

# DAY ONE PROJECT

Creating Transparency and Fairness  
in Automated Decision Systems for  
Administrative Agencies

Kevin Wei

November 2020

The Day One Project offers a platform for ideas that represent a broad range of perspectives across S&T disciplines. The views and opinions expressed in this proposal are those of the author(s) and do not reflect the views and opinions of the Day One Project or its S&T Leadership Council.

## Summary

Artificial intelligence is increasingly being used to make decisions about human welfare. Automated decision systems (ADS) administer U.S. social benefits programs—such as unemployment and disability benefits—across local, state, and Federal governments. While ADS have the potential to enable large gains in efficiency, they also run a high risk of reinforcing the class- and race-based inequities of the status quo.<sup>1</sup> Additionally, the use of these systems is not transparent, often leaving individuals with no meaningful recourse after a decision has been made. Individuals may not even know that ADS played a role in the decision-making process.

The Federal Government should take immediate action to promote the transparency and accountability of automated decision systems. Agencies must build internal technical capacity as well as data cultures centered around transparency, accountability, and fairness. The White House should require that agencies using ADS undertake a notice-and-comment process to disclose information about these systems to the public. Finally, in the long-term, Congress must pass comprehensive legislation to implement a single, national standard regulating the use of ADS across sectors and use cases.

## Challenge and Opportunity

From housing and food assistance to unemployment and healthcare benefits, governmental social programs in the U.S. are increasingly being administered by technology. These “automated decision systems” (ADS) are assisting or replacing human-operated tasks, making or helping make choices along the administrative pipeline in municipal, state, and federal agencies.

ADS are generally comprised of advanced computational methods such as machine learning, artificial intelligence, or data analytics. They often produce quantitative metrics about individuals or families. For instance, ADS have been used in Pennsylvania to detect child abuse and neglect by scoring families for risk of abuse or neglect.<sup>2</sup>

There are two major challenges posed by the deployment of these systems: ADS can often be biased or discriminatory, and the governmental use of ADS is generally not transparent to the public. The possibility of bias is particularly harmful because ADS have the potential to reinforce cycles of racial, gender, and economic inequality that have been coded into American life.

---

<sup>1</sup> Smith, Megan, DJ Patil, and Cecilia Muñoz. 2016. Rep. Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights. The White House: Executive Office of the President, May 4, 2016. <https://obamawhitehouse.archives.gov/blog/2016/05/04/big-risks-big-opportunities-intersection-big-data-and-civil-rights>.

<sup>2</sup> Eubanks, Virginia. 2020. “A Child Abuse Prediction Model Fails Poor Families.” *Wired*, January 15, 2020. <https://www.wired.com/story/excerpt-from-automating-inequality/>.

Many ADS are unfairly biased against low-income families and historically marginalized communities. An ADS used in Michigan to discover unemployment fraud, for instance, had a 70% error rate between 2013 and 2015. In other words, a majority of the unemployment benefits cases classified by Michigan's system as fraudulent were not in fact fraudulent; these errors caused 44,000 benefits recipients to be mistakenly fined for fraud.<sup>3</sup> The aforementioned system used to detect child abuse and neglect in Pennsylvania also disproportionately identified poor families as "high risk," leading to systematic profiling of low-income families.<sup>4</sup> Many, if not most, ADS face similar challenges.

Some states and local jurisdictions are beginning to recognize the harms created by ADS and have introduced legislation to regulate these systems.<sup>5</sup> Instead of allowing a patchwork of regulation to emerge, however, the federal government should take leadership to set the standard for ADS regulation nationally in order to prevent bias and to ensure transparency in and accountability of these systems.

## Plan of Action

**Action 1: Agencies that administer benefits programs should develop robust in-house technical capacity, as well as strong data cultures centering data transparency and accountability.** In order to assess and develop algorithms that minimize bias, agencies that administer benefits programs must have the technical talent to understand the algorithms. Rather than outsourcing to external vendors or using off-the-shelf solutions, agencies would ideally build their own systems, because in-house systems are more customized, allow for more internal testing, and avoid the trade secrets issues of commercial products.<sup>8</sup>

Furthermore, agencies must develop data cultures that center transparency, accountability, and fairness. Some actions that can be taken at the agency level include:

- Hiring data scientists or machine learning engineers for development, testing, and implementation of ADS within the agency.

