

ARTICLES

CONTRACTING FOR ALGORITHMIC ACCOUNTABILITY

CARY COGLIANESE* AND ERIK LAMPMANN**

I.	PROCUREMENT AS AI GOVERNANCE.....	181
II.	KEY ISSUES TO ADDRESS IN AI CONTRACTS.....	184
	<i>A. Transparency and Trade Secrets</i>	184
	<i>B. Data Privacy and Security</i>	189
	<i>C. Algorithmic Impact Statements and Audits</i>	192
	<i>D. Opportunities for Public Participation</i>	194
III.	SOCIAL PROCUREMENT IN THE ALGORITHMIC STATE	196
	CONCLUSION	198

At the local, state, and federal levels, governments increasingly use artificial intelligence (AI) tools to support decisionmaking about public services, government benefits, and legal punishments.¹ For instance, Penn-

* Edward B. Shils Professor of Law and Professor of Political Science, Director of the Penn Program on Regulation, University of Pennsylvania.

** Associate, Hogan Lovells US LLP

1. Cary Coglianese & Lavi M. Ben Dor, *AI in Adjudication and Administration*, BROOK. L. REV. (forthcoming 2021) (on file with authors). By the terms “artificial intelligence” or “AI tools,” we mean to encompass a broad class of computer programs or digital algorithms that travel under a variety of other names, such as machine learning, machine-learning algorithms, predictive analytics, algorithmic governance, and Big Data. This class of diverse tools can take many forms but each is linked by a common reliance on “an automated process of discovering correlations . . . between variables in a dataset, often to make predictions or estimates of some outcome.” David Lehr & Paul Ohm, *Playing with the Data: What Legal Scholars Should Learn About Machine Learning*, 51 U.C. DAVIS L. REV. 653, 671 (2017). Today, such tools can be found in ordinary smart phones or other off-the-shelf products. Although the issues we discuss here could technically apply to such out-of-the-box, mass-produced products, our principal focus is on custom digital products or other analytic services provided by consultants to meet specific agency needs, such as in identifying potential targets of inspections or enforcement actions, assisting with claims processing or other adjudications, or supporting

sylvania’s Allegheny County uses an AI assessment tool to prioritize follow-up responses to phone calls alleging child abuse and neglect.² AI also drives New York City’s “predictive policing” program by forecasting geographic locations of crime and supporting decisions about where to deploy police patrols.³ At the federal level, government agencies use AI tools to help identify possible cases of tax fraud or otherwise support regulatory enforcement. One study found that the federal government is using AI technology to drive over 150 projects at agencies as varied as the Securities and Exchange Commission and the Social Security Administration.⁴

These innovations have the potential to improve certain functions performed by the public sector by supporting more accurate, consistent, and timely decisions.⁵ The widespread use of AI by governments also raises a variety of accountability concerns. Some of these concerns arise because AI—also commonly referred to as machine learning—is driven by algorithms that are not intuitively easy to interpret.⁶ The relative opacity of machine-learning algorithms has earned them the moniker of “black box algorithms.”⁷ To the extent that pivotal decisionmaking relies upon what occurs only within a “black box”—that is, the internal workings of an algorithm—the public may look skeptically on governmental use of these

agency rulemaking. See generally Cary Coglianese & David Lehr, *Regulating by Robot: Administrative Decision Making in the Machine-Learning Era*, 105 GEO. L.J. 1147 (2017). Furthermore, although our focus in this Article is on government contracting, many of the same issues we discuss should be considered when private firms contract for AI services.

2. Dave Gershgorn, *Algorithms Can’t Fix Societal Problems—and Often Amplify Them*, QUARTZ (Oct. 17, 2018), <https://qz.com/1427159/algorithms-cant-fix-societal-problems-and-often-amplify-them/>. For context, see *The Allegheny Family Screening Tool*, ALLEGHENY CNTY., <https://www.alleghenycounty.us/Human-Services/News-Events/Accomplishments/Allegheny-Family-Screening-Tool.aspx> (last visited May 18, 2021).

3. The full extent of New York City’s “predictive policing” policies came to light in 2017 following a lawsuit by the Brennan Center for Justice. See Ali Winston, *Transparency Advocates Win Release of NYPD “Predictive Policing” Documents*, INTERCEPT (Jan. 27, 2018, 11:58 AM), <https://theintercept.com/2018/01/27/nypd-predictive-policing-documents-lawsuit-crime-forecasting-brennan/> (describing the full extent of the policies, including surveillance cameras and facial recognition technology).

4. See DAVID FREEMAN ENGSTROM, ET AL., GOVERNMENT BY ALGORITHM: ARTIFICIAL INTELLIGENCE IN FEDERAL ADMINISTRATIVE AGENCIES, ADMIN. CONF. U.S. (2020).

5. CARY COGLIANESE, A FRAMEWORK FOR GOVERNMENTAL USE OF MACHINE LEARNING, ADMIN. CONF. U.S. (2020).

6. Coglianese & Lehr, *supra* note 1, at 1156–60; Lehr & Ohm, *supra* note 1, at 655–56.

7. FRANK PASQUALE, THE BLACK BOX SOCIETY: THE SECRET ALGORITHMS THAT CONTROL MONEY AND INFORMATION 3 (2016).

tools.⁸ A lack of transparency over governmental uses of AI also exacerbates serious concerns about inaccuracies in algorithmic forecasts and the racial and gender biases that AI tools may perpetuate, especially when using data that already reflect such biases.⁹

Some advocacy organizations, such as the Electronic Privacy Information Center, have sued government agencies under the Freedom of Information Act, seeking details on the government's use of algorithms.¹⁰ Other litigants have filed lawsuits in both state and federal courts raising procedural due process claims and citing the lack of algorithmic transparency provided by government agencies.¹¹ Without question, agencies that choose to use AI tools need to be mindful of the possibility that their choices could later come under not just the spotlight of media attention but also the scrutiny of judicial review.

Some decisions by state and federal government officials to use AI technologies will also undoubtedly prompt calls for legislative oversight and standards. Local, state, and federal lawmakers have already entered the fray by proposing legislation that would address concerns about certain uses of AI tools.¹² In addition, the Administrative Conference of the United States

8. See, e.g., David Freeman Engstrom & Daniel E. Ho, *Algorithmic Accountability in the Administrative State*, 37 YALE J. REG. 800, 821 (2020); Ryan Calo & Danielle Keats Citron, *The Automated Administrative State: A Crisis of Legitimacy*, 70 EMORY L.J. 797 (2021). Even the use of machine-learning algorithms by private sector entities has created some alarm. See, e.g., PASQUALE, *supra* note 7 (offering a cautionary account of the use of machine-learning algorithms in different aspects of individuals' daily lives).

9. See, e.g., CATHY O'NEIL, WEAPONS OF MATH DESTRUCTION (1st ed. 2016); VIRGINIA EUBANKS, AUTOMATING INEQUALITY: HOW HIGH-TECH TOOLS PROFILE, POLICE, AND PUNISH THE POOR (2018); see also Joy Buolamwini & Timnit Gebru, *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification*, 81 PROC. MACH. LEARNING RES. 1, 8 (2018) (showing empirically that commercial facial recognition software programmed with AI is far less accurate when asked to identify women's faces or those of people of color).

10. See, e.g., *EPIC v. DOJ (Criminal Justice Algorithms)*, ELEC. PRIV. INFO. CTR., <https://epic.org/foia/doj/criminal-justice-algorithms/#foiadocuments> (last visited May 18, 2021) (cataloging documents released to the Electronic Privacy Information Center as the result of Freedom of Information Act litigation against the government concerning criminal justice algorithms).

11. See, e.g., *State v. Loomis*, 881 N.W.2d 749 (Wis. 2016); *State v. Rogers*, No. 14-0373, 2015 WL 869323 (W. Va. Jan. 9, 2015); *Malenchik v. State*, 928 N.E.2d 564 (Ind. 2010); *People v. Wakefield*, 107 N.Y.S.3d 487 (N.Y. App. Div. 2019); see generally Coglianese & Ben Dor, *supra* note 1, at 12–15.

