IBM Center for The Business of Government

Risk Management in the AI Era:

Navigating the Opportunities and Challenges of AI Tools in the Public Sector

Justin B. Bullock
Texas A&M University

Matthew M. Young
Syracuse University

IBM Center for
The Business
of Government
2020

Risk Management in the AI Era: Navigating the Opportunities and Challenges of AI Tools in the Public Sector

Justin B. Bullock
Texas A&M University

Matthew M. Young
Syracuse University

IBM Center for
The Business
of Government

2020

TABLE OF CONTENTS

FOREWORD

On behalf of the IBM Center for The Business of Government, we are pleased to release this report, *Risk Management in the AI Era: Navigating the Opportunities and Challenges of AI Tools in the Public Sector,* by Justin Bullock and Matthew Young.

Artificial Intelligence (AI) has moved into the mainstream of businesses and governments. Business leaders are rushing to take advantage of the benefits that can be brought to a wide array of industries to help increase productivity. Government leaders are also moving forward, but with appropriate caution. When considering the use and application of AI related technologies, government leaders weigh different factors than their private sector counterparts. These leaders must address key issues of accountability, transparency, ethics, equity, common good, effectiveness, efficiency, managerial capacity, and political legitimacy.

DANIEL J. CHENOK

The report authors put forth a threefold strategy to assist government leaders and public managers with how best to approach using AI.

First, the report reviews prior federal government studies on the use and application of AI. These reports reflect a number of important issues for agencies and stakeholders to consider as they begin incorporating AI; the studies also highlight the government's broad risk management approach to AI.

ZBYNEK KROBOT

Second, the report provides a risk management framework for when and how government can and should consider using AI tools, how to use these tools, and which organizational tasks and decisions may benefit from the use of AI tools. This section also discusses the key issues identified above.

Third, the report presents case studies of two innovative uses of AI tools to help manage risks from local governments: the City of Syracuse, New York, and the City of Bryan, Texas.

The authors close with a list of practical guidelines for government action in using AI tools to improve the overall quality of governance, while incorporating similar tools into their overall risk management strategy.

This report serves as an excellent companion piece to recent IBM Center reports that examine both AI and aspects of risk management that can help government agencies. The Center has recently collaborated with the Partnership for Public Service on several reports, which includes *Using Artificial Intelligence to Transform Government*; *More Than Meets AI*; and *More Than Meets AI: Part II*. These reports focus on the evolving use of AI in government, how this technology might affect federal employees, and how best to identify the risk associated with pursuing AI technologies. The report series also draws on lessons from companies and countries around the world that use AI. These organizations have identified and are addressing AI issues that include bias, security, transparency and job impact, and their insights can be instructive for federal agencies. A forthcoming Center report will focus on risk management in a law enforcement and judicial agencies.

Along with these AI-specific reports, the Center has a library of risk management research, which includes *Managing Risk in Government: An Introduction to Enterprise Risk Management* by Karen Hardy; *Managing Risk, Improving Results: Lessons for Improving Government Management from GAO's High Risk List* by Donald Kettl; *Improving Government Decision Making through Enterprise Risk Management* by Thomas Stanton and Douglas Webster; *Risk Management and Reducing Improper Payments: A Case Study of the U.S. Department of Labor* by Dr.Robert Greer and Justin B.Bullock; and *Managing Cybersecurity Risk in Government* by Anupam Kumar, James Haddow, and Rajni Goel.

We hope that this report provides a useful model for government agencies to adapt in managing AI risks.

Daniel J. Chenok
Executive Director
IBM Center for The Business of Government
chenokd@us.ibm.com

Zbynek Krobot
Partner, Digital, Shared Services, Strategy & Performance Analytics
IBM Global Business Services
zbynek.krobot@us.ibm.com

EXECUTIVE SUMMARY

Over the years, the growth in both the power and use of Artificial Intelligence (AI) has been staggering. This revolution fueled by significant advances in computing power and the explosion of digitization has led to a new golden era for AI.

Machine Learning (ML) continues to play an integral role in advancing this AI revolution. ML systems can "learn" and adapt in ways that the old expert systems that made up earlier forms of AI could not. There have been significant improvements in video and natural language recognition that seemed impossible only a few years ago. These ML techniques apply statistical inference to large input datasets, taking advantage of declining cost of computing power, to make predictions, create and analyze large data, and automate tasks that were either previously completed by humans or that were not even possible.[1]

The success of ML techniques enabled it to quickly spread throughout industries, including governments and the work of government agencies. Numerous reports highlight surveys of industry experts demonstrating the current and continued expected impact of AI tools, such as ML, natural language processing (NLP), robotic process automation (RPA), or artificial neural networks, in disrupting how private firms conduct their business.[2] AI tools are also augmenting and automating the work of governments.[3] Private industry faces its own challenges with adopting AI tools to make companies more profitable, but many more challenges are faced by government leaders and public managers. Those doing the work of government not only have to go about achieving their objectives and doing so in an efficient manner, but also they face

1. Matthew M Young, Justin B Bullock, Jesse D Lecy, Artificial Discretion as a Tool of Governance: A Framework for Understanding the Impact of Artificial Intelligence on Public Administration, *Perspectives on Public Management and Governance,* Volume 2, Issue 4, December 2019, Pages 301–313, https://doi.org/10.1093/ppmgov/gvz014
2. Drexler, K.E., "Reframing Superintelligence: Comprehensive AI Services as General Intelligence," Technical Report #2019-1, Future of Humanity Institute, University of Oxford. 2019.
3. President of the United States. "Artificial Intelligence, Automation, and the Economy." *Executive Office of the President,* 2016, Pages 18-19.

concerns of equity, accountability, transparency, and others that must be factored into the overall risk management strategy associated with the adoption of AI tools to improve overall governance.[4] It is to this challenge faced by government leaders and public managers that this report seeks to assist.

This report begins by highlighting the risk management approach that is currently being taken by the United States government to manage the opportunities and challenges of AI. This overview is not meant to be exhaustive, but rather to bring attention to a number of significant actions that have been taken by the U.S. federal government as a risk management approach to the development and utilization of AI within government and throughout society. This risk management approach is characterized by increasing the opportunities present in AI tools such as increases in task efficiency, decreases in task costs, improvements in decision-making quality, and overall improvements in the delivery of government services and doing more with less resources, and by addressing hazards and challenges in the forms of safety of use of AI tools, security of the tools, privacy concerns with gathering new types and amounts of data, lack of public trust, job loss, malicious hacking attacks, and biased data. This section will highlight work that began under President Barack Obama's administration to study and understand the impact of AI on automation and the economy and potential impacts on the U.S. workforce. This section continues on to describe the development of the National Artificial Intelligence Research and Development Strategic Plan, which sought to begin to develop a national response to the development and growth of AI.[5] Illustrating the bipartisan concern of these issues, President Donald Trump's administration has updated the National Artificial Intelligence Research and Development Strategic Plan for 2019[6] and issued Executive Order 13859 titled "Maintaining American Leadership in Artificial Intelligence."[7] Finally, building on these actions, we highlight the AI in Government Act of 2019, H.R. 2575,[8] which has recently been introduced into the U.S. House of Representatives and is currently being considered by the House Committee on Oversight and Reform. It attempts to codify much of what has been established across the previous reports and the Executive Order. This section is meant to be both a primer on the risk management strategy that the U.S. federal government is utilizing with respect to AI, and to offer one potentially viable approach for other governments to build upon.

Once the U.S. federal government's Risk Management strategy has been discussed, we turn our attention to providing a framework for government leaders and public managers for considering the role of AI tools within public organizations and the provision of public services.[9] That is, the provision of these services has traditionally been provided through a combination of legislation and human discretion. Human discretion refers to those situations in which a human has to make a decision that is relevant to the task at hand to be completed by a government agency. For some tasks little discretion is required, and for some tasks much discretion is required.[10] These decisions are made and tasks completed across all levels of organizations that provide public services. These levels range from those made on the ground

4. Justin B. Bullock. "Artificial Intelligence, Discretion, and Bureaucracy." *The American Review of Public Administration.* 2019. doi/10.1177/0275074019856123.
5. President of the United States. "Artificial Intelligence, Automation, and the Economy." *Executive Office of the President* (2016): 18-19.
6. Biegel, B., and J. F. Kurose. "The National Artificial Intelligence Research and Development Strategic Plan." *Washington, DC: White House* (2016): 8-40.
7. Executive Office of the President. Maintaining American leadership in artificial intelligence (E.O. 13859). 2019, February 11. Retrieved from https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining- american-leadership-in-artificial-intelligence.
8. "H.R. 2575—116th Congress: AI in Government Act of 2019." www.GovTrack.us. August 1, 2019. https://www.govtrack.us/congress/bills/116/hr2575.
9. Young, M., Bullock, J.B., and Lecy, J. 2019.
10. Ibid

at the micro-level—where decisions are made about how to directly provide services ranging from policing, teaching, health services, and an array of basic services from public utilities to transportation to zoning—to those decisions made within organizations at the organization meso-level. Here is where personnel decisions, department goals, internal resource deployment, and budgeting decisions are made. Finally, there is the institutional or macro-level, where the formulation of rules is made and networks are maintained. This framework will help public managers and government leaders identify the types of tasks and decisions that more clearly benefit from the use of AI tools while working to avoid task environments and decision spaces that are clearly better left to human discretion. Additionally, multiple types of tasks are identified for help from AI tools, including help gathering data, analyzing data for casual relationships, and the automation of certain tasks.[11] We conclude this section by examining how the use of AI tools may impact a number of evaluative concerns for government, particularly with respect to managing risks for government agencies.

Following the introduction of this guiding framework for government leaders and public managers, we turn our attention to two cases of adoption of AI tools by two local governments. We examine the risk management strategy of the City of Bryan, Texas which became one of the first municipalities to offer self-driving shuttle rides for residents as part of a partnership with industry and university researchers. We also examine how the City of Syracuse has both invested in the necessary infrastructure to support autonomous vehicles and other AI tools, and to proactively develop a comprehensive risk management strategy to guide future adoption and implementation decisions.

Finally, we offer recommendations for identifying both relevant opportunities and hazards presented by AI tools to government leaders and public managers as part of their overall risk management strategy. These recommendations will aid in decision making that maximizes the opportunities from AI tools while minimizing the challenges and hazards. While this list is certainly not exhaustive, it provides those doing the work of government with a starting point to quickly think through some of the most important considerations for adding AI tools to their broader toolkit of delivering government services effectively, efficiently, and equitably as possible.

11. Nabatchi, T., and Carboni, J. "Assessing the Past and Future of Public Administration: Reflections from the Minnowbrook at 50 Conference" *IBM Center for The Business of Government* (2019). http://www.businessofgovernment.org/report/assessing-past-and-future-public-administration-reflections-minnowbrook-50-conference#overlay- context=node/3464.

