain short on concrete solutions.[61] Additionally, Google's failure to send a representative to the Senate Select Committee on Intelligence hearing was seen as a sign that one of the tech community's biggest leaders does not take its role in this issue seriously.

Given the importance of exposure and awareness-raising to deterring and building resilience against this activity, the tech community must be transparent about continued malicious activity on its platforms as well as the steps being taken to combat this activity. Additionally, for most platforms, outside researchers are significantly limited in their ability to see and understand the malicious activity happening on them, therefore also limiting their ability to help develop solutions and provide accountability. It is essential that companies begin meaningfully sharing data with outside researchers in a manner that protects users' privacy. Further, just as government and tech companies need to develop means to share data, online information platforms also need to meaningfully share information with one another on the exploitation of their platforms by malicious actors, who work across platforms to manipulate the broader information ecosystem. This will raise awareness within the tech community and allow companies to get ahead of manipulation of their platforms. Recent efforts by social media platforms to take down foreign influence operations have shown progress in collaboration between companies and with law enforcement, but cooperative relationships should be institutionalized to ensure future coordination and information sharing.[62]

Malicious activity online can be targeted from one of two directions—the content (the end point) or the behavior/operations/structure (the origin). As previously mentioned, addressing information manipulation from the content perspective presents significant challenges in terms of the First Amendment and free speech, particularly when the content is political in nature, and also takes a reactive posture that inevitably ends in a game of whack-a-mole. Moreover, given that much of this activity is more about manipulation of the information space rather than spreading demonstrably false content (although there is plenty of that too), focusing on such content will only address a fraction of the problem. And if our interest is in protecting our democracy, steps that would undermine it—like policing speech—would only accelerate the efforts to weaken it.

Additionally, in the United States, unlike in Europe, Section 230 of the Communications Decency Act states that companies are not responsible for information on their platforms. Online platforms have used this provision to maintain that they are not publishers, rather simply the pipes through which information flows to users. But in reality, the algorithms that order and present information to their users DO dictate what those users see. Senator Ron Wyden,[63] the architect of Section 230, has underscored that the law was intended to work both as a shield and a sword, protecting companies from liability for content but requiring them to police their platforms. However, though tech companies have regularly used the shield, Wyden has suggested they have "sat on their hands with respect to the sword." While their legal obligations may be limited, these companies do have an interest in ensuring that their platforms are not co-opted for nefarious purposes, as authentic communications and presentation of accurate information is essential to maintaining users' trust.

The good news is there are increasingly discernible patterns to the actions of progenitors of information operations, and many of their behaviors are in direct violation of platforms' terms of service. Approaches that focus on the operational and structural aspects of information operations, such as Facebook's recent targeting of coordinated inauthentic behavior,[64] are likely to have a more sustainable result. New and developing technologies may provide effective tools for this approach. One technical approach could be to build on the hashing—or digital fingerprinting—mechanism that is used to identify and share information on terrorist content and child pornography.[65] This could help prevent the manipulation of images or the spread of disinformation that has already been flagged. Another example of technical improvements is the increasing removal of fake and automated accounts by major social media platforms, which is in part a result of the development and application of tools to improve back-end detection of such activity. As concerning as deepfakes are, there may be technical solutions to ensure that such content is identified before it can spread, such as giving images a "watermark" URL that allows comparison to the original to determine if it has been altered.[66] But technological approaches will never fully address the challenge, and ensuring sufficient human resources to oversee and address these challenges will be critical.

At the same time, tech companies will also need to empower platform users to reduce the effectiveness of information manipulation online. Ensuring that users have context about why they are seeing certain information, and the origins of it, allows for better assessment of its veracity. Making opt-in the default for data collection, instead of opt-out, will also give users better control over how they can be targeted by online advertisers. At the same time, fact-checking has had mixed results—with some evidence indicating that labeling information as disputed actually reinforces people's belief in it or drives additional traffic to it.[67]

Although these steps will help tech companies protect users from online manipulation, as long as targeted advertising remains the core driver of their business model, there will be limits to the effectiveness of policy and technological changes. A report from New America found that "the central problem of disinformation corrupting American political culture is not Russian spies or a particular social media platform. The central problem is that the entire industry is built to leverage sophisticated technology to aggregate user attention and sell advertising."[68] Any long-term solution to manipulation will require Silicon Valley to assess ways to reduce the manipulation inherent in this model. Alex Stamos, Facebook's departing CSO, advocated that Facebook "intentionally not collect data where possible, and to keep it only as long as we are using it to serve people," as well as "[deprioritize] short-term growth and revenue."[69]

It's on Us

It is critical—and urgent—that both government and the private sector take meaningful steps to address the challenges that technology poses to our democratic institutions. But that alone is not sufficient. As former Secretary of Homeland Security Michael Chertoff has underscored: "Russia could not plant seeds of dissension if there was not already fertile ground."[70] Addressing the underlying societal issues these operations exploit, and reducing division and polarization, will be critical to inoculating against these threats.

Over the long term, reinforcing critical thinking and media literacy skills will be essential to ensuring that consumers of information—facing an ever-increasing amount of it—have the necessary tools to be discerning. As Congressman Will Hurd has noted, "We all know stranger danger. Why are you listening to a stranger on social media?"[71] But as educators across the world have discovered, critical thinking is not always easy to teach.

We are likely only experiencing the beginning phase of these technologies fundamentally altering our democratic processes. The rapid advancement of technology means that its impact will only accelerate. In this dynamic environment, it is vital that governments, tech companies, and civil society take meaningful steps now to address these challenges, including creating new structures and ways of doing business when necessary, to ensure that technology can remain a democratizing force instead of one that undermines democracy itself.

Laura Rosenberger is the director of the Alliance for Securing Democracy and a senior fellow at The German Marshall Fund of the United States (GMF). Before she joined GMF, she was foreign policy advisor for Hillary for America, where she coordinated development of the campaign's national security policies, messaging, and strategy. Prior to that, she served in a range of positions at the State Department and the White House's National Security Council (NSC). As chief of staff to Deputy Secretary of State Tony Blinken and later Deputy National Security Advisor Blinken's senior advisor, she counseled on the full range of national security policy. In her role at the NSC, she also managed the interagency Deputies Committee, the US government's senior-level interagency decision-making forum on our country's most pressing national security issues. Ms. Rosenberger also has extensive background in the Asia-Pacific region, particularly Northeast Asia. She served as NSC director for China and Korea, managing and coordinating US policy on China and the Korean Peninsula, and in a variety of positions focused on the Asia-Pacific region at the Department of State, including managing US–China relations and addressing North Korea's nuclear programs. She also served as special assistant to Under Secretary of State for Political Affairs Bill Burns, advising him on Asia-Pacific affairs and on nonproliferation and arms control issues. Ms. Rosenberger first joined the State Department as a presidential management fellow.

[1] "Whiteboard," *Politico*, www.politico.com/politico44/perm/0211/obama_on_egypt_9cfd04c4-8b4a-4d5e-a7e3-efb791c90bf4.html.

[2] Sam Gustin, "Social Media Sparked, Accelerated Egypt's Revolutionary Fire," *Wired*, November 2, 2011, www.wired.com/2011/02/egypts-revolutionary-fire.

[3] Mark Warner and Marco Rubio, "Warner & Rubio: As Trump Meets Putin, We'll Spotlight and Resist Russian Aggression," *USA Today*, July 12, 2018, www.usatoday.com/story/opinion/2018/07/12/trump-putin-helsinki-summit-resist-russian-aggression-column/776617002/.

[4] "Do social media threaten democracy?" *The Economist*, November 4, 2017, www.economist.com/leaders/2017/11/04/do-social-media-threaten-democracy.

[5] Scott Shane, "The Fake Americans Russia Created to Influence the Election," *New York Times*, September 7, 2017, www.nytimes.com/2017/09/07/us/politics/russia-facebook-twitter-election.html.

[6] Bradley Hanlon and Grant Bennett, "Twitter Release Reveals the Kremlin's News Impersonation Game," Alliance for Securing Democracy, June 21, 2018, securingdemocracy.gmfus.org/twitter-release-reveals-the-kremlins-news-impersonation-game/.

[7] Bradley Hanlon, "From Nord Stream to Novichok: Kremlin Propaganda on Google's Front Page," Alliance for Securing Democracy, June 14, 2018, securingdemocracy.gmfus.org/from-nord-stream-to-novichok-kremlin-propaganda-on-googles-front-page/.

[8] Kirill Meleshevich and Bret Schafer, "Online Information Laundering: The role of social media," Alliance for Securing Democracy, January 9, 2018, securingdemocracy.gmfus.org/online-information-laundering-the-role-of-social-media/.

[9] Scott Shane, "How Unwitting Americans Encountered Russian Operatives Online," *New York Times*, February 18, 2018, www.nytimes.com/2018/02/18/us/politics/russian-operatives-facebook-twitter.html.

[10] USA v. Viktor Borisovich Netyksho, et al., Department of Justice, July 13, 2019, www.justice.gov/file/1080281/download.

[11] USA v. Viktor Borisovich Netyksho, et al., Department of Justice, July 13, 2019, www.justice.gov/file/1080281/download.

[12] Ellen Nakashima, Michael Birnbaum, and William Booth, "U.S. Indicts Russian Spies in Hacking Campaign Linked to Olympics Doping Scandal," *Washington Post*, October 4, 2018, www.washingtonpost.com/world/europe/britain-directly-blames-russian-military-intelligence-for-broad-range-of-cyberattacks/2018/10/04/13a3a1f8-c7b6-11e8-9158-09630a6d8725_story.html?utm_term=.1b9b1ae86327.

[13] Neil MacFarquhar, "A Powerful Russian Peapon: The Spread of False Stories," *New York Times*, August 28, 2016, www.nytimes.com/2016/08/29/world/europe/russia-sweden-disinformation.html.

[14] Adrian Sainz, "Cyberattack on Tennessee election website preceded outage," *Associated Press*, May 11, 2018, www.usnews.com/news/best-states/tennessee/articles/2018-05-11/firm-says-tennessee-election-web-site-may-have-been-attacked.

[15] Cornell University, "The Malicious Use of Artificial Intelligence: Forecasting, Prevention, and Mitigation," February 20, 2018, arxiv.org/abs/1802.07228.

[16] Louise Lucas and Richard Waters, "China and US Compete to Dominate Big Data," *Financial Times,* May 1, 2018, www.ft.com/content/e33a6994-447e-11e8-93cf-67ac3a6482fd.

[17] Paul Mozur, "Inside China's Dystopian Dreams: A.I., Shame, and Lots of Cameras," *New York Times,* July 8, 2018, www.nytimes.com/2018/07/08/business/china-surveillance-technology.html.

[18] Simina Misttreau, "Life Inside China's Social Credit Laboratory," *Foreign Policy*, April 3, 2018, foreignpolicy.com/2018/04/03/life-inside-chinas-social-credit-laboratory/.

[19] Alexandra Ma, "These Are the Things that Can Get You Punished Under China's Creepy 'Social Credit' System—From Fake News to Jaywalking," *Business Insider*, April 14, 2018, www.businessinsider.com/china-social-credit-system-things-you-can-do-wrong-and-punishments-2018-4.

[20] Human Rights Watch, "'Eradicating Ideological Viruses' | China's Campaign of Repression Against Xinjiang's Muslims," September 9, 2018, www.hrw.org/report/2018/09/09/eradicating-ideological-viruses/chinas-campaign-repression-against-xinjiangs.

[21] Aaron Wytze, "Killing Fake News Dead on Taiwan's Most Popular Messaging App," g0v.news, February 15, 2017, g0v.news/killing-fake-news-dead-on-taiwans-most-popular-messaging-app-c99d93582cbe.

[22] Abby Seiff, "Chinese State-Linked Hackers in Large Scale Operation to Monitor Cambodia's Upcoming Elections, Report Says," *Time*, July 11, 2018, time.com/5334262/chinese-hackers-cambodia-elections-report/.

[23] "Video: Rubio Chairs China Commission Hearing on Digital Authoritarianism & the Global Threat to Free Speech," April 26, 2018, www.rubio.senate.gov/public/index.cfm/press-releases?id=628537F9-1AC3-48F3-B63D-49AB2BF9F0C2.

[24] Lulu Yilun Chen, "WeChat Censoring Messages Even Outside China, Study Says," *Bloomberg*, December 1, 2016, www.bloomberg.com/news/articles/2016-12-01/wechat-censoring-user-messages-even-outside-china-study-says.