12. Some cities, for example, have responded to concerns about the use of facial recognition tools driven by artificial intelligence. See BERKELEY, CAL., MUN. CODE CH. 2 § 99 (2020); OAKLAND, CAL., CODE OF ORDINANCES CH. 9 § 64 (2021); S.F. ADMIN. CODE CH. 19.B (2021); SOMERVILLE, MASS., CODE OF ORDINANCES CH. 9 § 25 (2020). The state of California has adopted related privacy legislation with implications for the use of algorithmic

(ACUS) has issued a statement intended to guide federal agencies when using AI,¹³ and the National Academy of Public Administration has issued a report offering its own guiding principles.¹⁴ President Donald Trump even signed a comprehensive executive order setting out standards for governmental AI—an order that, so far at least, President Biden has yet to revoke.¹⁵

As government agencies turn more frequently to the use of AI, the pressure for legislative or judicial action to govern the use of these tools may only grow. But a rush to do *something* to regulate the rapid growth of AI—and in particular its use by government—could result in painting with too

tools. See California Privacy Rights Act, CAL. CIV. CODE § 1798.185(16) (West 2021). For a catalog of state legislation related to artificial intelligence, see *Legislation Related to Artificial Intelligence*, NAT'L CONF. ST. LEG., <https://www.ncsl.org/research/telecommunications-and-information-technology/2020-legislation-related-to-artificial-intelligence.aspx> (last updated Jan. 17, 2021) (tracking relevant state legislation, including whether the proposed legislation was enacted). At the federal level, legislation targeting the use of digital algorithms has been proposed but not adopted. See, e.g., Algorithmic Accountability Act of 2019, H.R. 2231, 116th Cong. (2019) (proposing that the Federal Trade Commission require private companies “that use, store, or share personal information to conduct automated decision system impact assessments and data protection impact assessments”). Abroad, the European Union has adopted the EU General Data Protection Regulation (GDPR): Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (Data Protection Directive), OJ 2016 L 119/1.

13. ACUS, Agency Use of Artificial Intelligence, 86 Fed. Reg. 6612, 6616–18 (Jan. 22, 2021). One of the authors of this Article served as a consultant to the Administrative Conference of the United States (ACUS) on its artificial intelligence project and was closely involved in drafting the agency’s statement. See COGLIANESE, *supra* note 5.

14. Joseph J. Avery et al., *Using Artificial Intelligence to Improve the Fairness and Equity of Government Decision Making*, PRINCETON PROJECT COMPUTATIONAL L. (2020), https://www.napawash.org/uploads/Using_AI_to_Improve_the_Fairness_and_Equity_of_Government_Decision_Making.pdf. For other guidelines for governmental use of AI tools, see, e.g., IBM CTR. FOR THE BUS. OF GOV'T & P'SHIP FOR PUB. SERV., MORE THAN MEETS AI: ASSESSING THE IMPACT OF ARTIFICIAL INTELLIGENCE ON THE WORK OF GOVERNMENT 4 (2019) (presenting a framework for understanding government use of artificial intelligence); IBM CTR. FOR THE BUS. OF GOV'T & P'SHIP FOR PUB. SERV., MORE THAN MEETS AI: PART II: BUILDING TRUST, MANAGING RISK 6 (2019) (operationalizing that framework, albeit with only a passing mention of the impact of procurement policy); Chris Bousquet & Stephen Goldsmith, *The Right Way to Regulate Algorithms*, BLOOMBERG CITYLAB (Mar. 20, 2018, 11:47 AM), <https://www.citylab.com/equity/2018/03/the-right-way-to-regulate-algorithms/555998/> (laying out four “basic requirements” for local governments when considering using artificial intelligence to automate certain municipal responsibilities, such as “[s]har[ing] the motivation for using an algorithm” and “[e]xplain[ing] what data went into the model and why”).

15. Exec. Order No. 13,960, 85 Fed. Reg. 78,939 (Dec. 3, 2020).

broad a brush and sacrificing its potential value. AI does have real potential to improve government services or reduce administrative costs.¹⁶ Yet some critics seem to paint machine learning as inherently unfair or categorically condemn all use of AI by government agencies. Any policy proposal that would altogether ban or significantly curtail the use of AI by government would deny society the benefits that the responsible use of such technology promises.¹⁷ We propose instead, at least as a starting point, that government agencies leverage the procurement process more effectively to reassure the public that government is using algorithmic tools responsibly.¹⁸

Not many government agencies possess the in-house expertise to develop and implement algorithmic software, which means that they often resort to purchasing such services—including technology support—from third-party

16. See generally Jane E. Fountain, *The Virtual State: Transforming American Government*, 90 NAT. CIVICS REV. 241 (2001) (discussing benefits from governmental reliance on digital technologies).

17. See, e.g., JOSHUA NEW & DANIEL CASTRO, CTR. FOR DATA INNOVATION, HOW POLICYMAKERS CAN FOSTER ALGORITHMIC ACCOUNTABILITY 1 (2018) (expressing skepticism about proposals to regulate algorithmic accountability in ways that “would lead to less AI use, thereby hindering social and economic progress”).

18. Although strikingly little attention has been paid to the use of procurement as an AI governance tool, we cannot claim to be entirely alone in pointing to its importance. The statement issued by ACUS, for example, acknowledges that government agencies may need to procure AI services from third-party vendors. 86 Fed. Reg. at 6617. The World Economic Forum has issued a set of guidelines for procuring AI systems. WORLD ECON. F., GUIDELINES FOR AI PROCUREMENT 5 (2019) [hereinafter WEF GUIDELINES] (outlining broad global regulatory principles regarding artificial intelligence procurement resulting from an industry consultation). See also World Econ. F., AI Procurement in a Box: AI Government Procurement Guidelines (June 2020), http://www3.weforum.org/docs/WEF_AI_Procurement_in_a_Box_AI_Government_Procurement_Guidelines_2020.pdf. Some governments have started to issue their own AI procurement guidelines. U.K. Office for Artificial Intelligence, Guidelines for AI Procurement (June 8, 2020), <https://www.gov.uk/government/publications/guidelines-for-ai-procurement/guidelines-for-ai-procurement>; New South Wales Government, Sourcing an AI Solution (Sept. 4, 2020), <https://www.digital.nsw.gov.au/policy/artificial-intelligence-ai/user-guide/sourcing-ai-solution>. In the legal literature, we are aware of only two extended works that discuss procurement of AI in any depth. See Deirdre K. Mulligan & Kenneth A. Bamberger, *Procurement as Policy: Administrative Process for Machine Learning*, 34 BERKELEY TECH. L.J. 773, 781, 791 (2019) (describing the deficiencies of normal government procurement process for providing public input into and oversight of automated systems); David S. Rubenstein, *Acquiring Ethical AI*, 73 FLA. L. REV. __ (Oct. 1, 2020) (unpublished manuscript on file with the authors) (proposing “a set of concrete regulatory reforms to center ethical AI throughout the procurement process: from acquisition planning through market solicitation, negotiation, and contractual award”). In the present Article, we intend to focus on key themes and principles that ought to guide government officials when contracting for AI services.

vendors.¹⁹ The contracts between government agencies and third-party vendors thus provide an important tool for ensuring both the responsible design of AI tools and overall accountability of their use.²⁰ A more deliberate approach to government contracting for AI systems could help protect agencies that use AI tools in the event of subsequent litigation, while also providing much-needed safeguards for protecting the interests of individuals and organizations subjected to the use of such tools.

A contracting approach to algorithmic accountability holds several advantages. First, contracting for algorithmic accountability is a highly actionable and adaptable approach, one that can be implemented immediately. Procurement officers do not need to wait for legislative action to ensure that contracts for AI services contain provisions that promote algorithmic accountability, such as by providing for appropriate public access to basic information about algorithms' design and functioning. Government contracts could also impose a series of responsible AI standards on vendors, such as provisions calling for data privacy, bias detection, transparency, algorithmic validation, and cybersecurity measures. In short, careful drafting of contracts for AI services paired with suitably robust public input over the provisions to be included in these contracts can allow procurement officers to assure the public that agencies are using AI tools responsibly.