The U.S. Federal Government's Risk Management Approach to AI

Artificial Intelligence (AI) continues to be a term that is greatly misunderstood among the general public. This misunderstanding likely spills over to government officials and managers. Between apocalyptic accounts that often occur in popular discussions and complete dismals of overhype by others, the public manager is often without a frame for understanding the real issues and opportunities that have been identified by experts. This section provides a detailed overview of the high-level approach the United States has taken in considering the opportunities and challenges generated by AI across the previous and current U.S. presidential administrations. We conclude the section by discussing how the multiple reports, executive orders, and bills act collectively as the U.S. federal government's risk management approach for applying AI tools to the work of government

A reasonable risk management approach to the use of AI tools by public managers must carefully consider both potential positive and negative outcomes, opportunities, and challenges associated with the use of these tools. Furthermore, public managers need some understanding of the relative likelihood or probability that the positive or negative outcomes will occur. A risk management strategy should attempt to maximize positive outcomes while minimizing the challenges and hazards to those positive outcomes. To help carefully identify the potential positive and negative outcomes, and how likely they may be, the U.S. federal government has begun to engage leading scholars and thinkers about AI, and AI R&D.

While the U.S. is arguably still the global leader in AI R&D, many countries—most notably China—are drastically increasing their AI capital and capabilities.[12] To manage both the challenges and opportunities that AI tools present, the U.S. is pursuing a federal risk management strategy that seeks to reap the continued benefits of AI while mitigating the hazards.[13] Beginning in 2016, under the Obama administration, the U.S. began to create a body of knowledge and to enact a systematic approach to managing AI for the betterment of society while working to minimize the challenges.[14] The work in this area under the Obama administration consisted of a major report on the impacts of AI with respect to the automation of tasks and the impact on the overall U.S. economy—and building from here towards a National Artificial Intelligence Research and Development Strategic Plan. As will be further highlighted below, investments in AI R&D, educating and training the American workforce to use AI tools, creating opportunities for human and AI tool collaboration, increasing data access, and expanding public/private partnerships are all clear strategies for maximizing the opportunities that AI tools present. On the other hand, developing technical standards for evaluation and benchmarks, creating robust and trustworthy AI tools, and addressing a myriad of legal, ethical, and corresponding social impacts all present strategies for addressing potential hazards of the misuse of AI tools.

Building upon this work, and underscoring AI's increasingly rare position as a topic of bi- partisan concern, the Trump administration issued Executive Order 13859 in February of 2019, titled "Maintaining American Leadership in Artificial Intelligence," to direct federal agencies and executive departments in how they should handle the opportunities and challenges related to AI in their day-to-day activities. Later that year the administration also updated the National Artificial Intelligence Research and Development Strategic Plan with new research and amended strategies. Congress also took up the issue of AI risk management in June of 2019 with the introduction of H.R. 2575, "The AI in Government Act of 2019," in the U.S. House of Representatives. As originally drafted, this bill would consolidate many of the efforts towards managing AI's opportunities and challenges under the umbrella of an "AI Center of

12. Husain, A. *The Sentient Machine: The Coming Age of Artificial Intelligence*. New York, NY. 2017.
13. Justin B. Bullock, Robert A. Greer, and Laurence J. O'Toole. Managing Risks in Public Organizations: A Conceptual Foundation and Research Agenda. *Perspectives on Public Management & Governance* 2(1):75-87. 2019. https://doi.org/10.1093/ppmgov/gvx016.
14. President of the United States. "Artificial Intelligence, Automation, and the Economy." 2016.

Excellence" to be housed within the General Services Administration (GSA). These documents are reviewed in detail below.

Taken together, these documents will help paint the landscape of potential opportunities and hazards for the use of AI tools in government and how we can work to realize the positive outcomes while minimizing the present challenges and hazards. As the following reports and complementary academic work will highlight, AI tools present opportunities to provide increases in task efficiency, decreases in task costs, improvements in decision making quality, and overall improvements in the delivery of government services. However, there are a number of challenges and hazards to the use of AI tools that threaten negative outcomes from AI tools, these include: safety of use of AI tools against misuse, security of the tools against malicious or even incompetent actors, privacy concerns with gathering new types and amounts of data, lack of public trust, job loss, malicious hacking attacks, and biased data.

The documents reviewed below highlight the strategy of the U.S. federal government with respect to AI tools and AI policy. After these documents are reviewed, we discuss how these documents implicitly form the U.S. risk management strategy with respect to the use of AI tools to augment and automate human discretion within the work of government.

AI, Automation, and the Economy—Executive Office of the President, December 2016[15]

The Obama administration published "Artificial Intelligence, Automation, and the Economy" in December of 2016. It was one of the administration's final publications, and was written by a team that included senior executives and staff from the Council of Economic Advisers, Office of Science and Technology Policy, the Domestic Policy Council, the Office of the Chief Technology Officer, and the National Economic Council. Rather than focusing directly on the implications of AI for the work of government, it addresses the broad-based potential impact of AI on the economy, society, and the future of work.

Its findings include predictions based on expert opinion that AI will continue to automate more tasks—and, ultimately, more jobs—that are currently completed by humans. The overall impact on employment is uncertain, but the report notes that experts expect the impacts of automation to predominantly affect lower-paid, lower-skilled, and less-educated workers for the immediate to short-term future. While the report notes the equity-harming implications of these structural changes to the national labor market, it also argues that advances in AI should contribute to overall economic growth, thereby increasing total social wealth. The authors highlight that advances in technology need not be deterministic in how they impact society and, furthermore, suggest three policy responses for managing the opportunities and risks brought about by advances in AI. These include:

1. Investing in and developing AI for its many benefits

2. Educating and training Americans for jobs of the future

3. Aiding workers in the transition and empowering them to ensure broadly shared growth

This report highlights the potential benefits to the overall economy of improvements in AI, but also warns of the potential detrimental impacts to workers. While it focuses primarily on pri-

15. AI, Automation, and the Economy—Executive Office of the President, December 2016.

vate sector AI implementation and corresponding labor market effects, it is also clear that government agencies and government workers also engage in work that will be automated and augmented by AI tools. Thus at least one directly relevant suggestion for public managers is to begin thinking carefully about how to educate and train future government workers while investing in AI for its benefits to effectiveness and efficiency. Many of the tasks currently completed by human government workers will be automated and augmented, and thus different worker needs and skills will characterize the future working space of government workers.

The National Artificial Intelligence Research and Development (R&D) Strategic Plan—National Science & Technology Council, 2016 with an update in 2019[16][17]

In 2016, The United States National Science & Technology Council (NSTC) published "The National Artificial Intelligence Research and Development Strategic Plan," and released an updated version in June of 2019. The 2019 revisions incorporate new research, and adds an additional strategic objective to the seven originally included in the 2016 edition. As of the 2019 revision, the objectives for this strategic plan include:

1. Make long-term investments in AI research. Prioritize investments in the next generation of AI that will drive discovery and insight and enable the United States to remain a world leader in AI. **2019 update: Sustaining long-term investments in fundamental AI research.**

2. Develop effective methods for human-AI collaboration. Increase understanding of how to create AI systems that effectively complement and augment human capabilities. **2019 update: Developing AI systems that complement and augment human capabilities, with increasing focus on the future of work.**

3. Understand and address the ethical, legal, and societal implications of AI. Research AI systems that incorporate ethical, legal, and societal concerns through technical mechanisms. **2019 update: Addressing ethical, legal, and societal issues in AI.**

4. Ensure the safety and security of AI systems. Advance knowledge of how to design AI systems that are reliable, dependable, safe, and trustworthy. **2019 update: Creating robust and trustworthy AI systems.**

5. Develop shared public databases and environments for AI training and testing. Develop and enable access to high-quality datasets and environments, as well as to testing and training resources. **2019 update: Increasing access to datasets and associated challenges.**

6. Measure and evaluate AI technologies through standards and benchmarks. Develop a broad spectrum of evaluative techniques for AI, including technical standards and benchmarks. **2019 update: Supporting development of AI technical standards and related tools.**

16. The National Artificial Intelligence Research and Development (R&D) Strategic Plan—National Science & Technology Council, 2016.
17. Ibid

7. Better understand the National AI R&D workforce needs. Improve opportunities for R&D workforce development to strategically foster an AI-ready workforce. **2019 update: Advancing the AI R&D workforce, including those working on AI systems and those working alongside them, to sustain U.S. leadership.**

8. **(2019 update addition) Expand Public-Private Partnerships to accelerate advances in AI. Promote opportunities for sustained investment in AI R&D and for transitioning advances into practical capabilities, in collaboration with academia, industry, international partners, and other nonfederal entities.**

These eight objectives illustrate the ambitious strategy the United States government is taking to capitalize on the opportunities presented by being a leader in AI while also working to mitigate hazards of concern. The strategies encourage long-term investment, improvements in data access and data sharing, setting standards, expanding partnerships, and anticipating AI R&D workforce needs, all of which are designed to increase the reach and utilization of AI by the U.S. and thus expand the opportunities presented by AI tools. On the other hand, attempts to address ethical, legal, and societal implications and to ensure the safety and security of AI point to concerns about the challenges and hazards presented to the U.S. by AI development. This report also offers a general set of strategies and concerns for other governments to consider in their own risk management towards increasing positive opportunities and mitigating hazards in the deployment of AI to the work of government.

Executive Order #13859: Maintaining American Leadership in Artificial Intelligence—Office of the President, February 11, 2019[18]

The Trump administration issued Executive Order (EO) 13859, titled "Maintaining American Leadership in Artificial Intelligence," on February 11, 2019. The EO declares that it is in the United States' economic and national security interests to remain a global leader in shaping the development and utilization of AI with U.S. values, policies, and principles. To this end, and building from the work completed towards the American AI Initiative, the EO specifies a set of principles and objectives for the U.S. to maintain American leadership in AI. The guiding principles include:

1. The United States must drive technological breakthroughs in AI across the federal government, industry, and academia in order to promote scientific discovery, economic competitiveness, and national security.

2. The United States must drive development of appropriate technical standards and reduce barriers to the safe testing and deployment of AI technologies in order to enable the creation of new AI-related industries and the adoption of AI by today's industries.

3. The United States must train current and future generations of American workers to develop and apply AI technologies to prepare them for today's economy and the future.

4. The United States must foster public trust and confidence in AI technologies and protect civil liberties, privacy, and American values in their application in order to fully realize the potential of AI technologies for the American people.

18. Executive Order 13859. https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence.

5. The United States must promote an international environment that supports American AI research and innovation and open markets for American AI industries, while protecting our technological advantage in AI and protecting our critical AI technologies from acquisition by strategic competitors and adversarial nations.[19]

These five principles in turn inform six specific strategic objectives for executive departments and agencies, including:

1. Promoting sustained investment in AI R&D in collaboration with industry, academia, international partners and allies, and other nonfederal entities to generate technological breakthroughs in AI and related technologies and to rapidly transition those breakthroughs into capabilities that contribute to our economic and national security.