[25] Paul Mozur, "China Presses its Internet Censorship Efforts Across The Globe," *New York Times*, March 2, 2018, www.nytimes.com/2018/03/02/technology/china-technology-censorship-borders-expansion.html.

[26] Shannon Liao, "Why Facebook's Secret Data-Sharing Deal with Huawei Has the US Concerned," *The Verge*, June 8, 2018, www.theverge.com/2018/6/8/17435764/facebook-data-sharing-huawei-cybersecurity.

[27] Paul Mozur, Daisuke Wakabayashi, and Nick Wingfield, "Apple Opening Data Center in China to Comply with Cybersecurity Law," *New York Times,* July 12, 2017, www.nytimes.com/2017/07/12/business/apple-china-data-center-cybersecurity.html.

[28] Li Yuan and Daisuke Wakabayashi, "Google, Seeking a Return to China, Is Said to Be Building a Censored Search Engine," *New York Times,* August 1, 2018, www.nytimes.com/2018/08/01/technology/china-google-censored-search-engine.html?partner=IFTTT.

[29] Josh Rogin, "Can the Chinese Government Now Get Access to Your Grindr Profile?" *Washington Post,* January 12, 2018, www.washingtonpost.com/news/josh-rogin/wp/2018/01/12/can-the-chinese-government-now-get-access-to-your-grindr-profile/?utm_term=.012ba397f869.

[30] Statista, "Leading Countries Based on Number of Twitter Users as of April 2018," www.statista.com/statistics/242606/number-of-active-twitter-users-in-selected-countries/.

[31] Bradley Hanlon, "The 'European Approach' to Fighting Disinformation: Lessons for the United States," Alliance for Securing Democracy, April 27, 2018, securingdemocracy.gmfus.org/the-european-approach-to-fighting-disinformation-lessons-for-the-united-states/.

[32] European Commission, "State of the Union 2018: European Commission Proposes Measures for Securing Free and Fair European Elections," September 12, 2018, europa.eu/rapid/press-release_IP-18-5681_en.htm.

[33] European Parliament and the Council of the European Union, "EU GDPR: Summary of key provisions," December 14, 2015, iapp.org/media/pdf/resource_center/Promontory_GDPR_compromise.pdf.

[34] EU Commission, "Code of Practice on Disinformation." ec.europa.eu/digital-single-market/en/news/code-practice-disinformation.

[35] Ivana Kottasová, "Europe Could Hit Tech Companies with Huge Fines over Terrorist Content," *CNN*, September 12, 2018, money.cnn.com/2018/09/12/technology/online-terrorist-content-eu/index.html.

[36] "Germany Starts Enforcing Hate Speech Law," *BBC News*, January 1, 2018, www.bbc.com/news/technology-42510868.

[37] Jacques Klopp, "France's Fake News Law Leaves Media Experts Uneasy," *Agence France Presse*, June 4, 2018, www.yahoo.com/news/frances-fake-news-law-leaves-media-experts-uneasy-145712000.html; Colleen de Bellefonds, "France Considers Criminalizing Fake News with Proposed Law," *US News & World Report,* July 25, 2018, www.usnews.com/news/best-countries/articles/2018-07-25/french-wrestle-with-tackling-fake-news-with-proposed-law.

[38] "US experts gird Finnish officials for information war," *yle*, January 22, 2016, yle.fi/uutiset/osasto/news/us_experts_gird_finnish_officials_for_information_war/8616336.

[39] Reid Standish, "Russia's Neighbors Respond to Putin's 'Hybrid War,'" *Foreign Policy*, October 12, 2017, foreignpolicy.com/2017/10/12/russias-neighbors-respond-to-putins-hybrid-warlatvia-estonia-lithuania-finland/.

[40] Kristine Berzina, "Sweden—Preparing for the Wolf, Not Crying Wolf: Anticipating and Tracking Influence Operations in Advance of Sweden's 2018 General Elections," Alliance For Securing Democracy, September 7, 2018, securingdemocracy.gmfus.org/sweden-preparing-for-the-wolf-not-crying-wolf-anticipating-and-tracking-influence-operations-in-advance-of-swedens-2018-general-elections/.

[41] Government Offices of Sweden, "A Practical Approach on How to Cope with Disinformation," October 6, 2017, www.government.se/articles/2017/10/a-practical-approach-on-how-to-cope-with-disinformation/.

[42] Michael Birnbaum, "Sweden Is Taking on Russian Meddling Ahead of Fall Elections. The White House Might Take Note," *Washington Post*, February 22, 2018, www.washingtonpost.com/world/europe/sweden-looks-at-russias-electoral-interference-in-the-us-and-takes-steps-not-to-be-another-victim/2018/02/21/9e58ee48-0768-11e8-aa61-f3391373867e_story.html?utm_term=.71ef99c0387e.

[43] Erik Brattberg and Tim Maurer, "How Sweden Is Preparing for Russia to Hack Its Election," *BBC News,* May 31, 2018, www.bbc.com/news/world-44070469.

[44] US Senate Select Committee on Intelligence, "Foreign Influence Operations' Use of Social Media Platforms," September 5, 2018, www.intelligence.senate.gov/hearings/open-hearing-foreign-influence-operations%E2%80%99-use-social-media-platforms-company-witnesses.

[45] Ryan Mac and Charlie Warzel, "Departing Facebook Security Officer's Memo: 'We Need to Be Willing to Pick Sides,'" *Buzzfeed*, July 24, 2018, www.buzzfeednews.com/article/ryanmac/facebook-alex-stamos-memo-cambridge-analytica-pick-sides.

[46] Alliance for Securing Democracy, "Report Launch: ASD's Policy Blueprint for Countering Authoritarian Interference in Democracies," June 26, 2018, www.gmfus.org/events/report-launch-asds-policy-blueprint-countering-authoritarian-interference-democracies.

[47] One example is the Global Internet Forum to Counter Terrorism (GIFCT), whose goal is to substantially disrupt terrorists' ability to promote terrorism, disseminate violent extremist propaganda, and exploit or glorify real-world acts of violence using our platforms by employing and leveraging technology; sharing knowledge, information, and best practices; and conducting and funding research. Also, the National Cyber Forensics and Training Alliance is a nonprofit partnership between industry, government, and academia to provide a neutral, trusted environment that enables two-way collaboration and cooperation to identify, mitigate, and disrupt cyber crime. Two models from the world of financial intelligence are the UK's Joint Money Laundering Intelligence Taskforce (JMLIT) and the United States' FinCEN Exchange.

[48] Sheera Frenkel and Matthew Rosenberg, "Top Tech Companies Meet with Intelligence Officials to Discuss Midterms," *New York Times*, January 25, 2018, www.nytimes.com/2018/06/25/technology/tech-meeting-midterm-elections.html.

[49] Chris Krebs, "Hearing Before the House Homeland Security Committee," July 11, 2018, www.washingtonpost.com/world/national-security/nsa-and-cyber-command-to-coordinate-actions-to-counter-russian-election-interference-in-2018-amid-absence-of-white-house-guidance/2018/07/17/baac95b2-8900-11e8-85ae-511bc1146b0b_story.html?utm_term=.8b285b052c10.

[50] Chris Krebs, "Hearing before the House Homeland Security Committee," July 11, 2018, www.washingtonpost.com/world/national-security/nsa-and-cyber-command-to-coordinate-actions-to-counter-russian-election-interference-in-2018-amid-absence-of-white-house-guidance/2018/07/17/baac95b2-8900-11e8-85ae-511bc1146b0b_story.html?utm_term=.8b285b052c10.

[51] House Armed Services Committee, "Reform and Rebuild: The Next Steps," accessed July 26, 2018, armedservices.house.gov/sites/republicans.armedservices.house.gov/files/wysiwyg_uploaded/FY19%20NDAA%20Conference%20Summary%20.pdf.

[52] Christopher Kirchhoff, "Reshaping National Security Institutions for Emerging Technology," in *America's National Security Architecture: Rebuilding the Foundation*, eds. Nicholas Burns and Jonathon Price (Washington, DC: Aspen Strategy Group, 2016).

[53] S.1989—Honest Ads Act, www.congress.gov/bill/115th-congress/senate-bill/1989.

[54] S.2261—Secure Elections Act, www.congress.gov/bill/115th-congress/senate-bill/2261.

[55] S.3127—Bot Disclosure and Accountability Act of 2018, www.congress.gov/bill/115th-congress/senate-bill/3127.

[56] Anne Gearan and Felicia Sonmez, "Trump Issues New Order Authorizing Additional Sanctions for Interfering in Upcoming US Elections," *Washington Post*, September 12, 2018, www.washingtonpost.com/politics/trump-issues-new-order-authorizing-additional-sanctions-for-interfering-in-upcoming-us-elections/2018/09/12/a90898a0-b6b0-11e8-a7b5-adaaa5b2a57f_story.html.

[57] S.2313—Defending Elections from Threats by Establishing Redlines Act of 2018, www.congress.gov/bill/115th-congress/senate-bill/2313/cosponsors.

[58] US House of Representatives, "John S. McCain National Defense Authorization Act for Fiscal Year 2019," July 25, 2018, www.congress.gov/115/crpt/hrpt874/CRPT-115hrpt874.pdf#page=613.

[59] NATO, "NATO's Response to Hybrid Threats," July 17, 2018, www.nato.int/cps/en/natohq/topics_156338.htm.

[60] Jamie Fly and Laura Rosenberger, "How Silicon Valley Can Protect US Democracy," *Foreign Affairs,* February 22, 2018, www.foreignaffairs.com/articles/united-states/2018-02-22/how-silicon-valley-can-protect-us-democracy.

[61] US Senate Select Committee on Intelligence, "Foreign Influence Operations' Use of Social Media Platforms," September 5, 2018, www.intelligence.senate.gov/hearings/open-hearing-foreign-influence-operations%E2%80%99-use-social-media-platforms-company-witnesses.

[62] Paresh Dave and Christopher Bing, "Facebook, Twitter Dismantle Disinformation Campaigns Tied to Iran and Russia," *Reuters*, August 21, 2018, www.reuters.com/article/us-facebook-russia-usa/facebook-twitter-dismantle-disinformation-campaigns-tied-to-iran-and-russia-idUSKCN1L62FD.

[63] Colin Lecher, "Senator Ron Wyden Reckons with the Internet He Helped Shape," *The Verge*, July 24, 2018, www.theverge.com/2018/7/24/17606974/oregon-senator-ron-wyden-interview-internet-section-230-net-neutrality.

[64] Facebook, "Taking Down More Coordinated Inauthentic Behavior," Facebook Newsroom, August 21, 2018, newsroom.fb.com/news/2018/08/more-coordinated-inauthentic-behavior/.

[65] Jeremy Kahn, "Tech Companies Identify, Remove 40,000 Terrorist Videos, Images," *Bloomberg*, December 4, 2017, www.bloomberg.com/news/articles/2017-12-04/tech-companies-identify-remove-40-000-terrorist-videos-images.

[66] Josh Constine, "Truepic Raises $8M to Expose Deepfakes, Verify Photos for Reddit," *TechCrunch*, June 20, 2018, techcrunch.com/2018/06/20/detect-deepfake/.

[67] Sam Levin, "Facebook Promised to Tackle Fake News. But the Evidence Shows It's Not Working," *Guardian*, May 16, 2017, www.politico.eu/article/fake-news-germany-elections-facebook-mark-zuckerberg-correctiv/.

[68] Dipayan Ghosh and Ben Scott, "Digital Deceit: The Technologies Behind Precision Propaganda on the Internet," New America, January 23, 2018, www.newamerica.org/public-interest-technology/policy-papers/digitaldeceit/.

[69] Ryan Mac and Charlie Warzel, "Departing Facebook Security Officer's Memo: 'We Need to Be Willing to Pick Sides,'" *Buzzfeed*, July 24, 2018, www.buzzfeednews.com/article/ryanmac/facebook-alex-stamos-memo-cambridge-analytica-pick-sides.

[70] Alliance for Securing Democracy, "Report Launch: ASD's Policy Blueprint for Countering Authoritarian Interference in Democracies," June 26, 2018, www.gmfus.org/events/report-launch-asds-policy-blueprint-countering-authoritarian-interference-democracies.

[71] German Marshall Fund, "Responding to Russia's Attacks on Democracy," January 10, 2018, www.youtube.com/watch?v=QTWGdy8gffw.

“The future of cyberwar is about the marriage of traditional hacking of systems and infrastructure with coordinated disinformation campaigns to hack the conversation.”