Second, unlike legislative efforts, which may require broad, sweeping regulatory language, a contracting approach can be tailored to meet the needs of specific agencies while balancing the accountability concerns posed by particular use cases. A contracting approach to algorithmic accountability is also more adaptable than a legislative approach; over time, agencies can experiment with new contract terms and gradually evolve procurement norms to map onto ever-changing technologies.

Finally, by imposing standards for transparency and responsible use of AI in public contracts, government agencies can advance a meaningful expressive purpose about responsible AI practices throughout the industry.²¹ When governments insist on contract terms that require robust practices related to algorithmic

19. Private vendors seem more than happy to meet this new demand. *See* Tod Newcombe, *Is Government Ready for AI?*, GOV'T TECH. (July/Aug. 2018), <https://www.govtech.com/products/Is-Government-Ready-for-AI.html> ("The vendor community is increasingly turning its attention to the government market, as far as AI is concerned.")

20. Throughout this Article, what we mean by accountability focuses on governmental entities' respect for important values implicated by their use of AI tools and well as these entities' ability to assure members of the public affected by AI tools, including litigants who seek judicial review, of the fair and responsible use of these tools.

21. OECD, *INTEGRATING RESPONSIBLE BUSINESS CONDUCT IN PUBLIC PROCUREMENT* 35 (2020).

accountability, vendors may well find that it becomes easier to shift to these practices for all of their clients. In this way, the expectations that governments insist upon in their procurement contracts can help set the bar for algorithmic accountability throughout the economy, promoting the diffusion of norms about responsible AI across both the public and private sectors.

I. PROCUREMENT AS AI GOVERNANCE

Public procurement is a lucrative part of the economy. In the United States, one recent estimate places the annual value of the federal government's procurement contracts at approximately \$500 billion.²² In addition, fifty state governments and many more municipal governments engage in their own extensive procurement of goods and services, making government procurement overall a substantial portion of the nation's economic activity in any given year.²³ Of course, only a fraction of government contracting today deals with AI tools. Nevertheless, the pervasiveness of contracting does mean that governments at all levels have established processes to ensure that government monies are spent wisely, vendors complete the work for which they were hired, and the procurement process is transparent and accountable to the public.²⁴ These same established processes can be used to help advance other social goals, including the responsible public sector use of AI tools.

At the federal level in the United States, vendors themselves already enforce accountability norms through an adversarial process known as a "bid protest."²⁵ This process allows "prospective bidders and offerors [to] challenge flawed solicitations before the [bidding deadline]."²⁶ It also permits "disappointed bidders or offerors [to] challenge flaws in the award process" after the award of a contract.²⁷ Bid protests are but one mechanism

22. Christopher R. Yukins, *The U.S. Federal Procurement System: An Introduction*, 3 PROCUREMENT L.J. 69, 69 n.1 (2017).

23. *Id.* at 69 (noting that the "procurement market is one of the largest in the world").

24. See NAT'L INST. OF GOV'T PURCHASING (NGIP); THE INST. FOR PUB. PROCUREMENT, THE PUBLIC PROCUREMENT GUIDE FOR ELECTED AND SENIOR GOVERNMENT OFFICIALS 1-2 (2016) (describing the importance of these mechanisms for the public officeholders allocating and administering government funds). For further background on federal and state procurement law, see Richard O'Duvall et al., *Public Procurement in the United States: Overview*, PRACTICAL LAW (Mar. 1, 2013); AM. BAR ASS'N SECTION OF PUB. CONT. L. & SECTION OF STATE & LOCAL GOV'T L., THE 2000 MODEL PROCUREMENT CODE FOR STATE AND LOCAL GOVERNMENTS (2000), https://www.americanbar.org/content/dam/aba/administrative/public_contract_law/pcl-model-02-2000-code-procurement.pdf.

25. See Yukins, *supra* note 22, at 87 (internal quotation marks omitted).

26. *Id.*

27. *Id.*

through which affected interests can keep government procurement processes in check. Overall, a focus on accountability is core to American procurement law.²⁸ Academics,²⁹ practitioners,³⁰ and members of the public³¹ have all voiced support for robust transparency requirements for public procurement processes. For this reason, procurement officers should presumably already understand, if not appreciate, the value of accountability and transparency when it comes to governmental use of AI.³²

Procurement officials should also already be familiar with the use of government contracting more generally as a tool to promote social goals—a practice sometimes called “social,” “green,” or “sustainable” procurement.³³ Social procurement seeks to ensure that government contractors do not discriminate, violate labor laws, harm the environment, or otherwise engage in undesirable business practices. Although procurement policies “were

28. See generally Steven L. Schooner, *Desiderata: Objectives for a System of Government Contract Law*, 11 PUB. PROCUREMENT L. REV. 103 (2002) (identifying and explaining pillars of U.S. procurement law).

29. See, e.g., Jennifer Jo Snider Smith, *Competition and Transparency: What Works for Public Procurement Reform*, 38 PUB. CONT. L.J. 85, 87 (2008) (emphasizing the importance of competition and transparency to a well-functioning procurement system); see generally Vinod Rege, *Transparency in Government Procurement*, 35 J. WORLD TRADE 489, 489 (2001) (exploring the impact of the World Trade Organization’s transparency proposals on developing countries).

30. See NGIP: INST. FOR PUB. PROCUREMENT, TRANSPARENCY IN GOVERNMENT; TRANSPARENCY IN GOVERNMENT PROCUREMENT 1 (2010) (providing observations on the importance of transparency in government contracting from the point of view of a trade association representing procurement professionals).

31. Public discussion about procurement even turned political under the Trump Administration, as procurement policy became a bit of a political football. See Steven Overly & Jacqueline Feldscher, *Trump Pledges to ‘Look’ at \$10B Pentagon Contract Amid Complaints About Amazon*, POLITICO (July 18, 2019, 5:57 PM), <https://www.politico.com/story/2019/07/18/trump-10b-pentagon-contract-amazon-1422011> (reporting that the White House was actively considering intervening in the Department of Defense’s decision to award an extremely large cloud computing contract to Amazon, a company owned by a political rival of the President); see also Dorothy Robyn, *A Disturbance in the Force: President Trump, DOD’s JEDI Procurement, and the Old Post Office*, BROOKINGS INST. (Aug. 23, 2019), <https://www.brookings.edu/blog/fixgov/2019/08/23/dods-jedi-procurement-and-the-old-post-office/> (covering the same set of events); Danny Vinik, *Trump’s \$440 Billion Weapon*, POLITICO (Dec. 22, 2016, 8:57 PM), <https://www.politico.com/agenda/story/2016/12/trump-federal-contracts-weapon-000262> (detailing the President’s stated interest in “the arcane, bureaucratic function of federal procurement”).

32. Cary Coglianese & David Lehr, *Transparency and Algorithmic Governance*, 71 ADMIN. L. REV. 1 (2019).

33. OECD, GOING GREEN: BEST PRACTICES FOR SUSTAINABLE PROCUREMENT (2015); Amy Ludlow, *Social Procurement: Policy and Practice*, 7 EUR. LAB. L.J. 479 *passim* (2016).

originally designed to achieve the best value for money by encouraging open competition” and fairness in bidding, over the years these policies have increasingly recognized that the public expects government officials to take “a full life-cycle approach, including the consideration of environmental and social costs.”³⁴ Local, state, and federal government procurement rules and programs have, for example, sought to advance social justice by promoting contracting with businesses owned by women and individuals from underrepresented racial communities.³⁵

Some research suggests that sustainable procurement has the secondary benefit of advancing social justice by diffusing social norms about best business practices throughout the private sector. For example, in California, cities that require their public construction contracts to comply with otherwise voluntary energy efficiency standards known as LEED see a greater proportion of private construction companies that also comply with LEED standards.³⁶ Even towns that surround these cities reportedly see an increase in LEED compliance.³⁷

Public procurement contracts at the local, state, and federal level could easily be revised to include provisions requiring compliance with otherwise voluntary standards for the responsible use of AI. For example, the professional engineering society, the Institute of Electrical and Electronics Engineers (IEEE), is developing a suite of voluntary standards that apply to both the technical and social aspects of AI development and deployment.³⁸ The organization AlgorithmWatch has compiled a list of over 150 different sets of voluntary standards or principles for the responsible use of AI.³⁹

34. OECD, *supra* note 33, at 19.

35. *Contracting Assistance Programs*, U.S. SMALL BUS. ADMIN., <https://www.sba.gov/federal-contracting/contracting-assistance-programs> (last visited May 18, 2021); Chris Burrell, *The Color of Public Money: Philadelphia and Massport Show Paths to Expanding Minority Contracts*, GBH NEWS (Jan. 4, 2020), <https://www.wgbh.org/news/local-news/2020/01/14/the-color-of-public-money-philadelphia-and-massport-show-paths-to-expanding-minority-contracts>; Courtney Buble, *Biden Administration Likely To Increase Contracting Opportunities for Small and Minority-Owned Businesses*, GOV'T EXEC. (Dec. 29, 2020), <https://www.govexec.com/management/2020/12/biden-administration-likely-increase-contracting-opportunities-small-and-minority-owned-businesses/171031/>.