2. Enhancing access to high-quality and fully traceable federal data, models, and computing resources to increase the value of such resources for AI R&D, while maintaining safety, security, privacy, and confidentiality protections consistent with applicable laws and policies.

3. Reducing barriers to the use of AI technologies to promote their innovative application while protecting American technology, economic and national security, civil liberties, privacy, and values.

4. Ensuring that technical standards minimize vulnerability to attacks from malicious actors and reflect federal priorities for innovation, public trust, and public confidence in systems that use AI technologies; and develop international standards to promote and protect those priorities.

5. Training the next generation of American AI researchers and users through apprenticeships; skills programs; and education in science, technology, engineering, and mathematics (STEM), with an emphasis on computer science, to ensure that American workers, including federal workers, are capable of taking full advantage of the opportunities of AI.

6. Developing and implementing an action plan, in accordance with the National Security Presidential Memorandum (NSPM) released on the same date, entitled "Protecting the United States' Advantage in Artificial Intelligence and Related Critical Technologies," to protect the advantage of the United States in AI and technology critical to United States economic and national security interests against strategic competitors and foreign adversaries.[20]

AI in Government Act of 2019 (H.R. 2575)—116th Congress (2019-2020), May 8[21]

The U.S. federal government has been studying AI and has begun to coalesce around guiding principles and strategic objectives. This work has culminated in the form of the previously discussed Executive Order instructing the National Science and Technology Council Select Committee on Artificial Intelligence to coordinate with federal agencies to take on these new government priorities and objectives. The AI in Government Act of 2019, in its current form,

19. Ibid
20. Ibid
21. H.R. 2575, 2019

seeks to complement this work and establish an "AI Center of Excellence" within the GSA to:

1. Advise and promote the efforts of the federal government in developing innovative uses of AI by the federal government to the benefit of the public

2. Improve cohesion and competency in the use of AI[22]

The duties of a Center of Excellence are further enumerated in the act and provide a framework for an organization that could help with risk management of AI within government to both build upon opportunities and mitigate hazards. These suggested duties include:

1. Regularly convening individuals from agencies, industry, federal laboratories, nonprofit organizations, institutions of higher education, and other entities to discuss recent developments in AI, including the dissemination of information regarding programs, pilots, and other initiatives at agencies, as well as recent trends and relevant information on AI

2. Advising federal government acquisition and use of AI through technical insight and expertise, as needed

3. Assisting agencies in applying the management and use of data in application of AI

4. Identifying and disseminating information regarding education and workforce development opportunities for agency employees relevant to AI topics, and leading those opportunities as needed

5. Studying economic, policy, legal, and ethical challenges and implications related to the use of AI by the federal government, including how the privacy, civil liberties, and civil rights of individuals are or will be affected by the use of AI by the federal government

6. Encouraging and assisting joint initiatives with state or local governments, regional organizations, private businesses, institutions of higher education, nonprofit organizations, and federal laboratories to advance the innovative use of AI in government

7. Assisting relevant agencies in developing and maintaining plans for the governance of agency AI systems[23]

For a risk management strategy to be effective, it not only needs research and attention by government officials, but also a resourced and organized organizational "home" to continue the work of implementation, evaluation, and adjustment. As written, H.R. 2575 represents the legislative branch's attempt to put this knowledge to practice by codifying the research, strategies, and recommendations enumerated in the Executive Office of the President report, the National AI R&D Strategic Plan, and Executive Order 13859. This Center could provide an organizational space chartered to implement the U.S. AI risk management strategy across and throughout the federal bureaucracy.

Discussion

As laid out in the Artificial Intelligence, Automation, and the Economy report, the National Artificial Intelligence Research and Development (R&D) Strategic Plan and the American AI

22. Ibid
23. Ibid

Initiative, and the AI in Government Act of 2019, the U.S. recognizes the immense opportunities and hazards presented by AI's current and potential capabilities and utilization. Viewed systemically, these documents spell out the United States' risk management approach to AI tools in the AI era. To accomplish this risk management, the U.S. federal government has provided a framework that is generally beneficial to other governments.

Opportunities and Hazards

The U.S. federal government highlights important opportunities for the development and use of AI tools including sustained investment in AI R&D, high-quality and accessible data for AI modeling, building domestic and international partnerships across governments, industry, and academia, and reducing barriers more generally to the utilization of AI. These documents also highlight a number of hazards to mitigate including public confidence, public trust, privacy, and civil liberties, technical risks, and potential malicious attacks, which are considered to be a significant threat in fast growing technologies such as AI tools. Additionally, the opportunity of maintaining the government's human capital in the area of AI by working to train the next generation of workers and scientists to be leaders in the utilization and development of AI is emphasized. Finally, the overall risk management strategy makes clear the importance of the opportunity of building partnerships across a variety of stakeholders, including other governments, academia, and industry as the challenges presented by AI requires the attention of a variety of stakeholders—an "all hands-on deck" approach.

Risk Management Process

This section has also illustrated a useful risk management process for identifying opportunities and challenges with AI tools in the AI era. First, the federal government created a research report summarizing what is known by the leading experts on AI. Second, this report fueled a National AI R&D strategic plan that set into motion an initial seven strategies for the federal government to begin to systematically approach this risk management challenge. Third, following the transition from a Democrat administration to a Republican administration, the new administration—also recognizing the importance of having a holistic risk management approach to the AI era—updated the strategic plan. Fourth, in an attempt to encourage more action on the strategic plan, the new administration issued an Executive Order for agencies to begin implementing the strategies and recommendations put forward in the strategic plan, all while encouraging the development of partnerships across the federal government and across the numerous relevant stakeholders. Fifth, and finally, the legislative branch has taken up this issue and offers an attempt to both codify parts of the strategic plan and authorize a Center of Excellence within the GSA to be the organizational home of implementing and maintaining this strategic plan.

This section has provided opportunities and challenges presented to the U.S federal government by AI tools and a basic overview of a process for developing a successful AI risk management strategy.

However, missing is a framework for how more specifically managers might navigate the space of utilizing AI within their own organization. In the next section, we introduce a guiding framework for how managers can utilize AI tools to augment and automate current tasks within their organizations to maximize the opportunities and mitigate the hazards. Managers need to take care when deciding how and where in their organization they choose to deploy AI tools.This section provides a guiding framework to aid managers in making those decisions and considering the potential impacts as part of an overall risk management strategy.

How can AI be used to Augment Human Decision Making?

Understanding how artificial intelligence interacts with human decision making requires some definitional clarity about what artificial intelligence is. In broad terms, AI refers to technology that allows machines to perform tasks that require a significant level of intelligence. This has made establishing a clear boundary around what technically constitutes AI difficult, both because of how difficult it is to define intelligence itself, and because artificial intelligence becomes a victim of its own success—applications that were once cutting-edge are now so seamless and ubiquitous that most people no longer consider them AI. For example, email spam filters and automatic spell-checking in word processing were once at the frontier of AI development. For our purposes however, the broad definition of AI as affording machines the ability to perform tasks that require some degree of intelligence is sufficient because it captures the issue at the center of adopting and implementing cutting-edge autonomous systems in the public sector.[24]

Modern autonomous systems employ a particular branch of AI often referred to as 'Deep Learning,' which is marketing shorthand for a specific form of AI design called Artificial Neural Networks (ANNs), which is in turn part of a branch of AI architecture known as 'Machine Learning.' ANNs mimic the biochemistry of neural pathways in animal brains; an input is passed through a series of decision nodes arranged in several interconnected layers, and eventually produces an output. The machine thus analyzes data by summing up representations of its different elements in a way that is roughly analogous to how we (often unconsciously) make determinations. Crucially, this mimicry includes the ability to adapt the network by reinforcing pathways and combinations that led to a correct output, and deprioritizing pathways that resulted in incorrect output. A consequence of this design approach is that Deep Learning-based AI needs to be 'trained' using a trial-and-error process with relevant data before it can perform reliably. In turn, the quantity and complexity of data used in both training and once the system is deployed improves its performance. This feature is what puts the 'learning' in Machine Learning. 'Deep' Learning simply refers to particularly complex ANN architectures with many nodes and many layers. Put simply, Deep Learning allows machines to make their own decisions about how best to make a determination instead of requiring explicitly defined processes in the form of human- written algorithms.

As the cost and capacity of the hardware, software, and data generation necessary for developing and running these systems decreased exponentially over time, their capacity to perform increasingly complex tasks as well or better than expert humans has increased. This leads to the description of AI as "data greedy." It is also a large motivating factor behind the market for consumer and user data at the center of companies like Facebook and Google's business models. And, as the public sector continues to expand and modernize its data-generating capacities, it also means that the use cases for Deep Learning-based AI in the public sector will continue to expand. The breadth, depth, and pace of these advances has made more and more jobs—including "white collar" professions that consist of intellectual rather than physical labor—highly vulnerable to elimination by automation. At the heart of this new frontier of intellectual labor automation is the decision-making capacity of Deep Learning-based AI systems.

For public managers and administrators, their relative degree of decision-making authority is referred to as their level of administrative or bureaucratic discretion. Exercising administration discretion (often described synonymously as "exercising professional judgement") involves making one or more determinations that classify objects or events, which may trigger contingent organizational actions. Recent work introduces a conceptual framework for understanding the use of decision-making AI in the public sector "to augment or automate the exercise of

24. Young, M., Bullock, J.B., and Lecy, J. 2019.

administrative discretion," which they define as "artificial discretion."[25] Public managers need to understand AI tools and the different contexts in which they may be implemented if they are serious about developing a risk management approach for these new technologies. The framework emphasizes two important characteristics whenever considering whether to use AI-based systems to augment or automate organizational tasks—first, the degree of discretion required to execute the task and, secondly, the level within the organization/institutional environment where the task takes place.[26]

Task Characteristics: Degree of Discretion and Organizational Location

Every organization requires countless individual tasks, whether they are physical (e.g., hauling trash) or mental (e.g., deciding how to optimally route trash removers across a city) in order to function. The degree of administrative discretion required for a given task is a function of the amount of inherent uncertainty involved in the tasks' execution. And just as the amount of uncertainty in any situation varies according to the situation's characteristics and some amount of random chance, so does the amount of discretion required of organizational agents to complete different types of tasks. At one end of this continuum, tasks have sufficiently little associated uncertainty that they can be fully specified ahead of time, allowing for relatively complete routinization. Agents performing these tasks exercise little to no discretion, as the set of operations required of them are predetermined. Other organizational tasks take place in less stable and certain environments.[27] In these cases, agents are empowered to exercise discretion in order to reduce organizational uncertainty and allow for any necessary contingent activities.