–JARED COHEN

Confronting Hybrid Warriors and the Disinformation Tactics They Use

Jared Cohen
CEO and Founder
Jigsaw

Societies cannot function without both physical and digital stability. Just as the physical state is destabilized by violence, the digital state is weakened by a combination of digital and cultural manipulation, a subversive type of sabotage that tries to hack not only power grids and networks, but also the public conversation. The relative ease with which a state or non-state actor can instigate a multidimensional conflict—i.e., cost/attribution benefits—means that all future wars will begin as cyberwars. Given the low cost, the lack of rules that govern responses, and the difficulties with attribution, the barrier of entry to cyber conflict makes the internet the newest and most confounding theater of war.

The future of cyberwar is about the marriage of traditional hacking of systems and infrastructure with coordinated disinformation campaigns to hack the conversation. We have a clear but evolving understanding of the former. Many people understand on some level that the world is already in a perpetual state of cyber conflict—with countries attacking each other's data and infrastructure in big and small ways every day—but it's hard for most observers to distinguish where the interstate digital conflict ends and where the "collateral damage" to civilians begins.

These wars are about information and disinformation, hacking, and manipulation. But more importantly they involve the digital deployment of people, new kinds of soldiers, militias, insurgents, and paramilitaries who are active across boundaries and reaching into sovereign territory to influence outcomes and foment chaos.

Whether we call it disinformation, a social media problem, or something else is irrelevant. What matters is that we understand the people involved and the tactics they deploy. This paper will describe four common tactics and two emerging threats.

Tactic 1: Digital Paramilitaries

When armies capture territory, they either do it overtly with uniformed soldiers, or they do it covertly with some combination of plausibly deniable assets—militias, paramilitaries, criminals, thugs, and mercenaries—who try to blend in with the local population and conceal their true purpose. The same is true online. What we call digital paramilitaries refers to a new category of hybrid warrior that builds elaborate fake identities online, cultivates influence through social networks, and then deploys that influence to manipulate public conversations or spread disinformation.

Acquire Accounts, Establish Cover Identities, Seed Influence

What do these militias actually look like in the digital context? A politically motivated actor identifies the trending topics and tensions in a society, combs social media for images and video of people who look the part of different demographics and constituencies, and then builds fake identities by cobbling it all together into various profiles on multiple platforms.

But simply having a synthetic militia that looks the part is not enough. Once that militia is deployed, it cannot successfully infiltrate without being seen as "authentic." Authenticity online often is conveyed through a user's discoverability and familiarity. Malicious actors may spend years carefully cultivating an individual persona to appear natural and credible, often through stealing profile pictures and other personal photos from unsuspecting online users to demonstrate the semblance of normal life. Another hallmark of online authenticity is the persona's consistent appearance across multiple online platforms. From posts on Twitter and Facebook to online blogs and websites to even courting interviews from real journalists, convincing digital militia have it all. The numerous references scattered across the internet ensure digital mercenaries are able to evade detection from even the most scrupulous online users.

For example, take one expertly crafted persona from the Internet Research Agency (IRA), that of "Jenn Abrams," an active conservative feminist blogger from the US. Jenn's postings appeared on a personal website, Twitter, and Medium and were even quoted in slew of news outlets, including *Mashable*, the *Washington Post, Buzzfeed, CNN*, and the *New York Times.* While much of the content was nonpolitical, some of the most popular posts were about slavery, the confederacy, and race that appeared designed to widen rifts within America's social fabric. The problem is Jenn was the product of one or more paid disinformation peddlers in Russia.

State and non-state actors like the IRA, a St. Petersburg–based private company conducting information operations for the Kremlin, employ professional disinformation peddlers whose compensation structures and key performance indicators would not be out of place at Western digital marketing firms: creation of fake profiles (often by stealing photos and images from other people's social media accounts) and daily quotas of two to three political posts, five nonpolitical posts, and 100 comments on other workers' posts. Those belonging to the "foreign desk" of the IRA were trained to know the nuances of American political issues on tax policies, LGBT rights, and the Second Amendment. They were armed with regularly refreshed troll farm manuals that outlined key disinformation narratives, together with hyperlinks to news stories and short summaries of how to comment on these articles, in order to incite American internet users and derail political discussions.

Once each individual member of a synthetic militia has earned its authentic status, the next step is for the adversarial state to command that militia, meaning get it to coordinate its activities. These man-made digital militias identify trending conversations (e.g., #blacklivesmatter, #bluelivesmatter, and even #TakeTheKnee, referring to the NFL) and are actively deployed to interact with people participating in these conversations to achieve a goal (e.g., exacerbate racial tensions). They draw people into conversations, stir up trouble, argue, manipulate, and organize physical activities that range from protests to rallies to martial arts classes.

Conduct Operations

Whether we call these militias digital insurgents, digital paramilitaries, or something else often depends on how much they want to blend into the crowd versus being seen as part of an organized group.

The cultivation of fake accounts aims to amplify partisan divisions on multiple sides of hot-button issues. For example, @tpartynews masqueraded as a pro-Trump American and attacked popular targets including Black Lives Matter. The account had nearly 22,000 followers on Twitter but was reportedly tied to Kremlin election interference. Russian-linked accounts also became active on the other side of this issue: #BlackLivesMatter IRA account @Blacktivist had 360,000 likes on Facebook (more than the verified Black Lives Matter account on Facebook!) and promoted real-life protests organized by real activists, such as one to commemorate the death of Freddie Gray in Baltimore. You would not have known these accounts' Russian connections from looking at them, though.

Facebook also reported that a Kremlin-linked source purchased Black Lives Matter–themed advertising to be shown to audiences in Baltimore and Ferguson, Missouri—both locations of recent racial incidents. Not only did the Kremlin create individuals and organizations on both sides of wedge issues, they also used targeted advertising to reach the audiences they believed would be most receptive to their distorted messages.

For example, the IRA created the aforementioned group called Black Fist, which had hundreds of thousands of followers across multiple social media platforms. They presented themselves as a spinoff of Black Lives Matter, created swag and compelling content, and actively engaged in public conversations. They paid black martial arts instructors in multiple boroughs of New York City to host classes sponsored by Black Fist and in exchange asked for images and video of the classes—presumably to have authentic content of black and white people hitting each other. In another example, a fake organization called BlackMattersUS was, in fact, a curated effort to appeal to the African American community by posting on issues related to police brutality and racism.

Like real-world fighters, synthetic militias get better with practice. But they still make mistakes—using a Russian phone number as the second factor in two-step verification, accidentally leaving geo-location on, or using the same operational tradecraft across dozens of accounts. Fake accounts that are coordinating within a broader organization may be working from a single instruction set, or they make common linguistic errors that can provide a semantic link between suspicious accounts. Accounts acting in concert may even exhibit unusual temporal alignment as they work together to push narratives at specific times. While no single fingerprint can conclusively tell us if an account is part of a foreign-sponsored disinformation campaign, taken together, these clues and more can help us identify the networks that may be working to subvert our democracy.

Tactic 2: Targeted Online Harassment, a.k.a. "Patriotic Trolling"

While some tactics involve indiscriminate, large-scale manipulation of public discourse, a second tactic is much more targeted, designed to take specific key influencers off the digital battlefield. Governments are increasingly leveraging the amplifying power of social media platforms to apply pressure to political opponents, often using the state security apparatus to coordinate distributed networks of social media accounts to target the same person. For the targets of this type of trolling—let's call it "patriotic trolling," since the harassment is often couched in terms of patriotism or nationalism against a "foreign" enemy—this can be a devastating tactic. Think of this as cyberbullying on steroids: when cyberbullying becomes better funded, better organized, and in some cases state sponsored. Many victims end up fleeing the country or narrowly escaping threats of violence. Many others are not so lucky.

Turkey provides an alarming case study into the evolving sophistication of so-called "patriotic trolling" attacks over just a few years. Fatih Tezcan, a prominent member of the ruling party, AKP, recently invited his 400,000 Twitter followers to incite an attack against *Cumhuriyet* reporter Ceyda Karan after she was convicted for incitement of hatred for publishing an image of the prophet Mohammed. Tezcan was also the most influential social media personality in a campaign against former *Cumhuriyet* editor-in-chief Can Dündar, who was targeted after he was convicted of publishing classified documents and narrowly escaped an assassination attempt.

Turkey has also used young government supporters to build its volunteer group of 6,000 "social media representatives," who are charged with promoting the party perspective and monitoring online discussions. In BBC World Service correspondent Selin Girit's case, bots using hashtags coined by AKP Ankara Mayor Melih Gökçek attacked Girit in more than 35,000 tweets and retweets. Girit received dozens of rape threats *per minute* from bots at the peak of the attacks.

These same tactics are also used to block key influencers from reaching virtual town squares, which prevents them from participating in the discourse. But the outcome is the same, which is to thwart key influencers from engaging the crowd. In the Philippines, a group of such trolls who call themselves Duterte Cyber Warriors (*Oplan Cyber Tokhang*), and operate a closed Facebook group, organized a "mass reporting" attack, designed to ensure that the legitimate Facebook accounts of Duterte critics were suspended or shut down.

Unfortunately, this practice is increasingly widespread. Freedom House published a recent study suggesting that online disinformation played an important role in at least eighteen recent elections, and in as many as thirty countries, the governments employed armies of trolls (the study refers to them as "opinion shapers") to spread government views, drive particular agendas, and engage—often hostilely—with critics of the government on social media.

Tactic 3: Social Chatbot Amplification

The deployment of synthetic militias, patriotic trolls, and the production of disinformation content are all significantly amplified with a third tactic, the automation of social network accounts to amplify a message beyond human scale.

For any victim of online mob harassment, the sheer volume of posts being directed toward you can be overwhelming. It is hard to believe that *this many* people all feel animosity toward one particular individual, especially if that person is not particularly famous or prominent. In many cases, that's because the accounts aren't real. Bots, or accounts that are controlled remotely, either with algorithms or with simple instructions to repost or modify other content, are a frequently used tool to amplify organized harassment campaigns. With basic technical knowledge and minimal resources, you can make ten people seem like 10,000 people online.

Even democracies have been exposed as using these tools to distort the public conversation. In 2012, South Korea's intelligence service posted more than 1.2 million deceptive Twitter messages supporting presidential candidate Park Geun-hye and criticizing her competitor. Park won the election, but the former director of South Korea's main spy agency was convicted and sent to prison when the deceptive practices were exposed.

In Mexico, protest leaders used the hashtag #RompeElMiedo ("break the fear") to coordinate a protest online, but soon the hashtag was overwhelmed by bogus posts from pro-government accounts and Enrique Peña–supporting bots known as "Peña bots"—making content that might have been useful to the protesters nearly impossible to find. The journalist Erin Gallagher, who has monitored this and other protests, has documented at least 75,000 automated Twitter accounts being used for these purposes in Mexico.

Tactic 4: Fake News

The tactic that gets the most public attention is the weaponization of what is often described as "fake news." Even before we discuss the varied and confused definitions of that term, it is important to distinguish between the production of bogus information online, which is primarily driven by "disinformation entrepreneurs" creating false news stories in order to make money from selling ads, from the more insidious and threatening practice of states using and amplifying that false information for political purposes.

As Jigsaw researchers learned from numerous interviews with fake news peddlers from Florida to Macedonia, fake news is primarily created and disseminated by opportunistic "entrepreneurs" who earn money from having these stories gain attention online and therefore drive advertising revenue. In other words, these entrepreneurs have a strong economic incentive to create such false material, and they are often working without any direct connection to the organizations that weaponize these stories for political purposes.

This false information often exploits popular myths or rumors to capture readers' attention online, which makes it a tempting weapon for digital propagandists looking to create confusion or propagate false narratives. These false stories can be algorithmically targeted to reach certain susceptible audiences in an attempt to "hack" the public conversation around certain issues.

The IRA specializes in this. The IRA spent more than a year creating dozens of social media accounts masquerading as local American news outlets. These accounts gradually amassed real followers, mostly from posting headlines from real news sites or by seeming to represent certain communities. They targeted a diverse cross-section of the United States both geographically and politically, with fake accounts representing people and organizations from Chicago, Los Angeles, Seattle, San Francisco, Boston, and beyond.

Despite the outsize public attention on the concept of fake news, it is unclear how effective it is as a tactic for propaganda and disinformation campaigns. Although fake news on the internet is arguably more sophisticated than other propaganda media (television, leaflets, print, radio), its efficacy is ultimately limited because it lacks the personal interaction that the other tactics exploit. Even if a piece of fake news "reached" 100 million people, that statistic doesn't necessarily indicate how many people actually saw that piece of information, much less let it affect their decision-making. But that's precisely the point—it is impossible to accurately measure its effect, so it is easy to overestimate it. Its very existence, compounded with our inability to measure it, destabilizes our culture.