---

<sup>3</sup> Egan, Paul. 2017. "Michigan Agency Review Finds 70% Error Rate in Jobless Fraud Findings." Detroit Free Press, August 12, 2017. <https://www.freep.com/story/news/local/michigan/2017/08/11/michigan-agency-review-finds-70-error-rate-fraud-findings/559880001/>.

<sup>4</sup> Eubanks 2020.

<sup>5</sup> Chae, Yoon. 2020. "S. AI Regulation Guide: Legislative Overview and Practical Considerations." The Journal of Robotics, Artificial Intelligence & Law 3, no. 1 (2020): 17–40. <https://www.bakermckenzie.com/-/media/files/people/chae-yoon/rail-us-ai-regulation-guide.pdf>. See also: Pierce, Jadzia, and Frank Broomell. 2020. "California Introduces Bill to Regulate Automated Decision Systems." Inside Privacy, February 19, 2020. <https://www.insideprivacy.com/artificial-intelligence/california-introduces-bill-to-regulate-automated-decision-systems/>. Tiedrich, Lee J., Nooree Lee, and T'Shae Sherman. 2020. "AI Update: U.S. State and Local Government Task Forces Continue to Examine Artificial Intelligence Trustworthiness." Lexology, May 15, 2020. <https://www.lexology.com/library/detail.aspx?g=7d7195de-02b7-4b62-b040-9905dd7a0ca3>.

<sup>8</sup> For a discussion of the costs and benefits of building in-house solutions vs. buying commercial solutions, see: Engstrom, David Freeman, Daniel E. Ho, Catherine M. Sharkey, and Mariano-Florentino Cuéllar. "Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies." Report Submitted to the Administrative Conference of the United States, 2020, 88-90. <https://doi.org/10.2139/ssrn.3551505>.

- Naming a dedicated Chief Data Officer with responsibility for overseeing automated decision systems within the agency.<sup>9</sup>
- Embedding non-technical subject-matter experts to participate in the ADS development, testing, and implementation processes. Data scientists and machine learning engineers are usually technicians and *methodological* experts who may know little or nothing about agency processes and programs. Treating non-technical subject-matter experts as equal participants throughout this process is critical not just to prevent bias in ADS or to define technical variables, but also to ensure that the ADS meets all agency requirements upon completion.
- Investing in data literacy training and education for agency employees.<sup>10</sup>
- Conducting regular, formal benchmarking on ADS and privacy / bias audits on ADS before and after deployment.

**Action 2: The White House should issue a Presidential Memorandum requiring that agencies submit proposed ADS to a notice-and-comment process before adoption.** The development and deployment of ADS by government agencies is currently not subject to notice-and-comment rulemaking procedures.<sup>11</sup> Rather than unilaterally implementing ADS systems, agencies should be required to undergo a notice-and-comment process for ADS that play substantial roles in the administration of public benefits programs. Disclosing information about these systems would build transparency, engender public trust in ADS, and give affected communities an opportunity to evaluate and reflect on systems that, if adopted, would impact their lives. Furthermore, this process provides an ex-ante check on proposed systems and could thereby prevent harms generated by ADS before they occur. A form of notice-and-comment has also been prioritized for ADS regulation by legal experts.<sup>12</sup>

Specifically, ADS should undergo this notice-and-comment process:

- If developed in-house by an agency: after development and before implementation.
- If procured from an external vendor: after a vendor is selected, but before the contract is signed and the ADS implemented. As part of the procurement process, vendors should be required to disclose information about their systems and development processes to the procuring agency as well as to the public.<sup>13</sup>

---

<sup>9</sup> For background information on Chief Data Officers in federal agencies, see: Chappellet-Lanier, Tajha. 2019. "Agencies Are Now Required to Have a Chief Data Officer. Do They?" FedScoop, August 19, 2019. <https://www.fedscoop.com/federal-chief-data-officer-evidence-based-policy-making-deadline/>. See also: Gittings, Heather, Nick Hart, Tracy Jones, Jeff Lawton, and Joe Willey. 2020. "Effective Data Governance: A Survey of Federal Chief Data Officers." Data Foundation, August 2020. <https://www.datafoundation.org/effective-data-governance-a-survey-of-federal-chief-data-officers-2020>.