Finally, at the opposite end of the continuum, agents are responsible for tasks that take place under extreme levels of uncertainty. These agents are therefore entrusted with a great degree of autonomy in which to exercise their discretion.

A second crucial task characteristic is its location within the organization and its broader institutional environment.[28] Task location affects the degree of interdependence and the amount and strength of constraints placed on the agent responsible for its execution. Similar to the discussion of discretion, the framework characterizes tasks at three distinct organizational levels: micro, meso, and macro.

Micro-level tasks are those that are carried out by frontline staff—often referred to in public organizations as "street-level bureaucrats"—in their standard duties. One example of this type of task is the filling of potholes by workers in a municipal public works department. On average these tasks have little to no contingent dependencies on other organizational agents or units, nor do they depend on extra-organizational advice or consent; these considerations have been "baked" into the predefined rules and requirements associated with the task, and agents are in turn entrusted to carry them out.

Meso-level tasks consist of activities that affect and/or involve multiple intraorganizational units. Examples of meso-level tasks range from facilities operations to developing and implementing performance management practices. These tasks are inherently more interdependent and constrained because they require coordination across different organizational units.

25. Ibid.
26. Ibid.
27. Bullock 2019
28. Young, M., Bullock, J.B., and Lecy, J. 2019.

Macro-level tasks extend the interdependencies and constraints imposed by intraorganizational coordination to the organization's external, or institutional, environment. Examples of macro-level tasks include managing or participating in multi-organizational collaborations, and representing organizational interests in higher-level policy formulation processes.

Table 1 provides examples of the interaction between the degree of discretion and location associated with different organizational tasks. It is important to note that the tripart distinction for both degree of discretion and task location is made arbitrarily for the sake of explanation. The location and degree of discretion required for organizational tasks is a continuum whose slope will vary from one organization to another. Furthermore, the slope will almost certainly not be monotonically increasing as organizational tasks shift from frontline workers to management.

Even executives are required to perform some tasks that are highly routinized, and frontline staff often find themselves in unique circumstances where predefined processes do not apply. That said, this discussion is not meant to be exhaustive of all possible scenarios—rather, it is presented as a heuristic for developing an informed understanding of the context associated with a task before deciding whether or how to augment or automate it using AI.

Table 1. Matrix of Task Analysis by Level of Analysis and Degree of Discretion

Degree of Discretion Level of Analysis	Low Discretion	Medium Discretion	High Discretion
Micro (Individual)	Data Entry, Issuing Licenses or Permits	Placing Children in Foster Care, Sentencing/Parole	Emergency Response
Meso (Organizational)	Facilities Operations	Hiring Processes, Performance Management	Goal Setting, Strategic Planning
Macro (Institutional)	Statutory Obligations	Policy Formulation	Crisis Response and Management

Source: Adapted from Young, Bullock, and Lecy (2019). Used with authors' permission.

AI Tools and Risk Management

Innovation and risk are inextricable from one another. Implementing AI tools in the public sector carries the additional hazard of novelty. As a relatively new technology with a simultaneously expanding range of applicability, there are fewer empirical examples to draw upon for evidence of success. This section builds upon the previous analysis of the nature of AI tools and important characteristics of the tasks these tools could be used for to create an evaluative framework for the use of AI tools to accomplish the work of government. There are two principle components to the framework. The first synthesizes the prior discussions of the technology behind AI tools and task characteristics to identify likely best-fit scenarios between technology and task. The second draws from seminal work in public administration theory by to provide a set of evaluative criteria for using AI tools.[29] These criteria are meant to be used assess both opportunities and hazards before adoption of an AI tool, and to design program evaluation processes after the tool's implementation.

29. Salamon, Lester M. *The Tools of Government: A Guide to the New Governance.* Oxford University Press. 2002.

Mapping Potential Uses of AI Tools to Tasks by Degree of Discretion

Because AI tools are, at their core, the automation of decision making, there are many different possible ways to implement it in public organizations. It can be used—as is commonly the focus of popular media when discussing modern AI—to fully automate tasks. But it can also be used to generate novel relational data from unstructured signal (e.g., video feed, still images, text) and reduce the complexity of high-dimensional data through pattern recognition and discovery of complex relationships that are otherwise imperceptible to human agents. And there are other uses still that hybridize the two. AI tools can also serve as a decision support tool, where human agents retain final decision-making authority but can leverage AI's scalability and analytic capacity to factor previously unavailable information into their decision.

While this versatility is an advantage in many ways, it also creates potential hazards. Some AI tools may prove less effective or worse than the status quo for a particular task. Decision support tools can effectively supplant the human decision makers if they are punished for deviating from the tool's recommendations. Large amounts of new data can both overwhelm decision makers and be ultimately unnecessary. And in addition to the social disruption that automation causes, it will inevitably make a mistake—which will carry heightened political consequences in addition to the direct consequences of task failure, an extremely important hazard to note and difficult to mitigate.

Those interested in implementing AI tools in their organizations can manage these risks in part by considering how the relationship between task characteristics and uses of AI tools interact. Drawing upon the previously established three-level conceptual model for task discretion, and on the current capacities of its underlying technology, we argue that the degree of discretion associated with a task is the strongest determinant of the most appropriate use for AI tools. Table 2 summarizes these pairings below.

Table 2. Potential Use of AI Tools for Tasks by Degree of Discretion.

Low Discretion	Medium Discretion	High Discretion
Automation	Decision Support Tool, Predictive Analytics	New Data Generation, Reduction of Data Complexity, Relationship Discovery

Source: Adapted from Young, Bullock, and Lecy (2019). Used with authors' permission.

Tasks with low levels of discretion are the most suitable for complete automation. Their correspondingly low levels of situational and outcome uncertainty make it comparatively easy to leverage AI's capabilities, so that it is likely to dominate human agents in terms of speed, error rate, and unit cost. Moreover, because they do not allow for much (if any) creativity on the part of agents, low-discretion tasks are often repetitious and dull. In addition to representing an under-utilization of human potential, this is also a recipe for increased error rates due to agent complacency or boredom. Automating these tasks can thus also free up human labor for more enriching and productive work.

As the degree of discretion required increases, automation becomes less feasible both because of task complexity and the social and political consequences associated with many of these tasks.[30] This is not to say that AI tools will necessarily make worse decisions for these tasks than their human counterparts, though that is certainly more likely than in low-discretion tasks. Rather, its use as a decision support tool is more appropriate because these cases frequently involve satisfying multiple competing, sometimes conflicting, objectives with less certainty around the first- and second-order consequences of the decision. This is not only a technical challenge for current AI systems, it is also a political challenge as both public organizations and the public at- large are likely to balk at ceding final decision-making authority to machines. As a decision support tool, AI tools can augment the human agent's process by increasing their ability to gather and process information, including scenario modeling and forecasting.

Finally, public organizations can deploy AI tools to collect and process new data in real-time and apply its strength in pattern recognition and relationship discovery to decrease the high degree of uncertainty that is endemic to tasks that require the most discretion.[31] These tasks can involve making time-sensitive decisions in situations with poor data, complex interconnected systems, and environments that are highly mutable over time. The social and political ramifications associated with these tasks make their automation highly unlikely. Moral calculations that create winners and losers under adverse—including life-or-death—conditions require human decision makers. Put another way, both people and organizations need the recognizable accountability of a human face associated with emergency response coordination in the immediate aftermath of a large-scale disaster, whether the response is adequate or not.

Evaluative Criteria for AI Tools

Seizing the opportunities and mitigating the hazards associated with adopting AI tools requires paying careful attention to the match between technology and task, understanding the organizational status quo with respect to task execution and performance, and having a clear idea of what constitutes success with respect to outputs and outcomes.[32] In this section we present a five-point evaluative framework adapted from Salamon's New Tools of Governance for public managers to use when considering whether or not to adopt AI tools. The five criteria are effectiveness, efficiency, equity, manageability, and political legitimacy.[33] Evaluating potential adoption cases against these criteria will reduce uncertainty with respect to both the technology and the status quo task environment, allowing for more informed decision making and improved risk management.

Estimating the likely effect of implementing AI tools for organizational tasks is more difficult when the use case itself is novel to the organization. One solution to this problem is to engage in interorganizational learning: identify peer organizations with similar characteristics and responsibilities that already perform the task and audit their processes, outputs, and outcomes as a proxy. Even when this approach is possible, however, it cannot account for uses where the task was simply not possible or feasible without using AI-augmented decision processes. In these cases, decision makers can identify potential hazards by identifying the intended—and, when possible, unintended—consequences.

30. Young, M., Bullock, J.B., and Lecy, J. 2019.
31. Ibid.
32. Bullock, Greer, O'Toole 2019
33. Salamon 2002

Efficiency

While the technical definition of efficiency—the ratio of output generated to input required—is constant across the private and public sectors, the latter adds additional policy and political constraints to its operational estimation. AI tools are extremely efficient once implemented because, as is the case with software in general, the marginal cost per task execution approaches zero over time.[34] What's more, its associated fixed costs at startup can be lower than traditional enterprise solutions because its physical infrastructure can use either preexisting systems or be offloaded to web-based services. On the other hand, organizations may need to substantively invest in data generation and/or storage capacity if they do not already have the data necessary for AI tools to perform well. That said, for any task that an AI tool can complete it is quite likely that they will dominate human agents with respect to efficiency.

Unfortunately, this dominance creates its own set of hazards. The most efficient solution is not necessarily most effective or equitable with respect to outcomes. Decision makers facing tight fiscal constraints can thus find themselves satisficing by prioritizing efficient solutions that counterintuitively end up costing the organization more in terms of its ability to meet its goals. With its high scalability, AI tools' overall efficiency appeal may lead to its adoption in instances where an alternative solution will make the organization better off. A more robust risk management approach to adopting AI tools should downweigh the value of efficiency with respect to the other criteria used in this framework. This is particularly important when the use case is new to the organization; in these instances, decision makers should follow the Ian Malcom Principle.[35]

Effectiveness

Effectiveness is the quantification of the degree of success associated with one or more attempts to meet a predefined objective. Accurate and consistent measures of effectiveness require a strong understanding of the task's objective, and what constitutes incremental success towards it. The pace of advances in the field of AI research, paired with the continued growth in digital data generation, makes the boundary between AI tools and human discretion a moving target. Taking the position that human discretion will remain indefinitely more effective for a given organizational task than AI tools, this requires strong assumptions that are ill suited for risk management approaches to innovation adoption.

Because of this inherent uncertainty, organizations should begin by carefully considering what the appropriate threshold for task quality ought to be. While it is natural to establish this threshold in terms of success rates, it is often more instructive for risk management to explicitly consider failure rates—simply the inverse of the success rate. Acceptable failure rates should account for both Type I (making the wrong choice) and Type II (not making the correct choice) errors. Once organizations determine these thresholds, they can consider small-scale or, ideally, controlled experiments comparing the effectiveness of an AI tool's task automation or augmentation against the base case of current processes. One crucial but easy to overlook element of this comparison is the level of variance in task effectiveness over multiple iterations within and between existing agents, whether human or machine. Once the average task quality of an AI tool is sufficiently high, its replicability and scalability may lead it to outperform unassisted human agents at a systemic level, even if the best-performing human agents are individually more effective.