Looking Ahead: Emerging Tactics on the Horizon

Understanding the current landscape of players and tactics in these multidimensional wars is a necessary but not entirely sufficient condition to developing policy prescriptions. We must also understand how these capabilities are likely to develop and how new capabilities might create unforeseen threats. It is worth subscribing to the view that once invented, if we can think of it, so too can our adversaries. This means considering the following hypothetical scenarios that combine emerging technological capabilities with existing tactics.

Falsified Video

Movie studios have long used computer-generated imagery to make science fiction films come to life, but advances in artificial intelligence suggest that this technology will soon be readily available for anyone to create realistic digital impersonations. This phenomenon is called "deepfakes," a portmanteau of deep learning (a sub-field of machine learning used to manufacture fake images) and fake videos. It is easy to imagine how this technology could wreak havoc on society.

We can expect the steep rise of politically motivated deepfakes. Imagine watching a video of the president of the United States making an important speech and not being able to tell whether the video you're watching in your social media feed is real. Or imagine during a contentious Congressional race, a falsified video emerged of one of the candidates going on a venomous, racist rant. Every piece of film would have to be carefully scrutinized, and it would be easy to cast doubt on the validity of any real footage by accusing it of being a deepfake. Without adequate and rapid means of detecting such falsified videos, any hobbyist or amateur filmmaker could create a convincing video that has real-world effects before anyone realized it was not real.

It is easy to envision more commercially motivated examples of the same technology being used to manipulate stock prices. Whether it is funds extracted through ransomware (WannaCry), hacking into and stealing from banks, or pursuing BTC, North Korea has proven its appetite for using its cyber capabilities for illicit fundraising. Imagine a scenario where North Korea creates a deepfake designed to make money from shorting stocks, say, for example, showing a CEO and CFO "secretly" caught talking over lunch about missing expectations in next week's quarterly earnings call. All they would have to do is put out the video on social media and tip off journalists to get the markets to (over)react. A few weeks or months prior, their affiliates who are active on the stock market would have bought puts, or options to sell during the period of the attack. So they buy at the plummeted price and exercise their options to sell at the pre-agreed price. They pocket the difference and are out of all positions by the time the "fake news" is debunked and the market recovers.

Buying War on the Black Market

The illegal arms trade has long provided a marketplace for deniable military operations around the world, as well as an easy way for militias to acquire military capability in a short period of time. The same is true for the digital side of conflict, where the dark web—the part of the internet that is not indexed by search engines and requires special software to access—serves as a thriving bazaar for nefarious tactics. Jigsaw's research has revealed a nascent but thriving online marketplace for purchasing anything from fake social media followers to entire propaganda campaigns. Vendors advertise their track record in both commercially and politically motivated campaigns. Prices are more affordable than you might think. In one case, 100,000 social media followers went for $700, and 10,000 posts on Twitter went for as little as $30. A complex propaganda campaign offering a variety of illegal services (i.e., email hacking, database scraping, intercepting television signals, falsifying security information) could be purchased for less than $10,000.

The providers of these types of services fall into two main categories. Nonpolitical actors provide a variety of services that could ostensibly be used for nonpolitical, commercial purposes, like paid social media followers to promote a product or cause. Political actors, on the other hand, specialize in unscrupulous tactics for overtly political purposes. One group claimed to have experience deploying these tactics in Georgia and Ukraine and were evaluating proposals to take them to Minnesota and Pennsylvania.

These dark web markets offer a wide range of services: writing political messages, disinformation content production, fake followers, identity theft, website hijacking, malware, distributed denial-of-service (DDoS) attacks, hacking of buildings (e.g., fire alarms) of political organizations, hacking WiFi networks, and hijacking email accounts, to name a few.

How Do We Strengthen Ourselves Against the Next Wave of Attacks?

Many of the strategic and policy changes required to confront this hybrid threat well into the future are long term and systemic. They will require marshalling the forces of the public and private sectors to shore up our infrastructural defenses, refine our understanding of social media platforms, and fortify our culture against attempts to manipulate public discourse. But there are also a number of shorter-term, tactical prescriptions that can provide significant insight into this new form of warfare and how we can be better positioned to meet future threats.

In order to prepare for the next generation of threats, we ought to study how hybrid warfare is transpiring in the most active digital war zones around the world. These active theaters of conflict allow us to examine these tactics in context—how they evolve and how effective our countermeasures can be. It is under these extreme circumstances that we'll develop the most innovative solutions.

NATO and its Western allies are understandably concerned about what Russia and other adversaries might do to attack the European Parliament elections in May 2019 and the US presidential election in November 2020. It remains unclear how active Russia has been in conducting operations designed to target the US midterm elections. But that does not mean that such operations are not occurring, nor should we interpret that as a diminishment in Russia's interest in leveraging hybrid warfare.

In any case, Ukraine remains an active theater where Russia tests its hybrid war strategy and tactics. We would be well-served to observe Ukraine as a case study to help inform how those attacks may manifest against other targets, and what can be done to defend ourselves. There is nothing that Russia will do to the US and Europe that they are not willing to do to Ukraine.

If Russia intends to use Ukraine as a proving ground for its latest cyber weapons, then the United States and its allies should use Ukraine as a testing ground for innovative defenses and enhanced measurement. This cannot be accomplished from afar, nor can it be accomplished by either the public or private sector alone. Rather, both public and private sector organizations with an interest in honing defenses against cyberattacks should go to Ukraine, meet with the local experts, and observe how these conflicts play out in real life. This will be as beneficial for the fledgling democracy in Ukraine as it will be for Western nations intent on defending future elections.

Jared Cohen is the founder and CEO of Jigsaw at Alphabet Inc. He also serves as an adjunct senior fellow at the Council on Foreign Relations. Prior to Alphabet, he was Google's director of ideas and an advisor to Google's executive chairman. From 2006 to 2010 he served as a member of the Secretary of State's Policy Planning Staff under both Condoleezza Rice and Hillary Clinton. Mr. Cohen is the *New York Times* bestselling author of four books. His books include *The New Digital Age: Transforming Nations, Business, and our Lives,* which he co-authored with Eric Schmidt. His other books include *Children of Jihad, One Hundred Days of Silence: America and the Rwanda Genocide,* and the forthcoming *The Accidental Presidents.* He has been named to the "TIME 100" list, *Foreign Policy*'s "Top 100 Global Thinkers," and *Vanity Fair*'s "Next Establishment." Mr. Cohen is on several advisory boards, including Allianz, Stanford University's Freeman-Spogli Institute, Rivet Ventures, FluidMarket, ASAPP, RizviTraverse, and NCTC. He speaks fluent Swahili. Mr. Cohen received his B.A. from Stanford University and his M.Phil in International Relations from the University of Oxford, where he studied as a Rhodes Scholar.

Sources

Zach Beauchamp, "Meet the Shady Putin Crony Funding Russia's Troll Farm and Mercenary Army," *Vox*, Feb 26, 2018, www.vox.com/world/2018/2/26/17044930/yevgheny-prigozhin-putin-mueller-troll-farm.

United States Department of the Treasury, "Treasury Designates Individuals and Entities Involved in the Ongoing Conflict in Ukraine," June 20, 2017, www.treasury.gov/press-center/press-releases/Pages/sm0114.aspx.

Claire Wardel, "Fake news. It's complicated," First Draft, February 16, 2017, medium.com/1st-draft/fake-news-its-complicated-d0f773766c79.

Freedom House, *Freedom on the Net 2017: Manipulating Social Media to Undermine Democracy,* freedomhouse.org/report/freedom-net/freedom-net-2017.

Choe Sang-Hun, "Prosecutors Detail Attempt to Sway South Korean Election," *New York Times,* November 21, 2013, www.nytimes.com/2013/11/22/world/asia/prosecutors-detail-bid-to-sway-south-korean-election.html?mcubz=0.

Adrien Chen, "The Agency," *New York Times Magazine*, updated February 2018, www.nytimes.com/2018/02/17/world/europe/russians-indicted-mueller.html.

Meduza, "An ex St. Petersburg 'Troll' Speaks Out Russian Independent TV Network Interviews Former Troll at the Internet Research Agency," October 15, 2017, meduza.io/en/feature/2017/10/15/an-ex-st-petersburg-troll-speaks-out. The subject of the interview is a man who allegedly worked for the IRA from 2014-2015.

Ben Collins and Joseph Cox, "Jenna Abrams, Russia's Clown Troll Princess, Duped the Mainstream Media and the World," *Daily Beast*, November 2, 2017, www.thedailybeast.com/jenna-abrams-russias-clown-troll-princess-duped-the-mainstream-media-and-the-world.

Sam Thielman, "Russian Trolls Promoted Trump While Trashing Black Lives Matter On Twitter," Talking Points Memo, September 13, 2017, talkingpointsmemo.com/muckraker/russian-trolls-tea-party-news-twitter-account.

Maya Kosoff, "The Russian Troll Farm that Weaponized Facebook had American Boots on the Ground," *Vanity Fair,* October 18, 2017, www.vanityfair.com/news/2017/10/the-russian-troll-farm-that-weaponized-facebook-had-american-boots-on-the-ground.

Carly Nyst and Nick Monaco, *State-Sponsored Trolling: How Governments Are Deploying Disinformation as Part of Broader Digital Harassment Campaigns*, Institute of the Future, 2018, www.iftf.org/fileadmin/user_upload/images/DigIntel/IFTF_State_sponsored_trolling_report.pdf.

Don Kevin Hapal, "Oplan Cyber Tokhang on Facebook: 'Extrajudicial Reporting,'" Rappler, December 1, 2016, www.rappler.com/newsbreak/investigative/154099-oplan-cyber-tokhang-facebook-security.

Robert Chesney and Danielle Citron, "Deep Fakes: A Looming Crisis for National Security, Democracy and Privacy?" Lawfare, February 21, 2018, www.lawfareblog.com/deep-fakes-looming-crisis-national-security-democracy-and-privacy.

Nicola Smith, "North Korea May Have Made as Much as $200 Million from Bitcoin, According to Expert," *Telegraph*, March 5, 2018, www.telegraph.co.uk/news/2018/03/05/north-korea-may-have-made-much-200-million-bitcoin-according/.

"If the United States attains its potential improvements in innovation performance, China's great leap forward will likely, at best, be a few steps toward closing the innovation leadership gap that the United States currently enjoys."

–JOHN DEUTCH

Assessing and Responding to China's Innovation Initiative‡

John Deutch
Professor
MIT

Introduction

The purpose of this paper is to describe China's innovation initiative, assess the threat it poses to global US innovation primacy, and suggest possible policy responses. There are three main conclusions: (1) Chinese innovative progress is the inevitable result of its economic growth and technological maturity; market forces require the United States and other OECD countries to adjust the balance of goods and services they offer in response; (2) the United States should adopt targeted protectionist measures that compensate for unfair trade practices until the two countries, through negotiation, reach agreement on stable and mutually beneficial market access, cross-border investment, and technology transfer; and (3) the United States must improve its innovation capability in order to compete effectively in global markets.

China's Plan

In March 2016, China's National People's Congress ratified its thirteenth Five-Year Plan (FYP) 2016–2020,[1] based on the July 2015 *Made in China 2025* strategic blueprint to achieve global manufacturing leadership through innovation.[2] The goal is to upgrade industry writ large, but the plan highlights ten priority sectors: (1) new advanced information technology; (2) automated machine tools and robotics; (3) aerospace and aeronautical equipment; (4) maritime equipment and high-tech shipping; (5) modern rail transport equipment; (6) new-energy vehicles and equipment; (7) power equipment; (8) agricultural equipment; (9) new materials; and (10) biopharma and

‡ This paper is based in large part on a longer article, "Is Innovation China's New Great Leap Forward?," in the Summer 2018 *Issues in Science and Technology.*

advanced medical products. While other documents support *China2025*, the blueprint for artificial intelligence announced by the State Council is especially telling because of its detail.[3]

One cannot help but admire the *China2025* document. It sets clear goals based on explicit strategic priorities, support mechanisms, and key performance indicators. Technology strategy documents issued by the United States do not compare. For example, the June 2016 Obama White House–released *Impact Report: 100 Examples of President Obama's Leadership in Science, Technology, and Innovation* recounts many valuable actions but no overall strategy and no comprehensive plan to strengthen federally supported innovation activities across the board.[4] It is unimaginable that the Trump administration will put forward an innovation strategy for the country.