<sup>10</sup> On the importance of data literacy within public agencies, see: Gittings, Heather, Hart, Jones, Lawton, & Willey 2020.

<sup>11</sup> Engstrom, David Freeman, and Daniel E Ho. 2020. "Algorithmic Accountability in the Administrative State." *Yale Journal on Regulation* 37, no. 800 (2020): 800–854. <https://digitalcommons.law.yale.edu/cgi/viewcontent.cgi?article=1563&context=yjreg>.

<sup>12</sup> Strandburg, Katherine J. 2019. "Rulemaking and Inscrutable Automated Decision Tools." *Columbia Law Review* 119, no. 7 (November 2019): 1851–86. <https://www.jstor.org/stable/26810852?seq=1>. See also: Engstrom and Ho 2020.

<sup>13</sup> Requirements for vendor disclosure can be achieved through contractual language to that effect. See: Brauneis, Robert, and Ellen P. Goodman. 2018. "Algorithmic Transparency for the Smart City." *Yale Journal of Law & Technology* 20, no. 103 (2018). <https://yjolt.org/algorithmic-transparency-smart-city>. 164-166.

The following details should be released as part of the notice-and-comment process for each system<sup>14</sup>:

- The intended purposes and goals of the ADS.
- A cost-benefit analysis of implementing the ADS.
- Descriptions of the processes used to build (“train”) and test (“validate”) the ADS.
- Details about how the ADS will be used, including:
  - Internal use policies outlining how agency officials may consider ADS classifications, scores, or results in decision-making.
  - Which portions of the decision-making process would or would not be automated after the introduction of the ADS.
- Technical definitions of the input and output variables of the ADS.
- A full description of the datasets used to train and test the ADS, including a list of variables in each dataset, as well as summary statistics of each variable.<sup>15</sup>
- Explanations of how the ADS arrived at its results, whenever possible.<sup>16</sup>
- The source(s) of the datasets used to train and test the ADS.
- Analyses of the racial, gender, and income-level composition of the datasets used to train and test the ADS.
- Performance metrics of the ADS, including metrics assessing accuracy (relative to the training and test sets) as well as bias along racial, gender, and class (e.g. error rates, differential fairness<sup>17</sup>, or inequality indices<sup>18</sup>).
- Which communities and individuals could be impacted by the adoption of the ADS, especially if any communities would be disproportionately affected.
- What methods of recourse are available to members of the public if they were to be adversely impacted by the ADS.
- Potential privacy or cybersecurity implications of the ADS deployment, as well as implications of the collection of individual-level data to be input into the system.

---

<sup>14</sup> For another description of how such a process could be implemented, see: Reisman, Dillon, Jason Schultz, Kate Crawford, and Meredith Whitaker. 2018. Rep. Algorithmic Impact Assessments: A Practical Framework for Public Agency Accountability. AI Now Institute, April 2018. <https://ainowinstitute.org/aiareport2018.pdf>. See also: Strandburg 2019.

<sup>15</sup> A variable is any field in the dataset, such as names, race, income, geographical location, etc. Summary statistics may include averages, ranges, and other information. While the datasets themselves should not be released, a list of variables or field names in each dataset should be publicly available.

<sup>16</sup> An explanation could be either “Human-interpretable information about the factors used in a decision and their relative weight...A list of the factors that went into a decision, ideally ordered by the significance to the output” or “A list of the factors that went into a decision, ideally ordered by the significance to the output” or “An answer to a counterfactual question...For example, we may want to know what effect a particular change to an input has on the output or, conversely, what change must be made to the input to change the output in a particular way.” Note that such an explanation or list of factors may not always be possible. See: Doshi-Velez, Finale, Mason Kortz, Ryan Budish, Christopher Bavitz, Samuel J. Gershman, David O'brien, Stuart Shieber, Jim Waldo, David Weinberger, and Alexandra Wood. 2017. Rep. Accountability of AI Under the Law: The Role of Explanation. Berkman Klein Center for Internet & Society, 2017. <https://doi.org/10.2139/ssrn.3064761>.