34. Young, M., Bullock, J.B., and Lecy, J. 2019.

35. As stated by the character Dr. Ian Malcom in the film adaptation of Michael Crichton's Jurassic Park (1993): ". . . your scientists were so preoccupied with whether or not they could that they didn't stop to think if they should."

Equity

Equity-related concerns with AI tools are perhaps the most challenging with respect to risk management. There are many different ways to define equity, especially as it relates to the distributional mission and effects of public organizations. Here, we use the two dimensions of equity employed by Salamon: equitable tools should treat all members of the public equally, and should be redistributive in favor of disadvantaged individuals and groups.[36] Several cases involving the use of AI tools in both the public and private sectors have raised serious concerns about its equity-harming effects. In particular, decision support tools used in local and state criminal justice systems have been found to recommend harsher decisions for black suspects and defendants than white ones, even when other strong predictors of recidivism suggested that white individuals were at greater risk of re-offending.

We argue that these cases highlight a number of hazards of implementing AI tools without paying deliberate and careful attention to the details of how the system was architected and trained. Administrative data are, fundamentally speaking, the documentation of information inputs, administrative decision making, and the associated outputs and outcomes. To the extent that these past decisions were prejudiced by structural or individual biases, these prejudices are embedded as statistical patterns in the data itself. Viewed in this context, AI tools trained on administrative data will identify and leverage these patterns to learn the optimal path to the 'correct' output because one of AI's greatest strengths is its capacity to identify structural correlations in complex, high-dimensional data that are not intuitively obvious to humans. And when training data is curated from a pooled source (e.g., images from security cameras for facial recognition), overrepresentation of particular population strata will cause the system to underperform when it encounters strata that were underrepresented (or not included at all).

Again, however, human agents' performance—and extreme variance across individuals—with respect to process-based equity are well-known problems in the status quo operation of public organizations. Comprehensive risk management strategies of adopting AI tools should use the observed performance of the organization and its agents as the base case, rather than an ideal state of perfect equality.

Turning to distributional equity, AI tool's potential impact is similar to other tools in that it largely depends on the purpose and manner in which it is implemented. The same structural concerns that lead to worsening inequality regarding access and ability to influence and shape technologies' use within and upon disadvantaged groups and communities apply to AI tools as well. As a tool, it can be put towards any normative goal, whether equity-enhancing or not. One particularly salient feature of AI tools for distributional equity is the capacity to increase psychological distance between public sector employees and the public. Because individuals on both sides are effectively anonymized to each other, AI tools can increase the hazard of dehumanizing or otherwise oppressive uses of power that might otherwise be too difficult for a human agent to engage in directly.

Manageability

Manageability is a measure of how simple or difficult a tool is to implement and operate.[37] Whether it is mentioned directly or not, one of the motivating principles behind public sector automation is the elimination of the hazard that bureaucrats will exercise their discretion in ways that do not comport with the goals of the organization or the desires of either elected officials, the public, or both. In certain respects, one can argue that transitioning a greater

36. Salamon, 2002.
37. Salamon 2002

share of organizational tasks from human to AI tools is likely to mitigate many of the hazards associated with human resource management and interpersonal dynamics. And even when AI tools are not used to completely automate tasks, its use in decision support and data generation and processing reduces uncertainty, which allows for increased standardization and decreased overall discretion.

Adopting AI tools still introduces hazards with respect to task manageability.[38] Its high scalability and efficiency, and ever-decreasing expertise- and price-based barriers to adoption can lead to overly ambitious implementations that are poor fits for either the context or the scope (or both) of the problem it is meant to solve. If the technology behind one of these systems is provided by a private firm via contract, there are not only the usual hazards associated with contracting but specific hazards with respect to intellectual property rights and the organization's ability to observe, monitor, and evaluate vendor performance. Researchers have also demonstrated that it is possible to design and deploy systems that are purpose-built to identify and exploit vulnerabilities in other AI tools to either induce failure or game its output to the exploiter's advantage. But perhaps the most vexing manageability-related hazard that AI tools introduce is that even though its decision-making process is modeled after the mammalian brain, we have no reliable way of understanding why these systems make a particular choice or determination. So, while we could ask a human agent to describe the thought process or logic model that led them to make a particular determination, AI tools are often "black boxes."

Legitimacy and Political Feasibility

It is impossible to overstate the importance of assessing whether an AI tool is likely to be viewed as legitimate by the public and those with veto authority within the organization and whether the adoption decision has sufficient support to make implementation feasible. While this is almost entirely contingent on the specific context associated with a given case, some generalizable points deserve attention. Each of the preceding criteria moderate perceptions of legitimacy and feasibility; all else equal increased effectiveness, efficiency, equity, and manageability ought to improve both and vice versa.

Similarly, the degree of discretion and organizational location associated with a potential adoption case are likely to condition an AI tool's perceived legitimacy and feasibility.[39] For example, the outright replacement of low-discretion public sector jobs by automation will necessarily be highly contentious even when there are clear efficiency and effectiveness improvements, both because of the real human and social costs that accompany eliminating jobs and because of the relative strength of both civil service and public sector union protections. At the same time, using an AI tool as a data generation and processing tool for disaster response and emergency management does not raise the same issues for legitimacy and feasibility because it constitutes a new process that does not disrupt existing social and labor relations. Yet this use may raise other concerns around privacy or cost. These are political questions, and such questions are endemic to adopting innovations like AI tools in public organizations.

38. Young, M., Bullock, J.B., and Lecy, J. 2019.
39. Ibid.

Case Studies in Local Government: Self-Driving Trolleys in Texas and AI Infrastructure Investments in New York

Self-Driving Trolleys in the City of Bryan, Texas

In mid-September of 2018, the City Council of Bryan, Texas, approved an interlocal agreement to become one of the first cities in the country to deploy self-driving electric trolleys on public rights-of-way. These electric trolleys utilize enormous amounts of input data from cameras and laser imaging, detection and ranging (LIDAR) sensors, to navigate roadways. This input data, in real time, allows the self-driving trolley to navigate public roadways. While fully autonomous, each trolley is staffed by two safety workers who can take control of the vehicle at any moment. The trolleys make roughly eight block trips, with five different passenger stops. In addition to the two safety workers, they can hold up to four passengers each for a total of six potential occupants.

This case is illustrative of an appropriate risk management approach taken by a local government. It explores the potential opportunities for improving the overall quality of life by applying AI tools in the form of gathering new data, analyzing that data, and automating some human discretion by operating the electric trolley. In gathering information for this case, one of the authors of this report was able to speak with multiple key players in this partnership including: City of Bryan's Mayor Andrew Nelson, City of Bryan Deputy City Manager Hugh Walker, City of Bryan Engineer Paul Kaspar, and Lead Researcher Texas A&M University's Associate Professor Srikanth Saripalli to better understand the process by which this project came to fruition.

The City of Bryan has leveraged its proximity to Texas A&M University to develop standing formal and informal relationships and formal partnerships. This regional advantage was not lost on current Mayor Andrew Nelson. In interviews for this report Mayor Nelson identified the importance of continuing to further develop the City's relationship with the university, particularly given that both the university and the City are currently experiencing growth.

Through its specific relationships with Texas A&M University's Department of Mechanical Engineering, Mayor Nelson was made aware of Professor Srikanth Saripalli's desire to create a research project where he could test out the utility of different sensors and AI technology in safe, but also "real world" conditions that could allow for slow moving trolleys to help move pedestrians throughout an urban environment.

Professor Sariapalli has long had interests in autonomous vehicles, unmanned systems, and their guidance and control systems. He has also spent significant time thinking and writing on how to safely introduce autonomous vehicles to American roads. While much of the funding for these projects comes from private industries, Professor Sariapalli was interested in a research project where the data would be freely available to public officials and policy makers in order to increase both transparency of how these trolleys operate and overall public awareness and acceptance. Professor Saripalli's desire to partner with a local government to help increase the social benefits of his research, and Mayor Nelson's desire to continue build relationships with Texas A&M and improve the lives of the residents of the City of Bryan, set the stage for a promising partnership for the use of AI tools in the form of autonomous transportation research.

The partners across Texas A&M and the City of Bryan worked proactively to manage the potential hazards associated with the self-driving trolleys. Much research and work had already been done to ensure the safety and efficacy of these trolleys. They operate at very low speeds—seven to 10 miles per hour—along with two safety workers who could take over the controls at any moment, and there was a further research need to gather more data to better understand the effectiveness of these vehicles completing trips in an urban environment. The previous research also included testing of the vehicles on the Texas A&M campus and shorter

routes in downtown Bryan. The ideas so far speak to concerns with effectiveness, efficiency, and equity of longer- term implementation of fully autonomous vehicles.

From the beginning and throughout the project, concerns about managerial capacity and political legitimacy were also at the forefront of project planning. Professor Saripalli and City officials wanted to learn how the general public felt about these trolleys, if they would use them, and ideas they had for improving them. Professor Saripalli noted these concerns in his earliest interviews concerning the project, and City officials highlighted these concerns in their interviews conducted with one of the authors of this report. To this end, the City of Bryan still has a survey available on their website for residents to give their feedback on their experience with and thoughts about the self- driving trolleys.

As the project negotiations continued, a number of other important actors and institutions participated in both the negotiations and bringing attention to the project, including Texas A&M Transportation Institute Director Dr. Greg Winfree, Vice Chancellor and Dean of Engineering Dr. M. Katherine Banks, and Texas A&M University System Chancellor John Sharp. The support of these powerful actors within their roles of powerful institutions aided in a successful partnership to move this project forward.

As implementation moved forward, Deputy City Manager Hugh Walker and City Engineer Paul Kaspar worked with Professor Saripalli and his team to work through remaining implementation challenges and potential hazards for use of the self-driving trolleys. For example, insurance had to be purchased as the vehicles were to be running on public roads, safety workers needed to be hired and trained, signage and roadway stops needed to be created and maintained for each of the pick-up and drop-off spots, and the vehicles needed to be stored and charged. These managerial challenges were tackled together by members from both the City team and the research team. Each of these hazards were considered along with actions for mitigating their potential impact.

The project began delivering rides to the public in late October of 2018, took a hiatus during much of December and into January of 2019, and continued to provide rides in the early spring of 2019. This project highlights a beneficial partnership between academic and local government to begin applying the benefits of AI to the delivery of government services. The project took a deliberate approach in creating a broad partnership, with multiple relevant stakeholders, and investing significantly in research and development. The project also allowed for the use of all three use types of AI tools: gathering new data from sensors, analyzing that new data, and automating the task of operating a vehicle. The data from this initial phase of the research project is being analyzed and future expansions of the project are under consideration.