US commentators uniformly express concern about *China2025*: the Center for Strategic and International Studies notes the challenge presented to multinational corporations in the priority sectors.[5] The American Institute of Physics warns that the United States is at risk of losing global leadership to China.[6] A Council on Foreign Relations blog post is titled "Why Does Everyone Hate Made in China 2025?"[7] The US Chamber of Commerce subtitles its review "Global Ambitions Built on Local Protections," asserting that *Made in China 2025* "aims to leverage the power of the state to alter competitive dynamics in industries core to competitive dynamics."[8] The US government report from the US-China Economic and Security Review Commission has the provocative title *China's Technonationalism Toolbox: A Primer.*[9]

The unanimous message is that China is adopting a predatory economic policy that presents threats to the US economy and national security.

The Political Context

This is not the first time the United States has appeared to lose global technical and innovative leadership. The Soviet Union's 1957 launch of Sputnik raised the specter that the United States was losing the "space race." In the 1980s, there was widespread concern that the United States had fallen irrevocably behind Japan in manufacturing, in areas such as autos, flat panel displays, and consumer electronics.

The present situation with respect to China is more serious for three reasons. First, the Chinese increase in technical capability is impressive. China's rate of increase in economic activity is greater than the rate of the United States and other OECD countries, and its level already surpasses the United States on a purchasing power parity (PPP) basis. Second, concerns about innovation and related unfair trade practices must be seen in the broader context of strained US-China economic and political relations. Xi Jinping is seen as reversing China's slow but steady evolution to more democratic governance and reverting to a centralized communist party system with a single leader making all key decisions.[10] Third, Washington and Beijing are divided on many issues: (a) trade and tariff barriers, (b) China's actions in the South China Sea, (c) China's modernization of its military forces, (d) the status of Taiwan and Tibet, (e) the importance of North Korean sanctions, and (f) human rights violations. In this climate, differences related to innovation are amplified and more difficult to resolve.

National Security Concerns

A significant part of public concern about China's innovation initiative is the perception that the Chinese have an organized effort to use illegal means to acquire and utilize technology and acquire intellectual property. A corollary belief is that absent access to illegal means, the Chinese innovation initiative could not succeed. Here is a partial list of concerns:

- China's cyber theft of intellectual property and illegal technology transfer.
- Foreign firm's restricted access to China and the requirement that foreign firms operating in China transfer technical know-how to China, in violation WTO rules.
- Unfair trade practices, including subsidies to key high-technology priority industries and offering key products in export markets at prices below domestic Chinese prices and costs (dumping).
- Nontraditional collection of advanced technology that targets and exploits professors and students in US universities and research centers.

A number of US official reports address these threats. For example, the Office of the US Trade Representative issued a lengthy report in March 2018: *Findings of the Investigation into China's Acts, Policies, and Practices Related to Technology Transfer, Intellectual Property, and Innovation*.[11] The Chinese have made some efforts to address these matters, but as Harvard Professor Martin Feldstein notes, these intentions have not eased concerns.[12]

The National Bureau of Asian Research established The Commission on the Theft of American Intellectual Property, chaired by former Director of National Intelligence

Admiral Dennis Blair and Ambassador Jon Huntsman Jr. The commission issued a report in 2013 and an update in 2017 on the extent of US intellectual property loss resulting from a range of Chinese illicit activities, including use of cyber.[13]

In the *2018 Worldwide Threat Assessment of the US Intelligence Community,* Director of National Intelligence and former US Senator Dan Coats states: "China, for example, has acquired proprietary technology and early-stage ideas through cyber enabled means. At the same time, some actors use largely legitimate, legal transfers and relationships to gain access to research fields, experts, and key enabling industrial processes that could, over time, erode America's long-term competitive advantages."

At a February 13, 2018, Senate Intelligence hearing, FBI Director Chris Wray was much more explicit about the Chinese threat to US universities (and the over 50 Chinese Ministry of Education Confucius Institutes on US university campuses), indicating the possibility of placing restrictions on historically open US university campuses. On June 6, 2018, the Senate Judiciary Committee held a hearing on whether US universities were inadvertently helping China gain access to early-stage innovation.[14] On June 11, the US Department of State restricted the length of visas from five years to one year for Chinese graduate students planning to study aviation, robotics, and advanced manufacturing.[15] On June 19, 2018, a bipartisan group of twenty-six lawmakers organized by Senator Marco Rubio (R-FL) and Representative Jim Banks (R-IN), wrote a letter to US Department of Education Secretary Betsy DeVos seeking an investigation of the support that Huawei Technologies' innovation research program has provided to a number of leading US research universities, warning that this program poses a threat of illicit transfer of technology. The Department of Energy's National Nuclear Security Administration requested that MIT bar foreign graduate students from participating in a grant to its Plasma Science and Fusion Center.

The record clearly establishes extensive Chinese illicit technology transfer behavior. Anyone with US national security experience does not need to be convinced. The Intellectual Property (IP) Commission Report estimates that the annual cost of IP theft to the US economy is between $225 and $600 billion, with "China being the world's principal infringer."[16] What is striking is the implied judgment that this illicit behavior has been and will continue to be *decisive* in the advance of Chinese innovative capability. Few, if any, voices are raised to say that significant improvement in Chinese innovation should be expected with the growth in the Chinese economy and the increased maturity of its indigenous S&T infrastructure *without any* illicit behavior.

What Does Economics Tell Us?

Innovation refers to the process of implementing new or improved existing technology and management practice to offer products and services with desirable performance at affordable cost. Innovation encompasses the entire pathway from early-stage idea creation through technology development and demonstration and finally to late-stage production and deployment. Innovation aims at increasing consumer satisfaction, economic productivity, growth, exports, and jobs.

Economics informs us on three important aspects of innovation.

First, public support for early stage idea creation is justified because the long lead times and uncertainty about application means private firms cannot be confident of capturing future returns. The result is underinvestment in early-stage idea creation absent government support. This proposition is entirely valid: public investment in early-stage support is essential to assure a wide range of future technology options. However, support for early-stage R&D is not sufficient to assure successful innovation.

Second, there has long been a debate about the government's role in assisting private firms in crossing the "valley of death," the part of the innovation pathway that requires hundreds of millions of dollars for technology demonstration of commercial viability in such areas as advanced nuclear reactors, carbon capture and sequestration, and smart electricity generation systems. Narrow economic reasoning says "no" because the government is no more able to assess technical and cost risk of a new technology than the private sector. History indicates that from the Clinch River Breeder Reactor Project of the 1970s to the recent failed FutureGen and Kemper County CO_2 capture projects, government management of large technology demonstration projects is problematic. Progressive economic reasoning says "yes" because market failures, most notably the absence of a CO_2 emission charge, means private investors are making decisions based on private rather than social cost. Economics does not show the way to resolve the industrial policy dilemma, and there is no useful or agreed upon guidance on which type of support mechanisms—direct contracting or indirect incentives, such as loan guarantees and production payments—is most efficient for advancing technology demonstration.

Third, economic analysis seeks to establish the relationship between innovation productivity and economic growth. An enormous amount of work has been done over the years to elucidate the relationship between output "Q" of an entity "n" and time "t" and the capital "K," labor hours worked "L," and a surrogate variable "R" intended to represent the intangible intellectual capital that a firm possesses as a result of its internal spending on R&D and its exploitation of knowledge in the public domain resulting largely from government-sponsored R&D activities.

Equation 1 $$Q_n(t)=A_n(t)F_n(K_n(t),L_n(t),R_n(t))$$

The quantity A(t) in Equation 1 denotes the "total factor productivity" (TFP) that describes the influence of innovation on productivity. This relationship is relevant to understanding the trends in innovation, productivity, and output for individual firms, industry sectors, and nations. Comparisons are made between sectors "n" or over time "t." In the case where the production $F_n(t)$ is of the Cobb-Douglas form, changes in the TFP (with coefficients α, β, and γ) are given by

Equation 2 $$\frac{\Delta A_n(t)}{A_n}=\frac{\Delta Q_n(t)}{Q_n}-\alpha\frac{\Delta K_n(t)}{K_n}-\beta\frac{\Delta L_n(t)}{L_n}-\gamma\frac{\Delta R_n(t)}{R_n}$$

This is not the place to discuss all the difficulties encountered in carrying out the analysis for any application. There are many studies that use ingenious methods to circumvent the complex difficulties encountered to arrive at results, but I must admit I find most of the results to be less than convincing.

Here is an example that illustrates the complexity of isolating innovation factors that affect productivity. Only a few years ago, watching a movie at home meant a trip to Blockbuster to rent a videotape or disk, bringing it home, putting the tape or disk into a video or DVD player to project onto a TV or computer screen to watch the film, and afterward returning the item to the store. Today, Netflix streams movies from the internet to your home TV or computer screen. Suppose the cost and price of viewing a film was the same, so that in either the traditional or modern system the same number of films was viewed. There can be little doubt that all consumers would (and did) shift to Netflix because of the convenience, but there would be no change in TFP, since TFP does not include change in consumer satisfaction; nor does the analysis include the spillover benefit of the allocation of the time freed to the consumer for other productive use. It is difficult to track the benefits of many changes in information systems and technology.

Innovation occurs as the result of know-how, which is the accumulation of intellectual capital from the stream of proprietary R&D investment and/or federal

R&D results in the public domain. Conceptually, the innovation process follows this path:

$$\begin{pmatrix}\text{Increase}\\ \text{R\&D}\end{pmatrix}\rightarrow\begin{pmatrix}\text{Increased knowledge of}\\ \text{capital stock accumulation}\end{pmatrix}\rightarrow\begin{pmatrix}\text{Increased}\\ \text{productivity}\end{pmatrix}\rightarrow\begin{pmatrix}\text{Higher}\\ \text{growth}\end{pmatrix}$$

We tend to focus on the first step when, in fact, it is the second step that is least well understood.

In 2013, the US Bureau of Economic Analysis revised the national income and product accounts (NIPAs) to treat R&D as an investment rather than an expenditure. This action implies that firms capitalize "purchased" or "own account" R&D expenses, and this intangible capital depreciates at a significant rate, which varies widely between firms and sectors. Advocates of increased R&D spending rarely relate how such spending will translate into innovation. For example, in the widely acclaimed 2015 Paris Climate Accord, twenty countries agreed to double spending on clean energy R&D over five years, a total incremental expenditure of $20 billion for the purpose of accelerating the reduction of greenhouse gas emissions. No analysis was provided of the likely effect of this technology initiative on emission reduction or economic output because of the difficulty of carrying an analysis from national income accounts where clean energy firms are spread over hundreds of standard industry codes, including utilities, mining, transportation, and information.

Comparison of the trends in TFP between the United States and China are shown in the following three graphs.[17] Graph A shows the well-known fact that China's economic growth has been greater than that of the US and on a PPP basis now exceeds the US level. Graph B shows that China's TFP performance has been below that of the United States, indicating that China's growth has been fueled by capital expansion rather than innovation. Graph C, scaling each country to unity in 2011, shows that China's TFP is increasing more rapidly than that of the United States, suggesting the United States has a reason to be concerned by its relatively low rate of increase in innovation-driven TFP and by China's new strategy that replaces capital-driven growth by innovation-driven growth. The McKinsey Global Institute analyzes four different aspects of innovation: customer focused, efficiency driven, engineering based, and science based. McKinsey concludes "that China has the potential to meet its 'innovation imperative' and to emerge as a driving force in innovation globally, implying that China can meet its ambitious innovation drive goals."[18]