<sup>17</sup> For technical details, see: Foulds, James R., Rashidul Islam, Kamrun Naher Keya, and Shimei Pan. 2020. “An Intersectional Definition of Fairness.” 2020 IEEE 36th International Conference on Data Engineering (ICDE), 2020. <https://doi.org/10.1109/icde48307.2020.00203>.

<sup>18</sup> For technical details, see: Speicher, Till, Hoda Heidari, Nina Grgic-Hlaca, Krishna P. Gummadi, Adish Singla, Adrian Weller, and Muhammad Bilal Zafar. 2018. “A Unified Approach to Quantifying Algorithmic Unfairness.” Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2018. <https://doi.org/10.1145/3219819.3220046>.

# DAY ONE PROJECT

- The contract and confidentiality terms surrounding externally-acquired systems and datasets, if any.

As part of this comment process, select third-party organizations (academics, think tanks, or advocacy organizations) should be allowed to conduct independent assessments or benchmarks of such systems. Pre-existing ADS that are currently in use should also be disclosed through this process.

It is also important to note that many federally-funded benefits programs are administered by state and local governments. Whenever possible, federal agencies that provide funding to sub-federal governments for these programs should require that funding recipients undergo similar disclosure processes before implementing ADS. Federal agencies should work with sub-federal or state agencies (e.g. by providing technical assistance) that administer federal programs to implement these disclosure processes at the state level.

Finally, agencies should *not* disclose certain information that could harm individual privacy or compromise the integrity of the ADS if implemented. This information would generally include systems' source code or datasets used to train and test the ADS, especially if those datasets contain personally identifiable information.

**Action 3: Congress should pass comprehensive legislation to regulate the use of ADS across domains including immigration & law enforcement, benefits programs, and others.** In the long term, Congress must codify regulations into law that would prevent the unchecked use and development of ADS in both the public and private sectors. Any such legislation should mandate transparency around and disclosure of the use of ADS, require continuous testing for bias, create mechanisms for accountability, regulate federal procurement processes for ADS, and create mechanisms for obtaining recourse for the public. One example of legislation that accomplishes some of these goals is the Algorithmic Accountability Act of 2019 (H.R. 2331).<sup>19</sup>

---

<sup>19</sup> Algorithmic Accountability Act of 2019. 2019. Bill, Congress.gov § (2019). <https://www.congress.gov/bill/116th-congress/house-bill/2231/text>.

## Frequently Asked Questions

### What technologies make up automated decision systems (ADS)?

“Automated decision systems” is one of many terms used to describe a variety of machine learning, artificial intelligence, data mining, data analytics, and predictive analytics technologies. These technologies take large datasets, find complex patterns or correlations in the data, and make predictions from those patterns or correlations. Predictions can take the forms of numerical scores (e.g. a number on a scale from 1-10) or classifications (e.g. “blue” vs. “red” vs. “green”, or “true” vs. “false”). For example, one ADS used in Arkansas used information about patient symptoms and medical diagnoses to determine how many caretaker hours a Medicaid recipient would receive.<sup>26</sup>

While these technologies are generally not called automated decision systems when used in other domains, such as law enforcement, the underlying technology is usually similar.

### Why target ADS in economic benefits programs?

Unlike the use of ADS in national security, law enforcement, and other domains, the use of ADS in the administration of governmental social programs has gone relatively unexamined. However, the impact of ADS in administering social programs could be much more far-reaching than the same systems in other domains, simply due to the sheer scale of benefits programs. In 2015, for instance, 52.2 million Americans (21.3% of the population) received assistance from Medicaid, the Supplemental Nutrition Assistance Program, Temporary Assistance for Needy Families, Supplemental Security Income, or housing assistance.<sup>27</sup> As of 2019, over 10 million Americans received benefits from Social Security Disability Insurance.<sup>28</sup>

### Where does bias in ADS come from? What are some examples of bias in ADS?

Bias can come from any of a number of sources. Some ways in which bias can appear in ADS include:

- Biases present in data used to train the ADS results in the algorithm returning biased results (“garbage in, garbage out”). For instance, the data could contain disproportionately fewer members of a minority demographic group resulting from bad data collection and sampling practices, or from historical inequities resulting in little data on that group. One example: in the early 2010s, Amazon attempted to develop a

---

<sup>26</sup> Lecher, Colin. 2018. “What Happens When an Algorithm Cuts Your Health Care.” The Verge, March 21, 2018. <https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy>.