36. Timothy Simcoe & Michael W. Toffel, *Government Green Procurement Spillovers: Evidence from Municipal Building Policies in California*, 68 J. ENV'T ECON. & MGMT. 411 *passim* (2014).

37. *Id.*

38. *Artificial Intelligence Systems (AIS) Related Standards*, IEEE STANDARDS ASS'N, <https://standards.ieee.org/initiatives/artificial-intelligence-systems/standards.html> (last visited May 18, 2021).

39. *About, AI ETHICS GUIDELINES GLOB. INVENTORY* (Apr. 2020), <https://inventory.algorithmwatch.org/about>.

Standards and principles such as these could be readily incorporated into public contracts for AI services, with procurement officials specifying compliance with such standards as an explicit contractual obligation of AI vendors.

II. KEY ISSUES TO ADDRESS IN AI CONTRACTS

We highlight four specific issues that government officials should consider when drafting contracts for AI tools and services: trade secrets, privacy and cybersecurity, auditing, and public participation. For each issue, we endeavor to provide a brief, actionable summary of relevant law as well as options for government officials to consider. We invite procurement officers and agency officials to take into account a suite of potential strategies that, if embedded in AI contracts, could meaningfully enhance algorithmic accountability in practice. Our aim is to offer suggestions—not necessarily to dictate precise contract terms. Which provisions should be added to new contracts, and the exact wording of any such contractual provision, will ultimately depend on agency practices, the technology at issue, the use to which it is being put, and the legal, policy, and even larger social and political context of a given agency and procurement process.⁴⁰

A. *Transparency and Trade Secrets*

One challenge for procurement officers involves striking the right balance between respect for vendors' legitimate trade secrets and the necessity of providing members of the public some form of visibility into automated government decisionmaking.⁴¹ For a piece of information to be a trade secret, it must not be widely known, its secrecy must generate commercial value, and it must be protected from disclosure.⁴² Naturally, technology and data analytics companies wish to keep information about their algorithms protected for competitive reasons.⁴³ By contrast, although government agencies might in some instances have their own reasons for keeping the inner workings of algorithms secret (such as when they are used for law enforcement purposes),⁴⁴ they will likely have much less reason to keep

40. COGLIANESE, *supra* note 5. See also Mulligan & Bamberger, *supra* note 18, at 812–813 (distinguishing between the transparency and reason-giving necessitated by “inward-facing” versus “public-facing” digital systems).

41. For more on the doctrine of trade secrets, see 18 U.S.C. § 1839(3) (defining “trade secrets”).

42. *Id.*

43. David S. Levine & Ted Sichelman, *Why Do Startups Use Trade Secrets*, 94 NOTRE DAME L. REV. 751, 755–770 (2019).

44. See, e.g., Joshua A. Kroll, et al., *Accountable Algorithms*, 165 U. PA. L. REV. 633, 658 (2017) (discussing potential for gaming government algorithms).

algorithmic secrets than will private firms. Indeed, in many cases, government agencies will often have an affirmative interest in making the inner workings of their algorithms open to the public so as to promote public confidence in these tools or to defend their use in court.⁴⁵ A tension over the transparency of algorithmic information will therefore often exist between the interests of government agencies and their vendors. Procurement officials need to be mindful of this potential conflict and proactively structure their contracts to include transparency mechanisms that can protect government interests while also respecting vendors' legitimate expectations to the secrecy of proprietary technologies.

If agencies do not take transparency into account when drafting contracts for AI services, they risk allowing private vendors to assert overly broad claims for trade secrets protection. Such broad claims have unfortunately been said too often to have "become the default way to protect algorithms."⁴⁶ To be sure, some protection against disclosure of proprietary details might well be valid, perhaps especially if full disclosure would entirely destroy a company's competitive advantage. But to claim that *any* disclosure, however modest, would destroy the company's value in its work product would overextend and abuse trade secret protection.⁴⁷ Yet vendors have made precisely these claims to shield their algorithms from discovery in lawsuits relating to local and state government eligibility determinations for social assistance benefits.⁴⁸ These vendors have with some consistency asserted trade secret protection and refused to allow plaintiffs to access information

45. Government agencies also have an interest in maintaining their own access to information regarding algorithms they hire vendors to program. See COGLIANESE, *supra* note 5, at 74 ("[W]hen contracting out for technical support and services in developing a machine learning system, agencies should take into account the need to have access to and be able to disclose sufficient information about the algorithm, the underlying data, and the validation results to satisfy transparency norms.").

46. Sonia K. Katyal, *Private Accountability in the Age of Artificial Intelligence*, 66 UCLAL REV. 54, 125 (2019).

47. Although no case presenting the question of the protectability of algorithmic information has reached the U.S. Supreme Court, the Court's decision in *Georgia v. Public.Resource.org, Inc.*, 140 S. Ct. 1498 (2020) might well point in the direction of default toward transparency. It suggests the possibility that any AI work performed under the supervision of government agencies that are related to or support authoritative decisions by those agencies will fall under the government edicts doctrine and not be protectable under copyright or perhaps other intellectual property laws. *Public.Resource.org, Inc.*, 140 S. Ct. at 1506.

48. See, e.g., *Hous. Fed'n of Teachers Local 2415 v. Hous. Indep. Sch. Dist.*, 251 F. Supp. 3d 1168 (S.D. Tex. 2017); *Michael T. v. Bowling*, No. 2:15-cv-09655, 2016 WL 4870284, at *1, *16 (S.D. W. Va. Sept. 13, 2016), *rev'd sub nom. Michael T. v. Crouch*, 2018 WL 1513295 (S.D. W. Va. Mar. 26, 2018); *K.W. v. Armstrong*, 298 F.R.D. 479 (D. Idaho Mar. 25, 2014).

about their algorithm—leaving the government vulnerable in due process litigation.⁴⁹

In *Houston Federation of Teachers Local 2415 v. Houston Independent School District*,⁵⁰ for instance, a school system used an algorithm created by a private firm as part of its teacher performance evaluation process.⁵¹ After the school district relied on the algorithm to determine whether to terminate teachers for poor performance, the local teachers' union filed suit raising due process objections. The court issued a preliminary decision allowing the plaintiff to take its case to a jury and emphasizing the lack of available information about the algorithm over which the private firm had claimed trade secret protection.⁵² The school district was thus unable to defend itself by providing the information that the plaintiff sought. The court explained that, without access to information the firm possessed about how the school district's algorithm operated, teachers lacked the ability to understand how the school board made decisions about their employment.

This tension between trade secrets and due process rights is hardly irresolvable. As it is, due process rights are far from absolute. Prevailing doctrine instead calls upon courts to balance competing factors when evaluating the adequacy of agency procedures under the Due Process Clause.⁵³ The choice, in other words, is not one between full disclosure and total trade secrecy. With attentiveness to the need for balance, procurement officers should be able to insist that vendors disclose information needed to assure the public and any litigants that government AI tools are working properly and fairly, even if firms wish to insist upon protecting their source code or other commercially valuable information.⁵⁴

Agencies can strike the right balance by requiring that their contractors disclose enough information to allow the public (and courts) to see what an AI tool has been designed to accomplish and to validate the tool's performance in meeting the desired objectives.⁵⁵ In particular, government

49. See Coglianesi & Ben Dor, *supra* note 1, at 29.

50. 251 F.Supp.3d 1168 (S.D. Tex. 2017).