Opportunities and Hazards

The City was able to seize upon a number of opportunities from its partnership with the university, branding opportunities for the city, and the potential to make public transportation more effective, efficient, and equitable as potential outcomes from the research project.

Additionally, the City was able to work with its partners to mitigate potential hazards including public confidence (partnering with leading scholars and engineers), public trust (transparency about the process and the project), privacy (protection of individualized data), and civil liberties (safety workers stationed on each trolley as part of the research process), technical risks (using multiple types of inputs to better assist the self-driving trolley's capacity for navigating its environment), and potential malicious attacks on the self-driving trolleys (limiting the trolleys to slow speeds).

Risk Management Process

This case highlights an overall risk management strategy similar to that of the U.S. government. For example, as with the U.S. government, the City of Bryan took a longer-term sustained investment in AI R&D. The project focused on creating high-quality and accessible data for AI modeling that would also be publicly available. It also focused on building partnerships across governments, industry, and academia, as well as reducing barriers more generally to the utilization of AI. This is what the City of Bryan was able to do by engaging researchers from Texas A&M University, by partnering with private actors, and by the city itself taking an active role in hazard mitigation by carefully planning the implementation and feedback of the project.

Artificial Intelligence Infrastructure Investments in the City of Syracuse, New York

Syracuse is a medium-sized city located in central New York State. Formerly home to a strong industrial manufacturing economy, the city over the last two decades has faced declining population and revenues, as well as severe and highly concentrated poverty. Investing in the technological infrastructure for autonomous systems and AI tools might seem out of place given the more immediate challenges that the city faces. A closer analysis, however, shows that the city has leveraged its relationships and a confluence of discrete opportunities over the past several years. The city has both developed the internal capacity necessary to design and manage these projects, offsetting the infrastructure's capital costs by bundling them into investments that will save the city money over time. Simultaneously, City managers have been very deliberate and strategic in their thinking about governance frameworks for autonomous vehicles and other AI tool-based systems, and what they want them to do—and not do—to help the city meet its objectives and improve the well-being of its residents and visitors.

Opportunities and Hazards

The direct catalyst for the city's investment was a change in state law that made it possible for the city to buy its streetlights back from the local utility, which had previously rented them back to the city at a cost of more than a million dollars per year. Once there was legal standing for the buyback initiative, Syracuse was able to affect the purchase by securing its largest bond offering in a generation and supplementary financing provided by the New York State power authority. Establishing property rights for the city's streetlights was the crucial first step; the next is the ongoing replacement of the existing incandescent lights with "smart" LED systems. In addition to saving the city a large amount of money in utility costs, the new lighting systems are also designed from the base up to be "smart." Each light is networked and can be outfitted with other modular sensors to provide the city with real-time data on localized weather, air pollution, and other phenomena.

The other crucial point of owning these streetlights is that the city now has the necessary property rights to sell easements for telecommunication companies to use the streetlights as platforms for wireless network infrastructure. Earlier in 2019, the Syracuse Common Council approved a contract to lease those easements to Verizon in order to rollout what will become the largest citywide 5G network in the U.S. This is also happening at the same time as the city is working with Onondaga County to be part of a planned "drone highway," a 50-mile-long stretch of airspace that will be used to test autonomous aerial vehicles.

The organizational vision and internal capacity necessary for Syracuse to even begin to think about adopting AI-based technology is due in part to its success in securing a Bloomberg Philanthropies Innovation Teams grant in late 2014. The grant provided funding for the city to

hire a team of innovation experts tasked with adopting and implementing new technologies and analytic techniques to improve the city's organizational performance. The value provided by this team was enough to lead the City to create and fund a permanent Office of Accountability, Performance, and Innovation (OAPI), led by the city's chief data officer, Sam Edelstein. But Syracuse also has external resources that it can rely on that are more local and more permanent; the city is home to Syracuse University, which boasts strong schools and faculty in the fields of information science and public administration and policy, among others.

Risk Management Processes

While there are no current fleets of autonomous vehicles or AI-augmented decision processes in Syracuse, the city is actively developing formal institutions and processes for evaluating prospective use cases that its current infrastructure investments will facilitate in the future. City officials as a whole see these as fundamentally high-risk technologies because there is so little known about how they will play out in practice. At the same time, they also recognize the risk of being left behind economically and socially in the future that would come from ignoring or resisting these systems wholesale. Instead, they have chosen an approach summarized by the city's CDO Sam Edelstein via a rhetorical question: "How do you balance immediate community needs with more long-term city needs, and both of those against what these technologies are actually capable of and what risks they introduce?" OAPI and the rest of the City government consider developing and using an evaluative framework to ensure completeness and consistency in their decision making central to this task.

In addition to and in support of their framework, Syracuse officials are also committed to performing extensive due diligence research on the current state of the art and future direction of AI technology. They want to proactively identify not only must-haves but also 'deal-breakers' that signal that a particular system or implementation is not a good fit for the city, whether due to technological limitations, cost, or issues of political legitimacy. They have also drawn upon the expertise of faculty at the SU information school to develop data security and privacy plans that complement the evaluation framework. For example, in interviews for this report that predated the City and County of San Francisco's announcement, Syracuse officials noted that they do not intend to adopt facial recognition technology for city security cameras and policy officer-worn body cameras due to internal and public concerns about equity and privacy.

Another example of the City's proactive risk management approach is its consideration of creative models of data ownership and management. As a relatively small and resource-light administration, Syracuse does not have the infrastructure or internal expertise to safely store and manage petabytes of real-time-generated sensor data. While several private firms offer these services via remote computing, these solutions present their own risks with respect to fiscal and transactions costs, as well as property rights with respect to the data themselves. One alternative option involves a collaborative governance regime or partnership with Syracuse University and one or more other local and regional governments that would have the University host the data with support from the city and other jurisdictions. This approach would be roughly modeled after the one taken by the City of Pittsburgh and Alleghany County in association with Pitt and Carnegie Mellon University.

Syracuse's use of risk management also extends to institutional design with respect to the decision-making process around adopting AI tools. Their approach emphasizes being as democratic, accountable, and transparent as possible. While this includes relying on the City's Common Council to provide democratic oversight, the city is also creating an outside evaluation panel of experts. Part of that panel's role will include using the city's evaluative framework to make recommendations that balance the various desires, needs, costs, and benefits associated with these systems.

Recent events demonstrate the value prioritizing democratic and transparent decision making. In the midst of negotiating the terms for the 5G wireless data network, there were some 11th-hour political challenges in the form of residents contacting their Common Council representatives to express concerns about potential community harm due to 5G's purported carcinogenic effects.

This led the Common Council to postpone the vote to approve the contract. Despite these claims being dismissible at face value, evidence demonstrating that there are no known carcinogenic effects generated by microwave radiation at the 5G bandwidth range was marshalled and made available to the public during the extended deliberation. The council members were able to receive sufficient information that clarified that these purported health risks are not supported by scientific evidence, and thus assuaged, they were able to approve the deal. At the same time, the transparent process by which they were convinced proactively discredits any future suggestion that the city willfully ignored or covered up these risks.