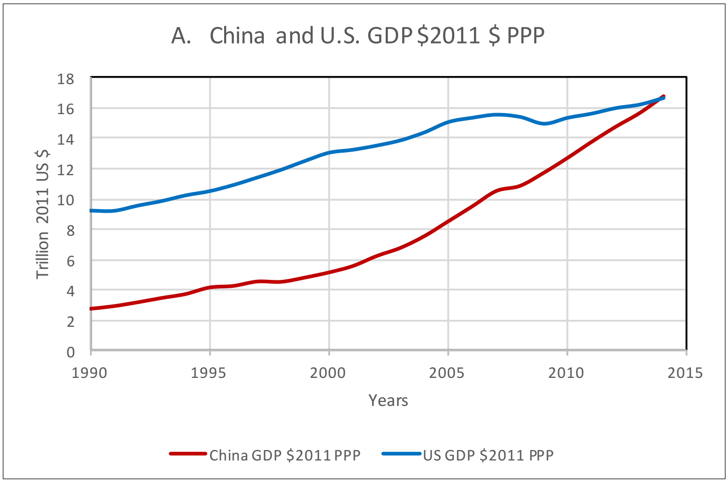
A. China and U.S. GDP $2011 $ PPP

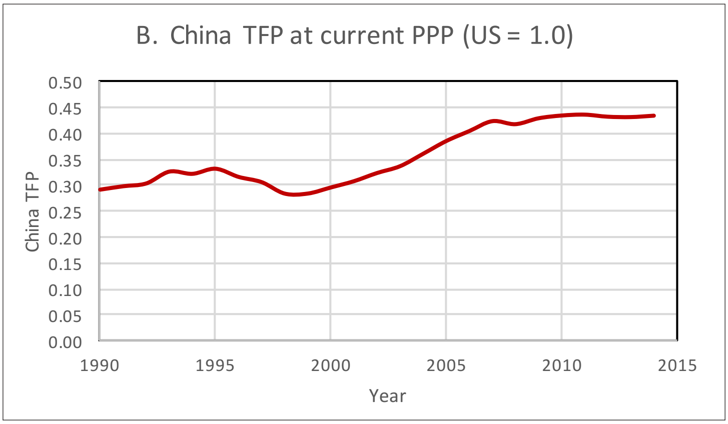
B. China TFP at current PPP (US = 1.0)

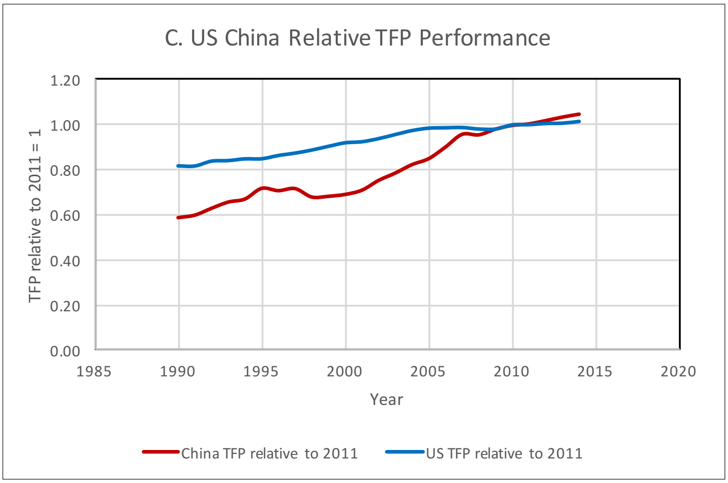
C. US China Relative TFP Performance

What Does the Data Say?

There are three sources for comparative country data that bear on innovation: the biannual US National Science Foundation's Science and Engineering Indicators,[19] the OECD R&D Statistics series,[20] and The World Bank Science & Technology Data.[21] These data share the shortcomings of uncertain data quality and reliance on purchasing power parity rather than market exchange rates to compare national efforts.[22] They bring to mind the caution of the late MIT economist Lester Thurow: "Never believe a number coming out of China."

Total R&D expenditures is the indicator most frequently cited.[23] China's total R&D expenditures are increasing more rapidly than those of the United States but have not yet reached the US level. As illustrated in Figure 1, if the United States has reason to be concerned with this trend, certainly the European Union and Japan should have even greater concern.

Figure 1

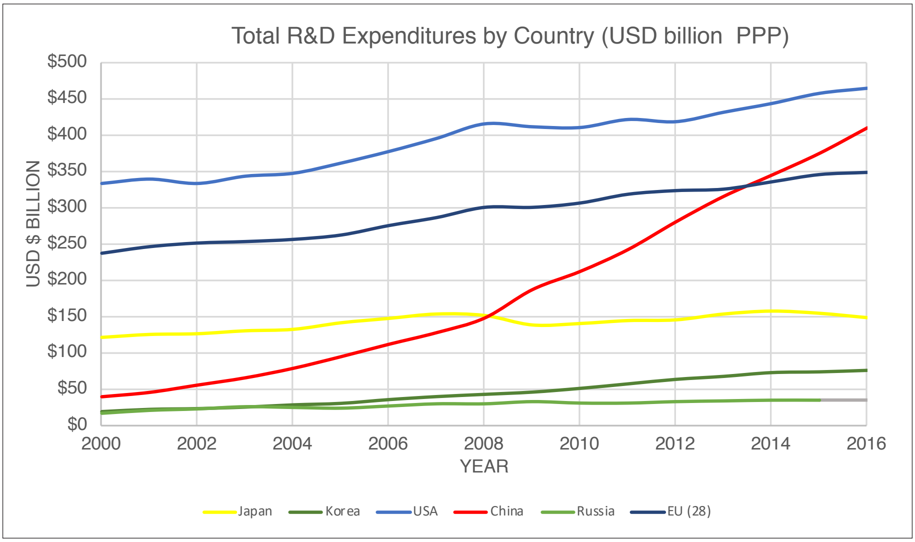

Other indicators of *inputs* to the innovation process are the size of the STEM workforce, expenditures on R&D plant and equipment, and the availability of venture capital. China's workforce exceeds that of the United States and its rate of increase is greater; however, the EU has a larger workforce than either country.[24] In 2014, China awarded about 1.65 million bachelor of science and engineering degrees, while the US total was about 742,000.

In 2015, the percentage of foreign workers in the US STEM workforce was approximately 30 percent,[25] with India and China the top two countries of origin at approximately 20 percent and 10 percent, respectively. There is widespread agreement that the US economy will need to rely on increasing numbers of foreign STEM workers, because the size of the US education pipeline is insufficient to meet anticipated demand. There is uncertainty about the trend in the net flow of foreign workers coming into the United States and foreign workers returning to their country of origin, as well as the implications for technology transfer of the reverse flow.

R&D plant and equipment[26] is another important *input* factor for innovation. Relatively little data or analysis is available to compare the inventory and the flow expenditures on different countries' R&D plant and equipment, such as laboratories and test facilities. It is likely that the R&D facilities inventory is higher in the United States and the EU than in China, but China's annual expenditure on R&D facilities is rising and may now exceed the US annual expenditures.

The balance of payment flows for intellectual property favors the United States by a factor of four, but payments to China are increasing.[27]

The preceding discussion addresses *inputs* to the innovation process. But, of course, the measurement of interest is *output* of the innovation process, in particular identifying the factors that influence process efficiency in different countries. Technical micro output indicators, such as patents, publication, citations, start-ups, and licensing agreements, do not tie to the innovation capacity of a firm, industry, or nation, and these micro indicators are even less useful for assessing the contribution that new technology or business practice makes to profitability or competitiveness.

Innovation Infrastructure

A country's innovation performance depends strongly on its underlying innovation infrastructure. This infrastructure has many components:

- the education of scientists and engineers who will enter the technical workforce;
- industry/university/government partnerships;
- laboratories and large-scale research facilities;
- standards for materials, products, safety, and subsystem interfaces;
- established patent, publication, and intellectual property rights;

- tax treatment for R&D activities;
- export controls on technology transfer and for participation of foreign scientists and engineers in the R&D enterprise; and
- access to venture capital.

The United States innovation infrastructure is the envy of the world, especially for early-stage R&D and for giving foreign students the opportunity to learn both technical and entrepreneurial aspects of innovation. China has taken a number of steps to improve its R&D infrastructure, but it still lags the United States, Europe, and Japan. The 2016 thirteenth five-year plan (2016-2020) includes increased R&D spending by China's Ministry of Science and Technology, its National Natural Science Foundation, and the Chinese Academy of Science to improve the science and engineering infrastructure base that fosters innovation. The United States also has a massive lead in private venture capital spending.[28]

Summing up, I believe China's comparative advantage will continue to be in manufacturing and its efficient supply chain. The US strength will continue to be its customer focus and developing new technologies for widespread application. The photovoltaic (PV) module experience is a helpful example. China has global leadership in low-cost manufacturing of PV modules based on conventional silicon solar cells. The impressive drop in module average sales price and accompanying demand growth is due to overcapacity and provincial rather than central government subsidies. Up to 2017, Chinese PV firms have not enjoyed profitability. In contrast, the United States maintains its lead in creating advanced PV technologies and developing production equipment and technology, which Chinese firms import and rely on. A joint US Department of Energy National Renewable Energy Laboratory and MIT study of the Chinese advantage in low-cost PV manufacturing states that the "price advantage of a China-based factory relative to a US-based factory is not driven by country-specific advantages, but instead by scale and supply-chain development."[29] This is likely to continue in the future, at least for the next decade or so, because of the relative strength and maturing of each country's innovation infrastructure. However, it is very unlikely that the Chinese will dominate the United States in key innovation areas such as artificial intelligence, robotics, machine learning, and CRISPR genetic editing.

Implications for US Policy

The United States should adopt four policies to respond to the innovation initiative at the heart of the *China2025* economic plan.

First, all US bilateral trade discussions with China should go beyond tariff negotiations to matters that directly affect innovation: market access, cross-border investment, and technology transfer.

Second, the United States must continue to monitor and expose illicit Chinese activities, track each documented incident, establish a process for confronting China with each case, and enforce US law by assessing penalties on violating firms. Present and former administrations have taken steps to protect the country from these illegal efforts, but much more needs to be done.[30]

Third, the United States must develop a new policy for engaging China on cross-border investments in high-technology firms and activities. This recommendation is motivated by two realities. First, the United States and China have very different objectives for cross-border investment. US firms are primarily interested in offering goods and services to the large and growing Chinese domestic market but fear the Chinese practice of hijacking technology in order to establish competitive indigenous capability. Chinese investments in the United States are primarily for access to high-technology–creating firms and to the advanced technologies outlined in *China2025*. Increasingly, these investments will be in start-up companies and in joint ventures that are creating key technologies for future innovation.

The second reality is that many of the key technologies, notably artificial intelligence and robotics, are inherently dual use, with important applications in both the commercial and national security sectors. The United States has long had an interagency process, the Committee on Foreign Investment in the United States (CFIUS), that reviews transactions that could result in foreign entity control to determine the effect of such transactions on US national security.[31] The original CFIUS mandate required it to focus "solely on any genuine national security concerns by a covered transaction, not on other national interests."[32] But over time CFIUS has been pressed to examine transactions that are perceived to affect US economic competitiveness, including foreign transactions that involve "critical technologies."[33] CFIUS is ill-equipped to assess the implications of start-ups and joint ventures in rapidly evolving key dual-use technologies of uncertain future application or success. If the United States wishes to reduce the ease with which China (and possibly other countries) can acquire US advanced technologies to fuel its innovation initiative, it is necessary to adopt controls that restrict access.

Following is a starting set of controls to consider for Chinese investment in US enterprises—start-up companies, joint ventures, and venture capital funds—in a defined set of key advanced technologies:

1. Require all investments to be registered with one of the CFIUS agencies.
2. Prohibit 100 percent interest in US advanced technology enterprises.
3. Require that an enterprise that has greater than a majority interest by a Chinese entity create an independent "security supervisory board" similar to those sometimes required by CFIUS to monitor and report all offshore technology transfer.

Federal agencies that fund technology development would be permitted to award contracts or grants only to Chinese enterprises that had operating subsidiaries in the United States.

China would not be barred from supporting research activities in US universities, provided that research results were publicly available. Chinese firms supporting research on US campuses would not be permitted preferential or exclusive licenses to any intellectual property produced.

Much effort is required to define precisely each of these suggested measures and a supporting administrative structure. The justification for such a set of protectionist measures is to confront the Chinese innovation initiative whose announced intention is to dominate world markets and whose progress depends to a significant degree on illicit technology transfer.

Fourth, the United States should not place restrictions on US universities and research centers to slow the leakage of technology to China in order to maintain US competitiveness.[34] Restrictive proposals include pre-publication clearance of research results supported by the US Department of Defense, applying the classification category of "sensitive but unclassified" on some government-sponsored research, and restrictions on foreign graduate students joining "sensitive" research projects and presenting research at international meetings. Each of these measures conflicts with the open structure of admission, research, and publication that keeps the US innovative ecosystem fresh, exciting, and agile.

Several reasons support this recommendation: (a) government agencies are unlikely to balance properly the effectiveness of a proposed restrictive measure with the adverse impact on admissions, (b) university faculty and administrators are ill-equipped to administer such restrictions, (c) a move toward US university restriction will inevitably slow the flow of students from China and other countries, and they are needed by US industry. If the federal government imposes restrictions on research it sponsors, it will weaken its link with the universities that have been so central to US innovation. The risk of loss is minor compared to the losses that will be incurred by restricting inquiry on university campuses.