51. *Id.* at 1171.

52. *Id.* at 1179 (“[W]ithout access to [the] proprietary information . . . [the teachers’ performance scores] will remain a mysterious ‘black box,’ impervious to challenge.”).

53. See *Mathews v. Eldridge*, 424 U.S. 319, 334–35 (1976) (providing the canonical articulation of this standard).

54. The notion of a type of “quasi” or “qualified” transparency of algorithms is discussed further by Kroll et al., *supra* note 44, at 641, and Frank Pasquale, *Beyond Innovation and Competition: The Need for Qualified Transparency in Internet Intermediaries*, 104 NW. UNIV. L. REV. 105, 163–64 (2010).

55. Coglianesi & Lehr, *supra* note 32, at 40, 42. This recommendation is also important for government procurement officers in Europe, where the law protects a “right to an

contracts should include language stipulating that contractors will provide information to show that “(a) an algorithmic system was constructed to advance a legally valid purpose by revealing the goal of an algorithm, (b) it is functioning correctly to advance that purpose (i.e., the program is not malfunctioning and it is producing validated results), and (c) it is being used as intended.”⁵⁶

Likewise, just because a private firm might desire to keep information secret does not mean that it is entitled to trade secret protection. Disclosing the goal of the algorithm—that is, its objective function—should surely never be treated as a trade secret, if only because the goal of the algorithm in question will likely have already been spelled out to some degree by the government in its request for proposals and will, in any case, never be the result of a firm’s business judgment.⁵⁷ Validation protocols and results presumably also should never reveal any trade secrets either, as they are simply methods used to assess how well the algorithm performs in relation to the stated objective, not how it was designed or developed.⁵⁸ Finally, it should also always be possible, without revealing any proprietary information, to provide affected individuals with their own specific data used by an AI tool to ensure it does not contain any errors. Procurement officers can thus justifiably specify in their contracts that all of this information will be disclosable by the contractor. Furthermore, such a contract could always provide that, for any information that is properly covered by trade secrets, the firm would consent

explanation” when algorithmic decision tools are used. *See* GDPR, *supra* note 12, at Recital 71 (making AI technologies “subject to suitable safeguards, which should include specific information to the data subject and the right to . . . an explanation of the decision reached after such assessment and to challenge the decision.”). European law also imposes other requirements or limitations on automated decisionmaking other than what the law currently does in the United States. *See id.* (“The data subject should have the right not to be subject to a decision . . . which is based solely on automated processing and which produces legal effects concerning him or her . . .”).

56. Coglianese & Lehr, *supra* note 32, at 47; *cf.* EUR. PARL., A GOVERNANCE FRAMEWORK FOR ALGORITHMIC ACCOUNTABILITY AND TRANSPARENCY, 2019, PE 624.262, at 58, [https://www.europarl.europa.eu/RegData/etudes/STUD/2019/624262/EPRS_STU\(2019\)624262_EN.pdf](https://www.europarl.europa.eu/RegData/etudes/STUD/2019/624262/EPRS_STU(2019)624262_EN.pdf) (indicating that trade secret protection should not extend to the existence of an algorithm or statements of the purposes for which it was developed).

57. We acknowledge that sometimes the full objective function could contain proprietary features, such as when it includes certain additions for a process of regularization. *See* Lehr & Ohm, *supra* note 1, at 704–05. But this does not preclude a government contract from requiring that the contractor disclose the core parts of the function related to the government’s goal or at least stipulating to in camera review, if litigation should arise. Coglianese & Lehr, *supra* note 32, at 49.

58. Coglianese & Lehr, *supra* note 32, at 49.

to allowing a court to review such information in camera in the event of litigation over the government's use of the algorithm in question.⁵⁹

Admittedly, the development and use of machine-learning algorithms produces information that could be solicited in litigation—including source code. But ultimately, government transparency doctrines—whether the procedural prong of the Due Process Clause or the Freedom of Information Act—are pragmatic, not absolute, in what they demand of government agencies to disclose.⁶⁰ In fact, that pragmatism has been on display in several cases already involving challenges to governmental use of algorithms. Some courts have taken a less demanding approach than the court did in the *Houston Federation of Teachers* case and have acknowledged that due process does not require the disclosure of *every* detail about an algorithm relied on by the government.⁶¹

In a case involving the use of a proprietary algorithm in criminal sentencing decisions, for example, the Wisconsin Supreme Court clarified that, even when a state's vendor merely makes publicly available “a full list of variables used” by its algorithm, but does not disclose other details about the algorithm, it does necessarily deprive any affected persons of their due process rights.⁶² That the court found limited disclosure passed constitutional muster is even more notable given the serious liberty interests at stake in any case where a criminal defendant has been subjected to a sentence based in part on the results of the algorithm.

Still, government agencies cannot always be sure how the courts will respond to demands by firms for trade secrets protection and due process challenges from members of the public affected by automated decisionmaking tools. Even apart from litigation, agencies ought to desire, for good government reasons, to provide as much transparency as they reasonably can over the pivotal algorithms they use. The way to do so is for government officials to demand that contractors waive any claims of trade secret privilege with respect to the types of information the government needs to demonstrate to the public that an algorithm is well-designed and fair.

59. *Id.* For a similar proposal to create a type of in camera review for algorithms, see Pasquale, *supra* note 54, at 164.

60. Coglianese & Lehr, *supra* note 32, at 34; *see also* Kroll et al., *supra* note 44, at 657 (characterizing full transparency of “source code as well as inputs and outputs for the relevant decisions” as “naïve”).

61. *See, e.g.*, *State v. Loomis*, 881 N.W.2d 749, 761 (Wis. 2016) (rejecting a due process challenge even though the state's contractor claimed its algorithm was “a proprietary instrument and a trade secret” and hence did “not disclose how the risk scores are determined or how the factors are weighed”).

62. *Loomis*, 881 N.W.2d at 760; *see also* Coglianese & Lehr, *supra* note 32, at 34.

Critics may argue that requiring procurement officers to negotiate such provisions asks too much of government officials while needlessly complicating straightforward procurement processes. However, like any contract negotiation, public procurement by design includes negotiations of all kinds—including over other issues of transparency or accountability. For instance, after obtaining a series of contracts for algorithmic services by open records requests, Robert Brauneis and Ellen Goodman concluded that in many cases government contracts for technological services already “do not . . . uniformly accede to contractor wishes for nondisclosure and data ownership.”⁶³ Brauneis and Goodman suggest that simply putting the onus on a vendor to redact specific information as protected by trade secrets helped in many cases decrease the number of overbroad assertions of trade secret protection.⁶⁴

Goodman has separately urged procurement officers to take responsibility for ensuring algorithmic transparency seriously, noting that government procurement officers have tremendous bargaining power in procurement transactions involving algorithms.⁶⁵ Moreover, interested members of the public may well be eager for these government officials to use that bargaining power to “push back on the universe of what is called propriety” by vendors.⁶⁶ Procurement officers should therefore negotiate for terms that allow the disclosure of the limited information needed by agencies to show both what their algorithms are supposed to do and why they are confident the algorithms work as intended.

B. Data Privacy and Security

AI implicates numerous data protection and privacy issues surrounding the collection, use, and disclosure of personal information. Contracting for algorithmic accountability is one way to manage the various privacy and security risks associated with compiling and working with the large databases used to train government algorithms and operate AI tools.⁶⁷

63. Robert Brauneis & Ellen P. Goodman, *Algorithmic Transparency for the Smart City*, 20 YALE J.L. & TECH. 103, 164 (2018).

64. *Id.*

65. Julia Powles, *New York City's Bold, Flawed Attempt to Make Algorithms Accountable*, NEW YORKER (Dec. 21, 2017), <https://www.newyorker.com/tech/annals-of-technology/new-york-citys-bold-flawed-attempt-to-make-algorithms-accountable>.