Conclusion

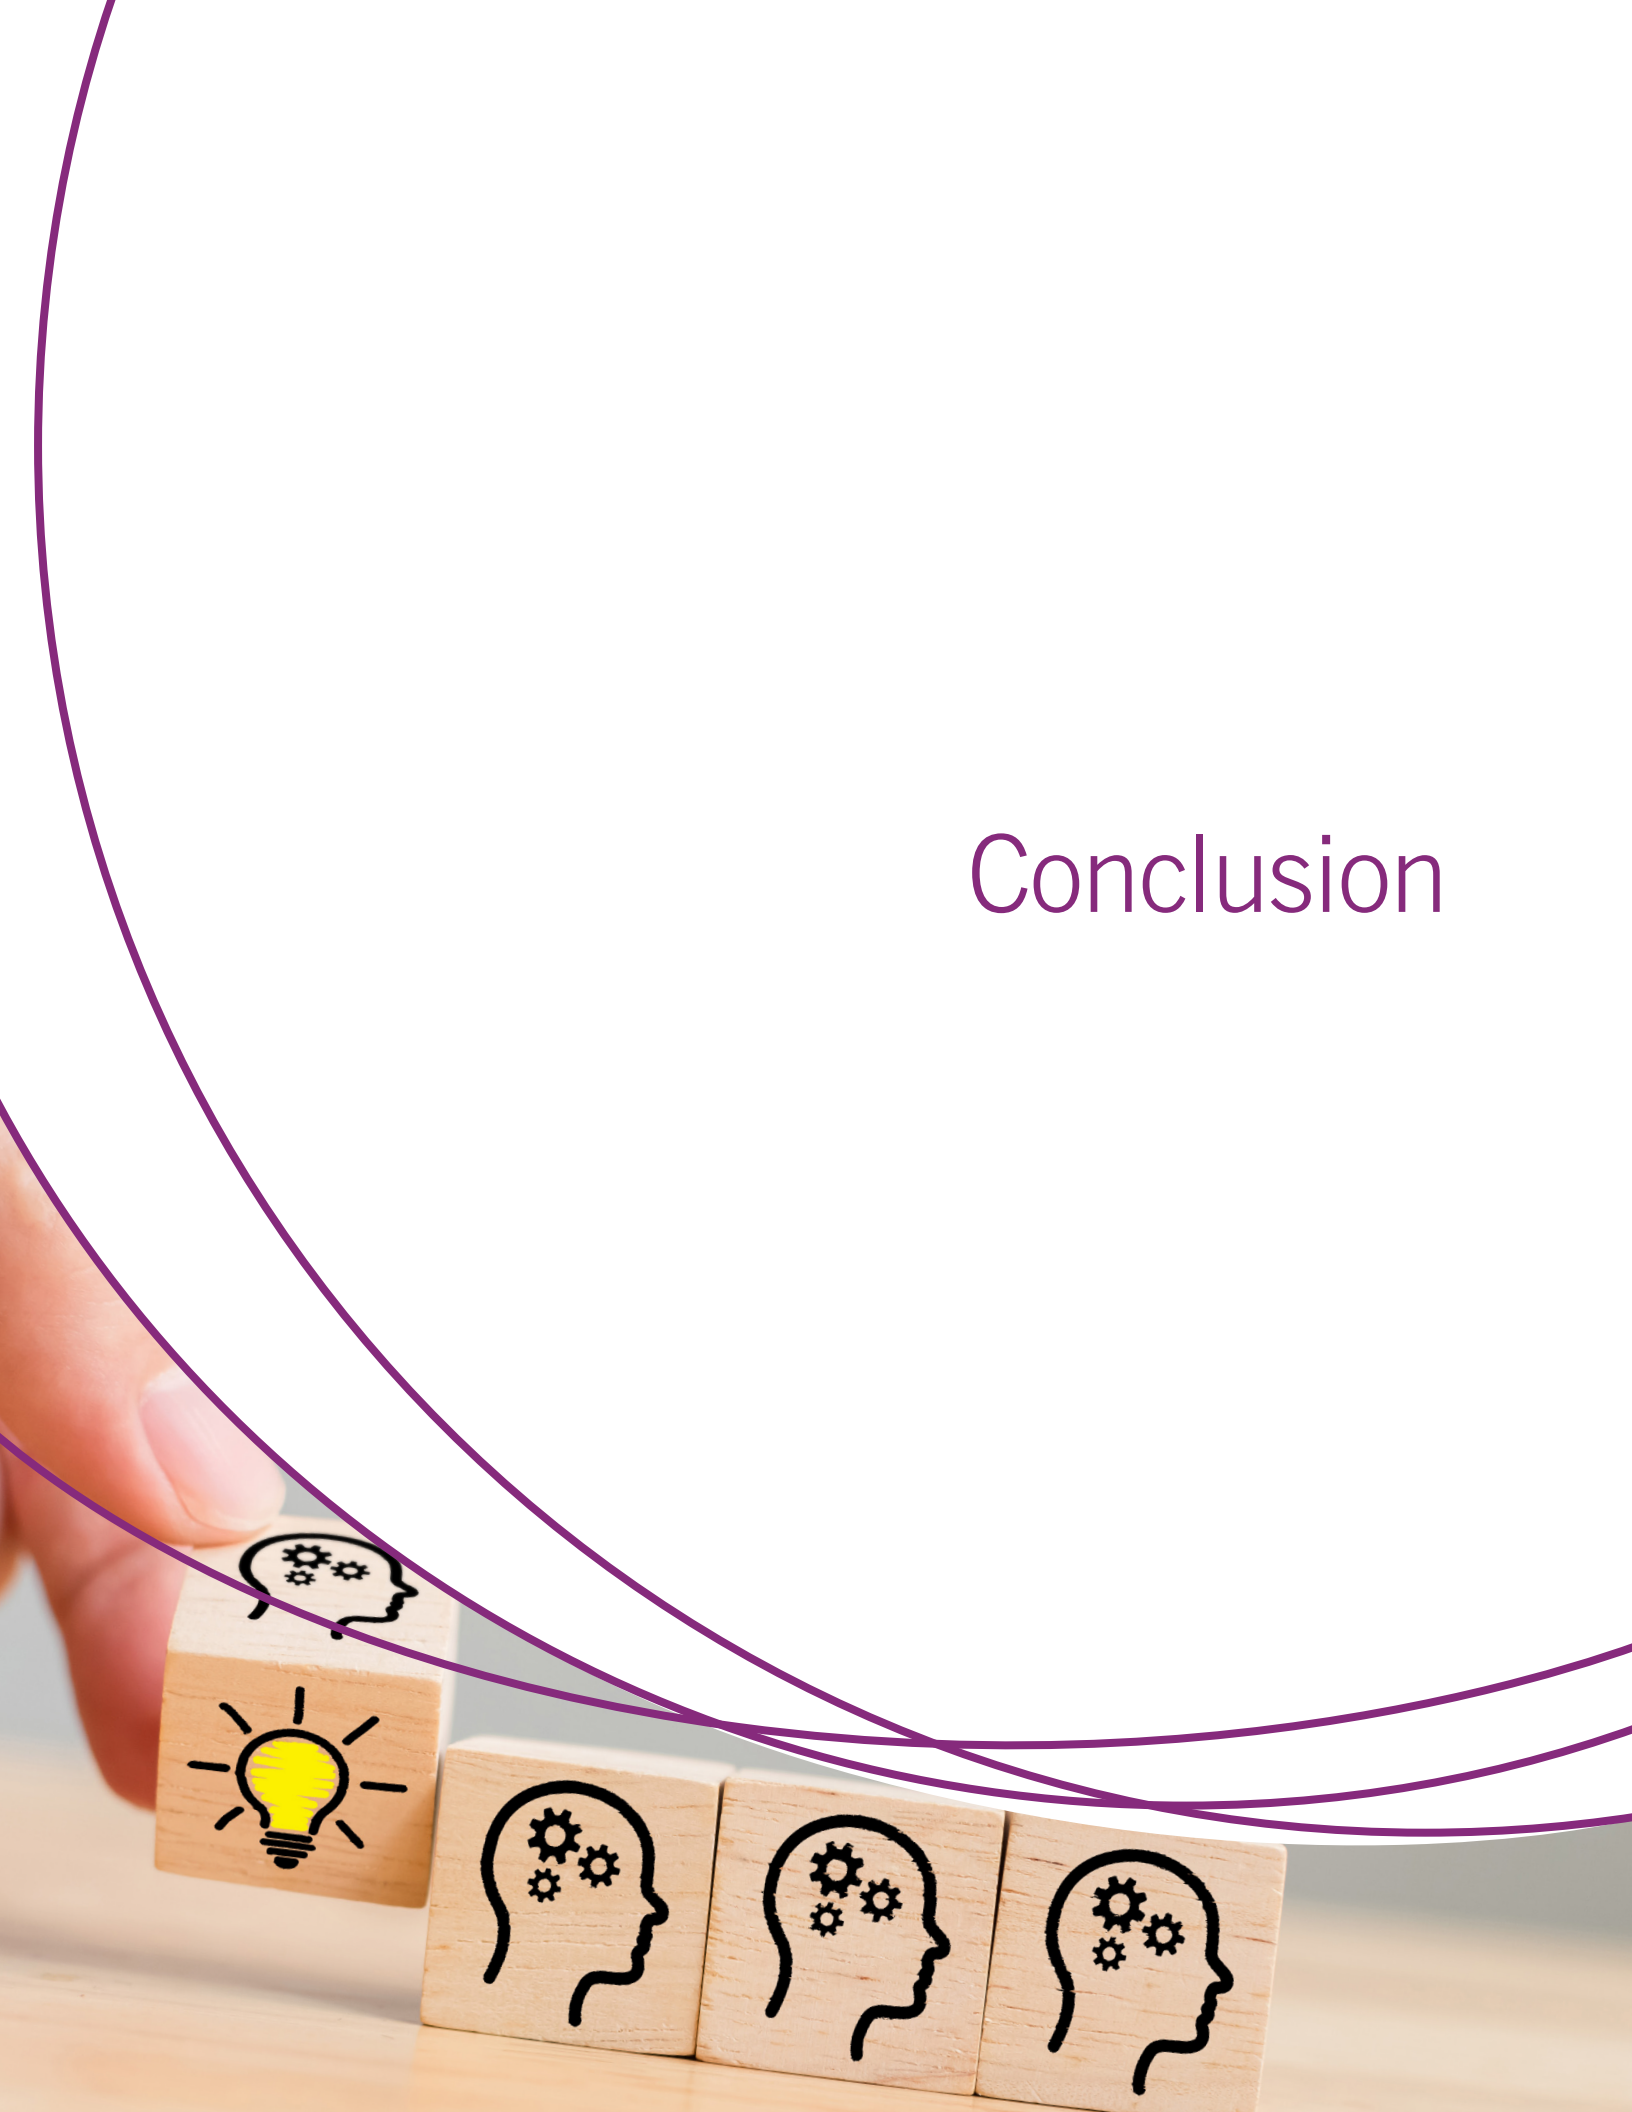

This report provides public managers with a framework for managing the risks associated with the public sector's entrance into the AI era. AI tools present both immense opportunities and challenges to good governance that should be considered and weighed carefully well before the adoption and implementation of the tools themselves. We began by reviewing the federal government's efforts to create a coherent risk management strategy that addresses AI's current and future capabilities. We trace these efforts to the 2016 National AI R&D Strategic Plan and the AI, Automation, and Economy report developed under the Obama administration. These efforts have subsequently been updated by the Trump administration in 2019 in the form of revisions to the National AI R&D Strategic Plan in 2019 and Executive Order #13859, which instructs federal agencies to implement the plan.

These has been followed by the introduction of the AI in Government Act of 2019 in Congress, which codifies the strategies and recommendations developed thus far, and establishes an AI Center of Excellence that would be responsible for guiding federal-level implementation. This section also highlighted a number of important strategies and recommendations for all government organizations to consider. Among these considerations are to continue investing in AI, improve data collection and data sharing, partner with a variety of stakeholders, prepare the workforce for the use of AI, develop technical standards that ensure safety, and only implement AI tools that carefully weight ethical concerns, privacy, and personal liberty.

We then introduce and discuss the scope of AI tools, and provide a framework developed by public administration scholars for understanding the potential use of AI in public sector administrative processes. We argue that AI tools can be utilized in three particular ways: (1) as a substitute for human discretion and labor, (2) as a decision support tool to find unseen patterns in large and multi-dimensional data, and (3) as an information generation and complexity elimination system. Depending on the nature of the task and the task environment, one of these three forms of implementation may be more useful or relevant than others, but they are not mutually exclusive. We then identify factors that will determine the potential risks and rewards from using these technologies, including the overall complexity of the task, the required degree of discretion, and the task's location within the organization and broader institutional environment. Finally, we introduce five evaluative criteria for public managers to use when developing their organization's risk management strategy for AI: effectiveness, efficiency, equity, manageability, and political legitimacy.

We also provided two simple illustrative cases of how local governments have carefully implemented a risk management strategy for the utilization of different types of AI tools. In the case of the City of Bryan, Texas, we present a successful risk management approach that includes a partnership across academia and a local government and proactively managing identified risks. In the case of Syracuse, New York, we describe a risk management approach that includes capitalizing on an otherwise unrelated opportunity, leveraging both external funding and relevant experts, and proactively managing identified risks.

To close, we distill our theoretical and empirical discussions of artificial intelligence tools in the public sector into recommendations for developing appropriate risk management processes. They are meant to help guide public managers as they consider, adopt, and implement AI tools to augment and automate tasks within their organizations. We base these recommendations on our analysis of existing federal policy, findings from the two case studies, and the existing academic research on innovation adoption in general and AI tools in particular. The recommendations are as follows:

Commit to up-front and ongoing investments in management capacity and analysis related to both AI and existing organizational processes.

The field of AI is evolving too rapidly to expect nonspecialists to stay at the cutting edge, but public managers must familiarize themselves with the fundamental and general concepts and capacities in order for any risk management approach to succeed. As part of this investment, consider dedicating resources—whether internal or contracted by a third party—to assess your current data and IT capacity, and determine what tasks or processes may be improved by incorporating AI tools. Build internal capacity by seeking out AI training for current employees. Stay informed on the latest AI tools, their cost, and their reliability.