It is futile to maintain US competitiveness and its lead in early-stage innovation by trying to keep others out or, for that matter, our ideas in. The only effective response to China's growing capability is to master the new intellectual frontiers and continue to recruit the most talented workforce able to rapidly translate new ideas into practice. In this regard, the United States should continue to welcome Chinese and other science and engineering graduates to US universities and liberalize immigration green card requirements to assure adequate supply for US industry.[35]

Fifth, in order to maintain its current relative position to China, the United States must dramatically increase its innovation effort, especially in manufacturing, to bring the key future technologies to market. The United States should not adopt, as China has, a single national strategy. Nor should the United States rely simply on increasing federal R&D support from traditional agencies. US innovative activity is tremendously dynamic due to individual entities applying new technical applications in unique ways. This dynamic process, distinctive to the United States, is possible because of its formidable innovation infrastructure, described above, and strong tradition of customer and application focus. The approach depends on the open and free character of US society; it is difficult to imagine such productive vitality existing in communist China, where freedom of expression and association is restricted.

The three prongs of US innovation—the federal government, industry, and academic research centers—need to follow distinct pathways to achieve a higher level of national innovation. Individual industry sectors have the most important role in bringing new technology and business practice to market. Industry associations need to convince their members of the urgency of increasing the pace of innovation and provide members with case studies of successful unconventional new innovation. Industry associations should launch efforts to spread best practices and selectively undertake public-private technology partnerships.

Universities and research centers have two important roles: first, to increase the flow of researchers with the motivation and experience to achieve innovations and, second, to greatly expand work on key technologies. The country needs a workforce able to contribute across the entire innovation pathway, including craftsmen, designers, and PhD researchers.

The federal government's role is to enable enhanced innovation in several ways: R&D support should prioritize work that stresses innovation (without endangering early-stage fundamental research) and should explore different mechanisms for providing support, such as the Department of Energy's ARPA-E program. Management of technology demonstration projects is a critical aspect of federal support for innovation. The Office of Management and Budget should undertake a thorough review to identify regulations that slow innovation in the areas of patents, security registration, tax provisions, and federal acquisition regulations and recommend changes that streamline the innovation process. The president should form an interagency council charged with overseeing the innovation efforts of different agencies and sharing best practices; the interagency council should also study and track the implications of an increasingly digital-based economy on the future of work and changing educational needs. The federal government must also continue to combat illegal theft and hacking of technology and know-how by China and other countries, as outlined above.

Finally, Congress should establish a national commission composed of political, industry, university, and public interest groups to communicate the nation's need for advancing innovation and to report progress. The Unites States' approach to improving its innovation capability is based on the core strength of its innovation infrastructure and the diverse and dynamic entrepreneurial enterprise incentivized to succeed. All elements of the public have a role to play in the process but should share a high-level vision of the importance of the task. If the United States attains its potential improvements in innovation performance, China's great leap forward will likely, at best, be a few steps toward closing the innovation leadership gap that the United States currently enjoys.

John Deutch is an institute professor at the Massachusetts Institute of Technology. Mr. Deutch has been a member of the MIT faculty since 1970 and has served as chairman of the Department of Chemistry, dean of science, and provost. Mr. Deutch has published over 160 technical publications in physical chemistry, as well as numerous publications on technology, energy, international security, and public policy issues. He served as director of central intelligence from May 1995-December 1996. From 1994-1995, he served as deputy secretary of defense and served as undersecretary of defense for acquisition and technology from 1993-1994. He has also served as Director of Energy Research (1977-1979), acting assistant secretary for energy technology (1979), and undersecretary (1979-80) in the United States Department of Energy. He is a member of the Aspen Strategy Group.

[1] *The 13th Five-Year Plan for Economic and Social Development of the People's Republic of China 2016–2020* is available at en.ndrc.gov.cn/newsrelease/201612/P020161207645765233498.pdf. The US government summary of China's 13th five-year plan is the February 2017 staff research report, *The 13th Five-Year Plan* by Katherine Koleski, US-China Economic and Security Review Commission; it is available at www.uscc.gov.

[2] English translation of Made in China 2025 by IoT ONE is available at community.iotone.com/t/report-made-in-china-2025-the-10-year-industrial-iot-roadmap-in-china/109.

[3] The news release "China issues guideline on artificial intelligence development" is available at english.gov.cn/policies/latest_releases/2017/07/20/content_281475742458322.htm. Paul Triolo, Elsa Kania, and Graham Webster, "Translation: Chinese government outlines AI ambitions through 2020," New American Foundation (blog), www.newamerica.org/cybersecurity-initiative/digichina/blog/translation-chinese-government-outlines-ai-ambitions-through-2020/.

[4] Available at obamawhitehouse.archives.gov/the-press-office/2016/06/21/impact-report-100-examples-president-obamas-leadership-science.

[5] Scott Kennedy, "Made in China 2025," Center for Strategic and International Studies, June 1, 2015, csis.org/analysis/made-china-2025.

[6] American Institute of Physics, "Biennial Report Shows US at Risk of Losing Global R&D Leadership to China," January 23, 2018, aip.org/fyi/2018/biennial-report-shows-us-risk-losing-global-rd-leadership-china.

[7] Lorand Laskai, "Why Does Everyone Hate Made in China 2025?" Council on Foreign Relations (blog), March 28, 2018, www.cfr.org/blog/why-does-everyone-hate-made-china-2025.

[8] US Chamber of Commerce, *Made in China: Global Ambitions Built on Local Protections,* 2017, www.uschamber.com/report/made-china-2025-global-ambitions-built-local-protections-0.

[9] Katherine Koleski and Nargiza Salidjanova, *China's Technonationalism Toolbox: A Primer*, US-China Economic Security Review Commission, March 2018, www.uscc.gov/sites/default/files/Research/China%27s%20Technonationalism.pdf.

[10] The April 2018 issue of the *Journal of Democracy* has a set of eight articles on "China in Xi's new era." Available at www.ned.org/journal-of-democracy-april-2018-issue-china-in-xis-new-era/.

[11] Office of the US Trade Representative, *Findings of the Investigation Into China's Acts, Policies, and Practices Related to Technology Transfer, Intellectual Property, and Innovation Under Section 301 of the Trade Act of 1974*, March 22, 2018, ustr.gov/sites/default/files/Section%20301%20FINAL.PDF.

[12] Martin Feldstein, "How to Make Trade Peace with China," *Wall Street Journal*, April 5, 2018.

[13] The Commission on the Theft of American Intellectual Property, *The Theft of American Intellectual Property: Reassessments of the Challenge and United States Policy*, 2017, www.ipcommission.org/report/.

[14] Senate Committee on the Judiciary, Subcommittee on Border Security and Immigration hearing, "Student Visa Integrity: Protecting Educational Opportunity and National Security," June 6, 2018, www.judiciary.senate.gov/meetings/a-thousand-talents-chinas-campaign-to-infiltrate-and-exploit-us-academia.

[15] Jeffrey Mervis, "More Restrictive US Policy on Chinese Graduate Student Visas Raises Alarm," *Science*, June 11, 2018, www.sciencemag.org/news/2018/06/more-restrictive-us-policy-chinese-graduate-student-visas-raises-alarm.

[16] The Commission on the Theft of American Intellectual Property, *The Theft of American Intellectual Property: Reassessments of the Challenge and United States Policy*, 2017, p. 13, www.ipcommission.org/report/.

[17] Robert C. Feenstra, Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," *American Economic Review* 105, no. 10 (2015): 3150-82, www.ggdc.net/pwt.

[18] McKinsey & Company, "The China Effect on Global Innovation," mckinseychina.com/the-china-effect-on-global-innovation/.

[19] National Science Board, *Science & Engineering Indicators 2018*, www.nsf.gov/statistics/2018/nsb20181/.

[20] OECD, Research and Development Statistics, www.oecd.org/innovation/inno/researchanddevelopment statisticsrds.htm.

[21] The World Bank, Science and Technology Data, data.worldbank.org/topic/science-and-technology?view=chart.

[22] Sean M. Dougherty, Robert Inklaar, Robert H. McGuckin, and Bart van Ark, *International Comparisons of R&D Expenditure: Does an R&D PPP Make a Difference?* National Bureau of Economic Research, Working Paper #12829, January 2007, www.nber.org/papers/w12829. This paper concluded that "the use of an R&D PPP will yield comparative costs and R&D intensities that vary substantially from the current practice of using GDP PPPs, likely increasing the real R&D performance of the comparison countries relative to the United States." The issue, and what if anything to do about it, remains unresolved.

[23] OECD, Gross Domestic Spending on R&D, data.oecd.org/rd/gross-domestic-spending-on-r-d.htm.

[24] National Science Board, "Global S&E Labor Force," in *Science & Engineering Indicators 2018*, www.nsf.gov/statistics/2018/nsb20181/report/sections/science-and-engineering-labor-force/global-s-e-labor-force.

[25] National Science Board, *Science & Engineering Indicators 2018*, Tables, www.nsf.gov/statistics/2018/nsb20181/data/tables; American Immigration Council, "Foreign-born STEM Workers in the United States," June 14, 2017, www.americanimmigrationcouncil.org/research/foreign-born-stem-workers-united-states.

[26] Michael Yamaner, "Federal Research and Development and R&D Plant Funding Drop by 9% in FY 2013," National Science Foundation, May 14, 2015, www.nsf.gov/statistics/2015/nsf15322/.

[27] Asia Society Policy Institute, "Innovation," Winter 2018, chinadashboard.asiasociety.org/winter-2018/page/innovation.

[28] See Figure 8-22 of National Science Board, "Invention, Knowledge Transfer, and Innovation," in *Science & Engineering Indicators 2018*, www.nsf.gov/statistics/2018/nsb20181/report/sections/invention-knowledge-transfer-and-innovation/innovation-indicators-united-states-and-other-major-economies#venture-capital.

[29] Alan C. Goodrich, Douglas M. Powell, Ted L. James, Michael Woodhousea, and Tonio Buonassisi, "Assessing the Drivers of Regional Trends in Solar Photovoltaic Manufacturing," *Energy and Environmental Science* 6, (2013): 2811-21.

[30] The Commission on the Theft of American Intellectual Property, *The Theft of American Intellectual Property: Reassessments of the Challenge and United States Policy*, 2017, Appendix, p. 16, www.ipcommission.org/report/.

[31] The Committee on Foreign Investment in the United States (CFIUS), US Department of the Treasury, www.treasury.gov/resource-center/international/Pages/Committee-on-Foreign-Investment-in-US.aspx.

[32] Fed. Reg. 74567 (Dec. 8, 2008).

[33] Section 721(m)(3) of the Defense Production Act, as amended, 50 USC. App. 2170.

[34] John Deutch and Condoleezza Rice, "Maintaining America's Lead in Creating and Applying New Technology," in *The World Turned Upside Down: Maintaining American Leadership in a Dangerous Age* (Washington, DC: Aspen Strategy Group, 2017), 117-21.

[35] The Commission on the Theft of American Intellectual Property, *The Theft of American Intellectual Property: Reassessments of the Challenge and United States Policy,* 2017, p. 5, www.ipcommission.org/report/.

"The creation of a triangular relationship between government, industry, and academia was, in its own way, one of the significant innovations that helped produce the technological revolution of the late 20th century."

–WALTER ISAACSON

The Sources of America's Innovative Edge

Walter Isaacson
Professor of History
Tulane University

For the past fifty years, the rational exuberance of the American economy has been propelled by the combination of three innovations: the computer, the microchip, and the internet.

The research and development that produced each came from a triangular alliance of government, academia, and private business. The first computers were funded by the military, built at the University of Pennsylvania and Harvard, then commercialized by companies such as Univac and IBM. Transistors were invented at Bell Labs, then federal funding for the space and strategic missile programs led private companies such as Fairchild and Intel to devise ways to etch thousands of them onto small silicon chips. And the internet was famously conceived by DARPA and built by research universities working with private contractors such as BBN.

This tripartite machine of government working with universities and private corporations was not merely a random array with each group pursuing its own aims. Instead, during and after World War II, the three groups had been purposely fused together into an innovation triangle.

The person most responsible for forging this assemblage was Vannevar Bush, an MIT professor who in 1931 had built an early analog computer. Bush was well-suited to this task because he was a star in all three camps: dean of the MIT School of Engineering, a founder of the electronics company Raytheon, and America's top military science administrator during World War II.

He was passionate about elevating the role of science and engineering in society at a time—the mid-1930s—when not much exciting seemed to be happening in either field. The most notable new inventions put into the time capsule at the New York 1939 World's Fair were a Mickey Mouse watch and a Gillette safety razor. The advent of World War II would change that, producing an outpouring of new technologies with Bush leading the way.