66. *Id.* Powles notes “[t]his is especially the case in New York, whose size, wealth, and high-quality demographic data make it a more desirable client than most cities.” She went on to note, “If New York doesn’t use that power to make systems accountable, who will?”

67. AI systems depend on access to vast databases. See generally Bernard Marr, *What is Deep Learning AI? A Simple Guide With 8 Practical Examples*, FORBES (Oct. 1, 2018, 12:16 AM),

To guard against the reputational, financial, and technical consequences of data breaches, government agencies using AI technologies should contractually require vendors to implement reasonable security measures and make their equipment and data available for routine conformity assessment with security standards.⁶⁸ Procurement officers may also insist on clauses requiring vendors to notify government officials in the event of such a breach and that vendors cooperate with government agencies in responding to any intrusion. The government should stipulate that it reserves the right to inspect proprietary information, such as algorithmic source code, following suspected or actual access by unauthorized threat actors.

Procurement officers would also do well to negotiate contractual guardrails on proper and improper uses and disclosures of personal information

<https://www.forbes.com/sites/bernardmarr/2018/10/01/what-is-deep-learning-ai-a-simple-guide-with-8-practical-examples/#19ebdb428d4b> (explaining “deep learning” as a process by which algorithms learn from a vast set of data). The existence of these databases almost invites attack—particularly when they contain personally identifying information. Attacks on massive public data sets appear increasingly common. *See, e.g.*, Alan Blinder & Nicole Perloth, *A Cyberattack Hobbles Atlanta, and Security Experts Shudder*, N.Y. TIMES (Mar. 27, 2018), <https://www.nytimes.com/2018/03/27/us/cyberattack-atlanta-ransomware.html> (reporting on a 2018 ransomware attack against the City of Atlanta, Georgia); U.S. Department of Homeland Security Office of Inspector General, Review of CBP’s Major Cybersecurity Incident During a 2019 Biometric Pilot (Sept. 21, 2020), <https://www.oig.dhs.gov/sites/default/files/assets/2020-09/OIG-20-71-Sep20.pdf> (reporting on the investigation of an unauthorized release of facial recognition data due to a cyberattack on a government subcontractor’s computer).

68. For examples of basic cybersecurity protections which vendors ought to employ, see FED. TRADE COMM’N, START WITH SECURITY: A GUIDE FOR BUSINESS, LESSONS LEARNED FROM FTC CASES (2015), <https://www.ftc.gov/system/files/documents/plain-language/pdf0205-startwithsecurity.pdf>. *See generally* Revision of OMB Circular No. A-130, “Managing Information as a Strategic Resource,” 81 Fed. Reg. 49,689 49,689 (July 28, 2016), <https://www.whitehouse.gov/sites/whitehouse.gov/files/omb/circulars/A130/a130revised.pdf> (detailing the Office of Management and Budget’s guidelines on information technology procurement, including requirements for securing personally identifiable information through measures such as access controls and data retention and deletion policies); *Cybersecurity Framework*, NAT. INST. FOR STANDARDS & TECH., <https://www.nist.gov/cyberframework> (last visited May 18, 2021) (describing the cybersecurity framework released by the Department of Commerce). Reports from the early months of the Biden Administration indicate that the President plans to release an executive order focused on cybersecurity, which may further affect which security measures are treated as “standard” by industry and government alike. *See, e.g.*, Jennifer Jacobs & Michael Riley, *Companies Must Quickly Report Hacks to U.S. Under Proposed Order*, BLOOMBERG (Mar. 31, 2021, 3:13 PM), <https://www.bloomberg.com/news/articles/2021-03-31/companies-must-report-hacks-to-u-s-within-days-in-draft-order>.

contained in relevant datasets.⁶⁹ Vendors should be expected to comply with all applicable federal laws and policies related to the handling of personal information.⁷⁰ Contractual provisions could also prohibit vendors from retaining personal information accessed while completing a government project for any other business purpose, even if that data is first de-identified.

AI tools have the ability to sort through different inputs (including live video feeds) and make forecasts about personal characteristics of individuals that governments should otherwise treat as private. One striking example of such profiling in the private sector involved a company that used AI to identify customers who were pregnant based on their purchase history and then sent these customers marketing materials about maternity products.⁷¹ The provisions of government contracts could expressly address such practices, prohibiting vendors from programming algorithms to make unnecessary forecasts about the personal characteristics of individuals who are subject to the algorithm's automated decisionmaking processes.⁷² Straightforward contractual safeguards like these could meaningfully insulate individuals' personal information from unauthorized access and disclosure, shore up public goodwill, and protect agencies against future privacy-related controversies.

69. Government contracts could also address what rights individuals affected by the use of automated tools will have to access their personal information from vendors and to dispute or correct any errors in such data.

70. The federal government is subject to at least two federal privacy laws constraining their use of personally identifiable information. E-Government Act of 2002, Pub. L. No. 107-347, §§ 511, 523, 116 Stat. 2899, 2965, 2968 (codified as amended at 44 U.S.C. § 3501); Privacy Act of 1974, 5 U.S.C. § 552(a). For additional guidance on privacy considerations, procurement officers might look to various federal legislative proposals on privacy. For an overview and comparative analysis of these proposals, see generally ELEC. PRIV. INFO. CTR., GRADING ON A CURVE: PRIVACY LEGISLATION IN THE 116TH CONGRESS (2019–2020), <https://www.epic.org/GradingOnACurve/EPIC-GradingOnACurve-Apr2020.pdf>. See also SEN. FEINSTEIN, ET AL., PRIVACY AND DATA PROTECTION FRAMEWORK 1 (2019), https://www.democrats.senate.gov/imo/media/doc/Final_CMTE%20Privacy%20Principles_11.14.19.pdf.

71. Charles Duhigg, *How Companies Learn Your Secrets*, N.Y. TIMES (Feb. 16, 2012), https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html?pagewanted=1&_r=1&hp. Although this example involves a private company rather than government agency, it dramatizes the type of arguably lawful but still invasive algorithmic processing that government contracts should anticipate.

72. The privacy of data subjects whose personal information is subject to automated decisionmaking is an important topic. See, e.g., EUR. COMM., ETHICS GUIDELINES FOR TRUSTWORTHY AI: HIGH-LEVEL EXPERT GROUP ON ARTIFICIAL INTELLIGENCE (2019), <https://ec.europa.eu/futurium/en/ai-alliance-consultation> (“Digital records of human behavior may allow AI systems to infer not only individuals’ preferences, but also their sexual orientation, age, gender, religious or political views.”).

C. Algorithmic Impact Statements and Audits

Procurement officers should also consider whether to negotiate contractual provisions that require vendors to follow substantive or process-based standards related to the responsible design and implementation of AI tools, whether those standards are developed by the contracting agency or by voluntary standards organizations, such as the IEEE.⁷³ If government contracts do make such standards contractually binding, they should also subject contractors to a requirement for regular auditing to assess conformity with the standards. After all, contractual terms are of little value if contractors do not adhere to them.⁷⁴

For instance, a procurement contract might require that vendors develop algorithmic “impact statements” or risk management plans.⁷⁵ Such management-based requirements would compel vendors to pause, reflect on possible adverse consequences of a particular course of action, and find ways to minimize those adverse effects.⁷⁶ Impact statements are widely used to address other types of concerns, such as the impacts of industrial pollution on the environment.⁷⁷ Adapting the impact statement framework to AI technology could involve government agencies requiring their vendors to agree to: (1) articulate the purpose of any contemplated use of artificial intelligence; (2) identify potential technical, transparency, equity, and other

73. Various industry groups and professional societies, including the Association for the Advancement of Artificial Intelligence, the Institute of Electrical and Electronics Engineers, and the Association of Computing Machinery (ACM), have developed or are developing ethical frameworks for technologists. See, e.g., ASS'N FOR COMPUTING MACH. U.S. PUB. POL'Y COUNCIL, STATEMENT ON ALGORITHMIC TRANSPARENCY AND ACCOUNTABILITY 2 (2017), http://www.acm.org/binaries/content/assets/public-policy/2017_usacm_statement_algorithms.pdf. The seven principles proposed by the ACM even include recommendations on a “right to an explanation” as well as a need to create software that can be easily and independently audited. *Id.*

74. See generally Kevin McGillivray, *FedRAMP, Contracts, and the U.S. Federal Government's Move to Cloud Computing: If an 800-Pound Gorilla Can't Tame the Cloud, Who Can?*, 17 COLUM. SCI. & TECH. L. REV. 336, 363–64 (2016) (chronicling federal government missteps in the acquisition of cloud computing solutions and recommending changes to the procurement process—namely through the use of more detailed standardized contractual clauses and more engaged internal compliance processes).