One example of this approach is the adaptation of existing risk identification procedures as a required element of any AI adoption decision. Many organizational decisions that predate AI have long required risk review and compliance checks from technical and/or legal experts. Investing in adapting these review processes minimizes duplication of effort and additional complexity from adding wholly novel routines. The Canadian government formalizes this approach in its 2019 Directive on Automated Decision-Making. A second example is adding organizational capacity by hiring candidates from the rapidly growing set of data analytic specializations and dual-degree options in existing public administration and policy graduate programs. Higher education programs have broadly acknowledged both supply- and demand-side pressures for graduates with expertise not only in the traditional facets of public administration and policy analysis, but modern analytic techniques and data management practices.

The case of the City of Syracuse's infrastructure investments exemplifies the value of this approach. By deciding to continue to staff and fund the innovation team initially established through nonprofit grant funding, Syracuse retained the in-house expertise necessary to identify and leverage the potential for the city to leapfrog to next generation, AI-enabling infrastructure. The City of Bryan also highlights this approach through its ongoing partnership with world-leading engineers at Texas A&M University as they pilot and refine the use of self-driving trolleys in an iterative process that both researchers and City officials have committed to.

At the same time, these skills are also a critical need for more senior, established personnel who represent the current generation of decision makers. One approach to meeting this need is incorporating recognition of ongoing professional training and education on both the technical and managerial implications of AI into performance management and personnel review programs. This training could include executive education programs or even Massively Open Online Course (MOOC)-style asynchronous online options, but integrating whichever approach is deemed best to performance plans where appropriate for the position will help align managerial and employee incentives around developing internal capacity to understand and manage AI-related risks.

Manage risk by maximizing goodness of fit between tool and task.

Avoid the temptation to integrate AI tools into existing or new processes as a method of virtue signaling how future-ready your organization is. Take the time and effort to seriously consider the consequences of a task that is being considered for automation. Manage upfront risk by looking to tasks that are routine, repetitive, and contain quality data as pilot cases for AI tool augmentation. Estimate the improvements in both effectiveness and efficiency from imple-

menting AI tools into existing or new processes ex ante, weigh them against the corresponding risk costs, and use robust sensitivity analysis to challenge every modeling assumption involved. Carefully consider all privacy, ethical, and personal liberty concerns with the use of data and the deployment of AI tools.

As with the adaption of risk review processes, the Canadian government's Directive on Automated Decision-Making is a clear example of how to put this recommendation into practice. It identifies four discrete levels of impact arising from the use of AI in administrative decision making, ranging from those that have little likely and/or brief and reversible impact on the rights, well-being, and economic interests of individuals and communities, to those that will create irreversible, permanent impacts. Each impact level is then assigned a minimum standard of review from internal and external experts and review boards. In practice, this approach is interdependent with both the preceding and following recommendations. Organizations must have both internal expertise and knowledge of and relationships with external experts for such review processes to be effective.

The case of the City of Bryan self-driving trolleys also highlights a process of risk management and maximizing goodness of fit between tool and task. As the self-driving trolleys are being iteratively tested and improved, their routes throughout the city are kept short, do not cross major roadways, and maintain a low speed. This way, the trolleys—even in the case of a driving error or software malfunction—are operating at such a slow speed that it would be unlikely that the passengers or other vehicles would suffer serious injury. As the trolleys are further tested and developed, their routes will be expanded as well as their general benefits to the residents of the city.

Don't bowl alone: Leverage existing relationships and foster new ones to share knowledge.

The details of prior successes—and failures—of peer governments, organizations, and agencies with using AI tools are the most valuable asset for public managers looking to minimize their own risks with similar projects. Consider using existing professional associations (e.g., ICMA, etc.) to create working groups, email lists, and other channels for knowledge sharing. Draw on existing partnerships and pursue new ones with those outside the public sector in industry, academia, and related nonprofits. Use the knowledge gained through these networks to avoid needless risks. Carefully negotiate contracts with private providers of AI tools, be aware of their records of accomplishment, their respect for citizen privacy, and the ability to retain levels of expertise and manageability in house.

Both of the local governments profiled in this report's case studies made direct and substantial use of this approach through their coordination with local research universities. The City of Bryan worked closely and iteratively with experts and world-leading engineers from Texas A&M University, from project ideation through implementation. In the case of the City of Syracuse, its Office of Accountability, Performance, & Innovation collaborated with Syracuse University's School of Information Studies and College of Engineering and Computer Science, as well as private technology vendors, in identifying and planning for future uses of its developing Smart City and 5G Wireless infrastructure as part of the city's "Smart Surge" initiative.

Smaller local governments that do not have a geographic or historic affiliation with a college or university could look to establish new connections with the alma maters of their senior leadership. They should also look for opportunities to solicit national-level partners across the public, private, and nonprofit sectors that are motivated to extend their expertise beyond major metropolitan centers. The reciprocal of this example is also true: federal and state governments should engage with local and regional officials not only in disseminating risk management strategies, but in their development. The breadth of variation in initial conditions, capacities, and constraints across localities make them a tremendous asset for reducing uncertainty. This can be achieved by improving policymakers' ex ante understanding of both what needs to go right and what can go wrong as AI implementation becomes more diffused and commonplace.

This list is not meant to serve as a standalone risk management strategy for public sector adoption of AI. Rather, we intend it to provide public managers with a starting set of basic considerations informed by the latest field research and practice to serve as a first step in developing an effective risk management framework for their organizations.

ABOUT THE AUTHORS

JUSTIN B. BULLOCK

Justin B. Bullock is an assistant professor in the Public Service and Administration department and a research fellow in the Institute for Science, Technology and Public Policy. Bullock earned his PhD in public administration and policy with a specialization in public management and public policy from the University of Georgia in 2014, as well as a master's in public administration and a bachelor's in business administration, also from the University of Georgia.

Dr. Bullock has a number of interests at the intersects of public administration, public management, artificial intelligence, digital governance, and digital discretion. He is interested in the effects on governance, administration, and society of the rapid changes in decision making tools and data. He is also interested in what new opportunities, challenges, and questions these rapid changes pose for governance and society.

MATTHEW M. YOUNG

Matthew M. Young is an Assistant Professor of Public Administration and International Affairs at the Maxwell School at Syracuse University. His research interests include public management, innovation and technology adoption and use by the public sector, and service delivery. Dr. Young's work on in the public sector technological innovation has been published in top public administration and management journals, including *Public Administration Review, Public Management Review, and Perspectives on Public Management and Governance,* among others. He earned his PhD in Public Policy and Management as well as a Master of Public Policy at the University of Southern California Sol Price School of Public Policy, and a Bachelor of Arts degree in Political Science at the University of California at Berkeley. Before joining the academy, Dr. Young worked for more than a decade in Silicon Valley in software engineering, product management, and consulting.

KEY CONTACT INFORMATION

To contact the authors:

Justin B. Bullock

Assistant Professor,
Bush School of Government and Public Service
Texas A&M University

Phone: (979) 458-8032

jbull14@tamu.edu

Matthew M. Young

Assistant Professor of Public Administration
and International Affairs
Syracuse University

myoung10@syr.edu

REPORTS FROM THE IBM CENTER FOR THE BUSINESS OF GOVERNMENT

For a full listing of our publications, visit www.businessofgovernment.org

Recent reports available on the website include:

Agility:

Transforming How Government Operates: Four Methods of Change by Andrew B. Whitford
Agile Problem Solving in Government: A Case Study of The Opportunity Project by Joel Gurin and Katarina Rebello
Applying Design Thinking To Public Service Delivery by Jeanne Liedtka, Randall Salzman

Digital:

More Than Meets AI: Part II by the Partnership for Public Service and The IBM Center for The Business of Government
Financial Management for The Future: How Government Can Evolve to Meet the Demands of a Digital World by Angela Carrington, Ira Gebler
More Than Meets AI by The Partnership for Public Service and The IBM Center for The Business of Government
The Impact of Blockchain for Government: Insights on Identity, Payments, and Supply Chain by Thomas Hardjono
A Roadmap for IT Modernization in Government by Dr. Gregory S. Dawson
Delivering Artificial Intelligence in Government: Challenges and Opportunities by Kevin C. Desouza
Using Artificial Intelligence to Transform Government by Partnership for Public Service, The IBM Center for The Business of Government

Effectiveness:

Mobilizing Capital Investment to Modernize Government by Steve Redburn, Kenneth J. Buck, G. Edward DeSeve
Scaling Evidence-Based Programs in Child Welfare by Patrick Lester
Reducing Administrative Burden in Federal Research Grants to Universities by Lisa Mosley, Jeremy Forsberg, David Ngo
Emulating Value Chains of Consumer Goods to Save Lives: A Case Study of ColaLife's Work in Zambia by Paulo Savaget, Cassi Henderson, Steve Evans
Responding to Global Health Crises: Lessons from the U.S. Response to the 2014-2016 West Africa Ebola Outbreak by Jennifer Widner
Buying as One: Category Management Lessons From the United Kingdom by Anne Laurent
Improving the Delivery of Services and Care for Veterans by Matthew Hidek, Nathaniel Birnbaum, Nicholas Armstrong, Zachary S. Huitink

Insight:

Integrating Big Data and Thick Data to Transform Public Services Delivery by Yuen Yuen Ang
A Practitioner's Framework for Measuring Results: Using "C-Stat" at the Colorado Department of Human Services by Melissa Wavelet
Data-Driven Government: The Role of Chief Data Officers by Jane Wiseman
Integrating and Analyzing Data Across Governments—the Key to 21st Century Security by Douglas Lute, Frank Taylor

People:

Preparing the Next Generation of Federal Leaders: Agency-Based Leadership Development Programs by Gordon Abner, Jenny Knowles Morrison, James Perry, and Bill Valdez
Assessing the Past and Future of Public Administration: Reflections from the Minnowbrook at 50 Conference by Tina Nabatchi and Julia L. Carboni
Off to a Running State Capital Start: A Transition Guide for New Governors and Their Teams by Katherine Barrett and Richard Greene

Risk:

Managing Cybersecurity Risk in Government by Anupam Kumar, James Haddow, Rajni Goel

About the IBM Center for The Business of Government

Through research stipends and events, the IBM Center for The Business of Government stimulates research and facilitates discussion of new approaches to improving the effectiveness of government at the federal, state, local, and international levels.

About IBM Global Business Services

With consultants and professional staff in more than 160 countries globally, IBM Global Business Services is the world's largest consulting services organization. IBM Global Business Services provides clients with business process and industry expertise, a deep understanding of technology solutions that address specific industry issues, and the ability to design, build, and run those solutions in a way that delivers bottom-line value. To learn more visit ibm.com.

For more information:
Daniel J. Chenok
Executive Director
IBM Center for The Business of Government

600 14th Street NW
Second Floor
Washington, DC 20005
202-551-9342

website: www.businessofgovernment.org
e-mail: businessofgovernment@us.ibm.com

Stay connected with the IBM Center on:

or, send us your name and e-mail to receive our newsletters.

IBM Center for
The Business of Government