Worried that America's military was lagging in technology, he mobilized Harvard President James Bryant Conant and other scientific leaders to convince President Roosevelt to form the National Defense Research Committee and then the military's Office of Scientific Research and Development, both of which he headed. With an ever-present pipe in his mouth and a pencil in his hand, he oversaw the Manhattan Project to build the atom bomb as well as projects to develop radar and air-defense systems.

When the war ended, Bush produced a report in July 1945 at the behest of President Roosevelt (which ended up being delivered to President Truman) that advocated government funding of basic research in partnership with universities and industry. Bush chose an evocative and quintessentially American title for his report: "Science, The Endless Frontier." His introduction deserves to be reread whenever politicians threaten to defund the research needed for future innovation. "Basic research leads to new knowledge," Bush wrote. "It provides scientific capital. It creates the fund from which the practical applications of knowledge must be drawn."

The war, Bush wrote, had made it "clear beyond all doubt" that basic science—discovering the fundamentals of nuclear physics, lasers, semiconducting materials, computer science, radar—"is absolutely essential to national security." It was also, he added, crucial for America's economic security. "New products and new processes do not appear full-grown. They are founded on new principles and new conceptions, which in turn are painstakingly developed by research in the purest realms of science. A nation which depends upon others for its new basic scientific knowledge will be slow in its industrial progress and weak in its competitive position in world trade."

Bush's description of how basic research provided the seed corn for practical inventions became known as the "linear model of innovation." Based on this report, Congress established the National Science Foundation.

Most important, government spending was not funneled into government-run labs, as had happened with the Manhattan Project. Instead, government research funding went to universities and private contractors. "No American had greater influence in the growth of science and technology than Vannevar Bush," MIT President Jerome Wiesner proclaimed, adding that his "most significant innovation was the plan by which, instead of building large government laboratories, contracts were made with universities and industrial laboratories."

The creation of a triangular relationship between government, industry, and academia was, in its own way, one of the significant innovations that helped produce the technological revolution of the late twentieth century. The Department of Defense soon became the prime funder of much of America's basic research. By 1965, 23 percent of the federal government's funding for university science came from the Pentagon—almost twice as much as from the National Science Foundation. The return on that investment was huge, leading not only to the internet, but to many of the pillars of America's postwar innovation and economic boom.

A few corporate research centers, most notably Bell Labs, existed before the war. Bell Labs brought together theoreticians, materials scientists, metallurgists, engineers, and even telephone-pole climbers. Bell Labs showed how sustained innovation could occur when people with a variety of talents were brought together, preferably in close physical proximity where they could have frequent meetings and serendipitous encounters.

After Bush's clarion call produced government contracts, other corporate research centers began to proliferate. Xerox created the Palo Alto Research Center, known as Xerox PARC, that had as one of its leaders Bob Taylor, who had helped create the internet while running DARPA's Information Processing Techniques Office. Xerox PARC developed the graphical user interface now used on personal computers, the ethernet, and dozens of other innovations that became part of the digital revolution.

In addition, hybrid labs combining government, academia, and industry were launched. Among the most notable were the RAND Corporation, originally formed to provide research and development (hence the name) to the Air Force, and Stanford Research Institute (SRI).

Many of America's most important new private corporations were spawned by the three-way relationship of Bush's innovation triangle.

Take Google, for example. Larry Page's father was a professor of computer science and artificial intelligence at Michigan State, the recipient of large federal research grants. Sergey Brin's parents were refugees from Russia who received visas to come to the US. His father became a math professor at the University of Maryland, where the Department of Defense funded ways to calculate missile trajectories, and his mother became a researcher at the nearby NASA Goddard Space Flight Center.

Both Larry and Sergey ended up at Stanford as graduate students in a government-funded program called the Digital Libraries Initiative. The money came from the National Science Foundation and a consortium of other federal agencies. With their tuition paid by this program, they came up with systems called BackRub and PageRank that indexed the World Wide Web. Thus was Google born.

Another great innovation spurring the American economy and competitiveness was likewise funded by the federal government through universities and corporate labs. Beginning with the presidency of George H.W. Bush, the Human Genome Project sequenced DNA and launched a revolution in biomedicine that will produce the most important innovations and discoveries of the twenty-first century. "Through it all, the federal government invested heavily in basic research," says Eric Lander, one of the leaders of the genome project. "That policy made American universities engines of discovery that attracted the best talent to our shores and sparked the world's most innovative companies."

The question now, Lander says, is "whether America will yield its position as the world's leader in science and technology. For the first time since World War II, our primacy is in jeopardy."

A recent report from the Atlantic Council echoed Vannevar Bush's phrasing when it called such examples of federally funded basic research at university and corporate labs "the nation's scientific seed corn, enabling basic, pre-competitive R&D that will mature into harvestable technologies in the future." However, the report noted, "federal R&D spending has shrunk significantly over the last few decades; once the world leader, the United States now ranks twelfth in government-funded R&D spending as a percentage of GDP." Federal R&D spending has declined from about 1.2 percent of GDP in 1976 to less than 0.8 percent in 2016. This is the lowest level since the pre-Sputnik era, and in the Aspen Strategy Group, there may still be a few people who know what the pre-Sputnik era was.

Some of this decrease in federal funding has been replaced by an increase in corporate research, especially in sectors such as the pharmaceutical industry, where it is clear that research can lead directly to valuable products. In the 1960s, around 70 percent of total R&D was federally funded, with 30 percent coming from the private sector. Now those figures are reversed.

Corporate funding tends to be more focused on products. As the balance has shifted away from government funding at university research labs, there has been a reduction in basic scientific research that is aimed at creating the fundamental theoretical knowledge that can produce the seed corn that will eventually lead to great innovations.

This decline in scientific investment in basic research and university labs is not a partisan phenomenon or a product of the Trump administration. For almost twenty-five years, federal funding for university research and state funding for higher education has been in decline. Between 2011 and 2015, during the Obama administration, federal investment in university research declined by 13 percent.

But it's now getting even worse. In the latest proposed budgets from House Republicans and the Trump administration, science and technology research federal funding would be cut by an additional 15 percent. And in the 5,000 or so tweets by President Trump since he came to office, the words "science" and "technology" have never appeared.

In addition, despite the launch of some corporate labs such as GoogleX, private corporations have largely dismantled the research institutes like Bell Labs and Xerox PARC, partly in the face of challenges from short-term investors who demand a shorter time horizon in returns on investment.

The potential economic and security ramifications can be foreshadowed by looking at the opposite approach now being taken by China, which is heavily funding basic scientific research, including in vital fields such as artificial intelligence (AI) and genetic engineering.

Take the AI sector, for example. In its thirteenth Five-Year Plan released in 2016, China's leadership announced its ambition to transform China into a "nation of innovation" by launching fifteen "Science and Technology Innovation 2030 Megaprojects." These included big data, intelligent manufacturing, and robotics. It was a steroid-charged version of Bush's 1945 paper urging America to combine federal dollars with university and corporate labs. A year ago, in May 2017, China added "Artificial Intelligence 2.0" as the sixteenth mega-project.

The goal of this project is audacious yet simple: to make China the world leader in AI by 2030. Combining government dollars with corporate and academic initiatives, China is now building an ecosystem that would transcend even Bush's wildest dreams.

The local government of Tianjin, a city two hours from Beijing, is raising a $5 billion fund to support AI development, and the central government is building a $2.1 billion AI technology park in Beijing's western suburbs.

Guided by the government's vision, money is also flowing into the Chinese private sector. Venture funds and other private funds invested $4.5 billion into more than 200 Chinese AI companies between 2012 and 2017, according to Kai-Fu Lee, a former Google and Microsoft executive who now leads a venture capital firm, Sinovation Ventures. The AI start-up SenseTime raised $600 million in a deal led by Alibaba, giving SenseTime an implied valuation of more than $3 billion. CB Insights reports that, by certain types of measurement, China has overtaken the US in the funding of AI start-ups. For example, China accounted for 48 percent of the world's AI start-up funding in 2017, compared to 38 percent for the US.

The funding and investments are already paying off. China's students and programmers are now routinely winning international competitions in AI and machine learning. Baidu is at the forefront of AI, with 2,000 researchers, including in offices in Silicon Valley and Seattle. It now rivals Google as a global leader in AI research and boasts the most accurate and powerful program for speech recognition. According to the White House's National Artificial Intelligence Research and Development Strategic Plan, in AI research, China has surpassed the US in the number of journal articles that mention "deep learning" or "deep neural network." At the annual Association for the Advancement of Artificial Intelligence conference, the percentage of Chinese authors of AI research papers presented grew from 10 percent to 23 percent between 2012 and 2017, while the percentage of US authors declined from 41 percent to 34 percent.

China has one other advantage that the US should not envy. It has fewer restrictions on data collection and less compunction about violating personal privacy. This is unnerving, but it is also an advantage because big data will fuel many AI advances. China sits on a growing reservoir of big data, making it, as *The Economist* put it, "the Saudi Arabia of data."

China's version of Vannevar Bush's innovation triangle is a "military-civil fusion" that encourages the collection of data on citizens. It uses facial recognition technology for domestic surveillance. In Shenzhen, for example, there are cameras on poles with reminders saying, "Jaywalkers will be captured using facial-recognition technology." Cross the street improperly, and your face and name are likely to be displayed publicly on a nearby screen and put into a database.

Enter a search query into Google, and the company may gather the data to improve its algorithm and market products to you; make the same search on Baidu, and your data also goes into a government-controlled database. The same data collection policies apply every time someone in China uses a WeChat wallet, shops online on Taobao, or hails a ride with Didi. Baidu uses facial recognition of its employees to open the security gates in its lobby, and the technology allows customers at Kentucky Fried Chicken to authorize a payment via facial scan. The technology is also used to recognize passengers at airport security gates. When US Customs and Border Protection last year floated a plan to do the same to verify the identity of people boarding certain flights in the US, a controversy erupted.

As these examples show, there are elements of China's technology and innovation initiatives that the US will not wish to emulate. That is true of facial recognition, AI, and big data, and it is also true in gene editing, cloning, and other types of biotechnology where China has fewer ethical and policy restrictions.

But the political and ethical restrictions in the US make it even more important for the US to stay ahead of China in other ways, most notably funding basic research into science and investing in university and corporate labs.

A good place to start would be revitalizing our investments in research universities, now being decimated by cuts and other challenges. The US has thirty-two of the top fifty universities in the world, magnets for the world's best students. "But America seems increasingly unwelcoming to foreign students, whose applications this year have fallen by as much as 30 percent in some programs," says Lander, who was a leader of President Obama's President's Council of Advisors on Science and Technology. "Will the next generation of entrepreneurs and leaders from around the world study elsewhere? At the same time, the Trump administration has proposed slashing funding for basic research. Congress came within a hair's breadth of taxing graduate student fellowships and did impose a tax on university endowments, which help fund costs not covered by tuition."

Reversing such policies is the critical first step to creating, once again, the research breakthroughs that will lead to future innovations, rather than continuing on America's new path of destroying our seed corn before the next harvest.

Walter Isaacson is a professor of history at Tulane and an advisory partner at Perella Weinberg, a financial services firm. He is the past CEO of the Aspen Institute where he is now a distinguished fellow and has been the chairman of CNN and the editor of TIME magazine. Mr. Isaacson's most recent biography, *Leonardo da Vinci* (2017), offers new discoveries about Leonardo's life and work, weaving a narrative that connects his art to his science. He is also the author of *The Innovators: How a Group of Hackers, Geniuses, and Geeks Created the Digital Revolution* (2014), *Steve Jobs* (2011), *Einstein: His Life and Universe* (2007), *Benjamin Franklin: An American Life* (2003), and *Kissinger: A Biography* (1992), and coauthor of *The Wise Men: Six Friends and the World They Made* (1986). He is chair emeritus of Teach for America. From 2005-2007 he was the vice-chair of the Louisiana Recovery Authority, which oversaw the rebuilding after Hurricane Katrina. He was appointed by President Barack Obama to serve as the chairman of the Broadcasting Board of Governors, a position he held from 2009 to 2012. He serves on the board of United Airlines, the New Orleans City Planning Commission, the New Orleans Tricentennial Commission, Bloomberg Philanthropies, the Rockefeller Foundation, the Society of American Historians, the US Defense Department Innovation Board, the Carnegie Institution for Science, and My Brother's Keeper Alliance. Mr. Isaacson is a graduate of Harvard College and of Pembroke College of Oxford University, where he was a Rhodes Scholar.