75. See, e.g., Cary Coglianese & David Lazer, *Management-Based Regulation: Prescribing Private Management to Achieve Public Goals*, 37 L. & SOC. REV. 691 (2003) (discussing regulatory strategy according to which entities are required to engage in their own risk management).

76. See Katyal, *supra* note 46, at 112 (summarizing existing applications of the impact statement framework).

77. *Id.* at 112.

problems to avoid; and (3) describe ways the vendor plans to address those problems and manage any risks associated with the use of the algorithm.⁷⁸

To help ensure vendors take impact statement requirements seriously, their statements, management plans, and all other documentation produced in conformity with the contract provisions should serve as the foundation for third-party or agency-conducted audits of government AI systems. Procurement officers should ensure that AI contracts call for such audits to occur at regular intervals—perhaps annually or semiannually—and that vendors fully cooperate with the auditors.⁷⁹

Because algorithms, especially those used by government entities, implicate a variety of concerns beyond simple mathematical accuracy, audits will likely need to be interdisciplinary, taking into consideration not only statistical analysis but also ethical considerations.⁸⁰ Thorough audits will demand time and expense, but, with greater experience, AI auditors should be expected to develop more routinized protocols to address the ethical considerations. An AI audit industry may even emerge in time, one “with proper credentialing, standards of practice, disciplinary procedures, ties to academia, continuing education, and training in ethics, regulation, and professionalism.”⁸¹

78. For analysis of what an impact statement might look like in the context of algorithmic accountability, see generally *id.* at 112–15. For another proposal on public agency algorithmic impact assessments, see DILLON REISMAN, ET AL., AI NOW, ALGORITHMIC IMPACT ASSESSMENTS: A PRACTICAL FRAMEWORK FOR PUBLIC AGENCY ACCOUNTABILITY 4 (2018), <https://ainowinstitute.org/aiareport2018.pdf> (identifying five elements which ought to be included in algorithmic impact assessments, such as “conduct[ing] a self-assessment of existing and proposed automated systems, evaluating potential impacts on fairness, justice, bias, or other concerns . . .” and “provid[ing] notice to the public disclosing their definition of ‘automated decision system[s],’ . . . and any related self-assessments and researcher review processed before the system has been acquired”).

79. Government officials and their auditors need to be mindful of the risk that impact statements and their audits could turn into exercises of “pencil whipping” or “Potemkin villages”—that is, with vendors simply going through the motions without attending seriously to potential problems with their algorithms. See NAT. ACAD. OF SCI., ENG’G, & MED., DESIGNING SAFETY REGULATIONS FOR HIGH-HAZARD INDUSTRIES (2018); Garry C. Gray & Susan S. Silbey, *Governing Inside the Organization: Interpreting Regulation and Compliance*, 120 AM. J. SOCIO. 96 (2014). See generally Cary Coglianese & Jennifer Nash, *Compliance Management Systems: Do They Make a Difference?*, in CAMBRIDGE HANDBOOK OF COMPLIANCE (D. Daniel Sokol & Benjamin van Rooij, eds.) (forthcoming 2021).

80. James Guszczka et al., *Why We Need to Audit Algorithms*, HARV. BUS. REV. (Nov. 28, 2018), <https://hbr.org/2018/11/why-we-need-to-audit-algorithms>.

81. *Id.* For discussion of codes of conduct and self-regulation within the AI community, see Katyal, *supra* note 46, at 108–11.

Requiring audits to ensure compliance with industry standards is far from a novel concept. Auditing of financial statements, for example, is a well-accepted good business practice.⁸² The rationale behind regular financial auditing has a certain symmetry with the rationale for AI audits: just as algorithms can be opaque, “companies’ internal operations appear as ‘black boxes’ to those on the outside.”⁸³ Requiring regular, independent audits of the ways algorithms actually work could help provide government agencies (and the members of the public that they serve) greater assurances that vendors are designing and deploying AI responsibly.⁸⁴ Although even the most thorough audit framework does not guarantee compliance, auditing is nevertheless valuable, if only to the extent it can help improve accountability and strengthen public trust in government use of AI.

D. Opportunities for Public Participation

In designing and implementing AI tools, as well as developing contracts about them, agency officials and procurement officers have a variety of consultative processes at their disposal to draw on the knowledge of expert and non-expert members of the public.⁸⁵ Leveraging these processes could assist both procurement officers tasked with navigating the field of AI as well as help agency officials and vendors improve their design and use of AI tools.

When government procures AI services, its contracts should thus consider providing for vendors to cooperate with appropriate opportunities for public participation. Public participation can be useful at the beginning of the design and development of algorithmic tools as well as during ongoing evaluation.⁸⁶ These tools demand that someone make choices about their goals and the

82. Guszczca et al., *supra* note 80.

83. *Id.*

84. *Id.* Moreover, independent auditors could “provide reasonable assurance that the reports coming from the ‘black box’ [of the algorithm] are free of material misstatement.” *Id.* Of course, we have no illusions that audits by themselves, especially if they are paid for by the contractors who are being audited, will prove a panacea. See Max H. Bazerman, George Loewenstein, & Don A. Moore, *Why Good Accountants Do Bad Audits*, HARV. BUS. REV. (2002), <https://hbr.org/2002/11/why-good-accountants-do-bad-audits> (describing auditors’ incentives to produce favorable audits and explaining how unconscious bias impacts audits); Coglianesi & Nash, *supra* note 79 (reporting research showing compliance management systems’ limited discernible impact).

85. For a general discussion of different participatory processes and their potential benefits, see OECD, FOCUS ON CITIZENS: PUBLIC ENGAGEMENT FOR BETTER POLICY AND SERVICES (2019), <https://www.oecd.org/governance/regulatory-policy/focusoncitizenspublicengagementforbetterpolicyandservices.htm>.

86. See Mulligan & Bamberger, *supra* note 18, at 845–850.

various tradeoffs that they can present, such as between accuracy and equity.⁸⁷ These are choices that data scientists alone cannot answer; government officials, informed by suitable modes of public engagement, must be involved. By negotiating for contractual terms that require vendors to participate in public participation processes, government agencies can help ensure that AI tools better match the values and priority of the public and that they enjoy greater public acceptance when put into use.⁸⁸

Contracts might provide, for example, for vendors to cooperate with agencies in providing “notice-and-comment” opportunities about key choices to be made in designing new AI tools. Contractual provisions could also call for vendors to become involved with members of the public in still more sustained, ongoing discussions around the use of AI. For instance, the United Kingdom has experimented with the use a “‘citizens’ jury’ to explore the use [of] AI to make, or help make, decisions.”⁸⁹ A citizens’ jury approach could take a variety of forms—from one-off consultative meetings with focus groups to a more sustained series of deliberative dialogues between a randomly-selected, representative group of participants, expert technologists, and policymakers.⁹⁰ By anticipating in advance, during the contracting stage, these possibilities for public engagement, government officials can ensure that their vendors will view their participation in such sessions as integral to their work.

Recommendations from citizen juries such as these need not be treated as binding to provide meaningful feedback from the public.⁹¹ The City of Bos-

87. Sandra G. Mayson, *Bias In, Bias Out*, 128 YALE L.J. 2218, 2248–2250 (2019).

88. Of course, creating robust public engagement protocols presupposes the existence of an informed, curious, and critical public. See generally Jakko Kemper & Daan Kolkman, *Transparent to Whom? No Algorithmic Accountability Without a Critical Audience*, 22 INFO., COMM. & SOC’Y 2081 (2019) (emphasizing the need for a critical audience in order for many algorithmic accountability schemes to function as designed).

89. ROYAL SOC’Y FOR THE ENCOURAGEMENT OF ARTS, MFRS., & COM., ARTIFICIAL INTELLIGENCE: REAL PUBLIC ENGAGEMENT (2018), https://www.thersa.org/globalassets/pdfs/reports/rsa_artificial-intelligence---real-public-engagement.pdf.

90. According to proponents of the longer-term, sustained citizens’ jury, this type of effort is necessary because of “the normative nature of policymaking and, thus, the need for integrating deliberative dialogue in governance alongside empirical analysis and logical reasoning.” *Id.* at 18. For a review of different modes of public participation, see Michael Sant’Ambrogio & Glen Staszewski, *Democratizing Rule Development*, 98 WASH. UNIV. L. REV. 793 (2021).

91. The goal of public engagement also need not be to find unanimity among members of the public, which may be impossible. See Cary Coglianese, *Is Consensus an Appropriate Basis for Regulatory Policy?*, in ERIC ORTS AND KURT DEKETELAERE, EDS., ENVIRONMENTAL CONTRACTS: COMPARATIVE APPROACHES TO REGULATORY INNOVATION IN THE UNITED STATES AND EUROPE 93–113 (2001) (discussing the pitfalls of public participation processes that seek unanimity, or near unanimity).

ton, for example, launched a major AI-driven reconfiguration of the city's school bus schedules, only to encounter major public backlash against the AI-generated schedules, which were not sufficiently designed to optimize the values held by the city's parents and students.⁹² The city would have been well-served if it had organized listening sessions or focus groups with parents and students—and if it had included provisions in its contract to have the technologists who developed the new schedules assist in organizing such sessions, or at least to participate in them.

III. SOCIAL PROCUREMENT IN THE ALGORITHMIC STATE

Industry, government, and the broader public are grappling with both the promise and the problems of the ever-changing world of AI.⁹³ Most observers would seem to agree that machine learning offers government a great potential for improved efficiencies,⁹⁴ even if they disagree on whether and how that potential should be harnessed. One way to govern the use of AI tools by government would be to establish overarching rules. But rules can be blunt instruments. And fears about the dangers of unaccountable algorithms could well lead to undue restrictions and even blunter bans altogether on governmental use of machine-learning algorithms.

In 2017, for example, a member of the New York City Council introduced a bill that would have required every vendor who built an algorithm for the city government to make the underlying source code publicly available.⁹⁵ In short order after the bill's introduction, private firms objected, citing potential harm to their competitive advantage and urging an incremental approach to transparency. Policy experts raised the specter of cybersecurity vulnerabilities and warned against “giv[ing] bad actors an easy way to game the public-benefits system.”⁹⁶ Even city administrators declared as a non-starter the adoption of an overly broad requirement that AI firms publish all

92. Ellen P. Goodman, *Defining Equity in Algorithmic Change*, REG. REV. (Feb. 11, 2019), <https://www.theregreview.org/2019/02/11/goodman-defining-equity-algorithmic-change/>; Ellen P. Goodman, *Smart Algorithmic Change Requires a Collaborative Political Process*, REG. REV. (Feb. 12, 2019), <https://www.theregreview.org/2019/02/12/goodman-smart-algorithmic-change-requires-collaborative-political-process/>.

93. See generally Lehr & Ohm, *supra* note 1 (seeking to provide lawyers and technologists with a shared vocabulary for discussing machine learning and technology policy). For additional information regarding AI and machine learning policy developments after this Article has gone to print, see generally *Algorithmic Transparency: End Secret Profiling*, ELEC. PRIV. INFO. CTR., <https://epic.org/algorithmic-transparency/> (last visited May 18, 2021).

94. Coglianese & Lehr, *supra* note 1, at 1160.

95. Powles, *supra* note 65.

96. *Id.*

of their proprietary information on an algorithm.⁹⁷ A spokesperson for the mayor, for example, said that “[p]ublishing the proprietary information of a company with whom we contract would not only violate our agreement, it would also prohibit other companies from ever doing business with us, which would prevent us from trying innovative solutions to solve everyday problems through technology.”⁹⁸ What eventually followed was a more balanced approach: the creation of a task force charged with considering, in part, how to “mak[e] technical information [concerning algorithms . . . publicly available where appropriate.”⁹⁹

Our suggestion in this Article also pursues a more balanced approach—and one that is ultimately incremental. We suggest that governments give greater attention to how they design and structure their contracts for services to develop and operate AI tools. Using contracting as a tool for algorithmic governance can allow governments and society to benefit from the improvements that AI tools can offer, while also helping ensure that these tools will be designed and deployed responsibly. The public deserves to know about the algorithms that affect their lives and interests, and public knowledge about algorithms—or even just the potential for litigation seeking to review an algorithm—can itself provide some constraint on ill-considered and unfair algorithmic practices. After all, if AI vendors can operate in total secrecy, protected by trade secrets, and never expected to meet basic standards for data security, privacy, algorithmic fairness, or public participation, then algorithmic accountability is nothing but a myth. For this reason, procurement officers and the government officials they serve need to ensure that proper contractual duties and restraints are imposed on vendors’ development and deployment of algorithmic tools.

Admittedly, the contracting process can already become protracted, and adding more terms to negotiate may only make procurement more cumbersome. But when the stakes for society’s and individuals’ interests are high, the extra time and effort needed to craft suitably protective standards can certainly be worth the investment.¹⁰⁰ As they move toward automating

97. *Id.*

98. *Id.*

99. *Id.* The Task Force released its report in November 2019. N.Y.C. AUTOMATED DECISION SYS. TASK FORCE, AUTOMATED DECISION SYSTEMS TASK FORCE REPORT 1 (2019), <https://www1.nyc.gov/assets/adstaskforce/downloads/pdf/ADS-Report-11192019.pdf>. The final report included a set of recommendations called for by the original City Council bill. *Id.* at 17.

100. A draft regulatory proposal released by the European Union in April 2021 accords with this notion of proportionality to governance when the stakes are high, as it proposes to place heightened standards and controls on “high-risk AI-systems.” *Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence*

a wider array of consequential governmental functions, agencies would do well to spend additional time and thought during the procurement process to consider the ramifications of new AI tools and their intended uses—as well as their potentially unintended consequences.

CONCLUSION

As governments come to rely increasingly on AI tools to support tasks and functions that affect individuals and organizations, the demands for algorithmic accountability will only grow. Important concerns have already emerged about the fairness and transparency of AI technologies used by some governmental authorities, and it is evident that these concerns can be exacerbated when government agencies contract out for the design, testing, and operation of AI tools. Private contractors may possess the analytic capacity that government needs for developing and running AI systems, but private firms' connections with the public will surely be more attenuated than will public agencies' and their motives may not align as well with the delivery of public value. Private vendors also tend to prefer to conduct their work with less oversight and disclosure, often claiming trade secret protection over their algorithmic tools.

Nevertheless, government contracting itself can operate as an important, tractable governance strategy. To use AI tools responsibly, agencies should seek to contract responsibly for the support and technology need to create such tools. Specifically, agency officials and procurement officers should attend to four key issues. First, they should ensure that AI contracts are drafted to ensure sufficient public transparency and to prevent vendors from claiming trade secret protection over all of their work. Second, government contracts should obligate AI vendors to follow accepted privacy and security protocols—and to allow the government to access information needed to ensure those protocols are followed. Third, agencies should consider negotiating contracts that include substantive standards for responsible AI and that insist vendors follow procedures, such as periodic audits, to document their compliance with such standards. Finally, whenever agencies anticipate the need for public participation to inform the design and operation of AI tools, their AI contracts should obligate private vendors to cooperate in the agency's process of public engagement.

Contracting for algorithmic accountability is an immediately feasible strategy for governing a rapidly evolving and highly varied set of technological innovations. Government contracts can be designed and

(Artificial Intelligence Act) and Amending Certain Union Legislative Acts, COM (2021) 206 final (Apr. 21, 2021).

adapted so that they address the important, practical needs that lead government agencies to develop AI tools, while also respecting society's desire for public accountability and engagement. By using the procurement process to achieve greater algorithmic accountability, public officials can help provide a path toward a future in which AI is deployed responsibly to improve governmental performance.

* * *