<sup>27</sup> US Census Bureau. 2015. “21.3% Of US Participates in Government Assistance Programs Each Month.” May 28, 2015. <https://www.census.gov/newsroom/press-releases/2015/cb15-97.html>.

<sup>28</sup> “Chart Book: Social Security Disability Insurance.” 2019. Center on Budget and Policy Priorities. September 6, 2019. <https://www.cbpp.org/research/social-security/chart-book-social-security-disability-insurance>.

# DAY ONE PROJECT

machine learning model that would rate job applicants on a scale of 1-5. They trained the algorithm on resumes submitted to Amazon over 10 years, but because the majority of those job applicants were men, the algorithm learned to discriminate against women applicants.<sup>29</sup>

- Biases arising from unpredictable correlations in the data used to train the ADS, causing the algorithm to generate false associations that do not hold in the real world. One simple example of this principle: while developing an algorithm to diagnose skin cancer, scientists discovered at one point that the algorithm learned to rate photos of skin lesions next to rulers as more likely to be cancerous than photos of skin lesions without rulers (because doctors are more likely to measure lesions that they deem risky).<sup>30</sup>
- Biased or unfair conceptual definitions, measurements, or evaluation metrics in the ADS. These biases are normally highly specific to the algorithm at hand. Here is one example: Alleghany County's child neglect and abuse detection algorithm included a biased definition of "child maltreatment." The algorithm defined maltreatment as the higher number of either a "community re-referral, when a call to the hotline about a child was initially screened out but [the agency] receives another call on the same child within two years" or "child placement, when a call to the hotline about a child is screened in and results in the child being placed in foster care within two years."<sup>31</sup> However, this definition resulted in a disproportionate number of black and biracial families being flagged by the system because "mandated reporters and other members of the community call child abuse and neglect hotlines about black and biracial families three and a half times more often as they call about white families."<sup>32</sup>

## Why use the notice-and-comment process?

Notice-and-comment provides two substantial benefits over other regulatory mechanisms. First, it prevents harm from occurring rather than seeking to repair harms that have already been committed. Given that some benefits programs may be life-or-death for recipients (e.g., healthcare programs), a regulatory scheme that prioritizes preventing harms would be most appropriate.

Secondly, the notice-and-comment process provides a pre-existing structure through which ADS can be disclosed to the public while simultaneously allowing for robust public input. Many of the decisions made by or assisted by ADS can be viewed as "rulemaking" in a sense, and legal

---

<sup>29</sup> Dastin, Jeffrey. 2018. "Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women." Reuters, October 10, 2018. <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.

<sup>30</sup> Patel, Neel V. 2017. "Why Doctors Aren't Afraid of Better, More Efficient AI Diagnosing Cancer." The Daily Beast, December 11, 2017. <https://www.thedailybeast.com/why-doctors-arent-afraid-of-better-more-efficient-ai-diagnosing-cancer>.

<sup>31</sup> Eubanks 2020.

<sup>32</sup> Ibid.



commentators have already begun analyzing the question of adapting the notice-and-comment process for automated decision systems.<sup>33</sup>

Notice-and-comment may also be particularly suited to analyzing the use of ADS in government benefits programs. Stanford professors David Engstrom and Daniel Ho write that “notice-and-comment is more appropriate when AI adoption involves considerable distributive consequences. For instance, when [the ADS] expedites benefits to a distinct demographic group, the decision presents larger policy questions best suited for notice and comment.”<sup>34</sup>

## What are some examples of how ADS are used in the context of benefits programs?

Automated decision systems have been deployed

- at the Social Security Administration to prioritize the review of administrative law cases concerning disability benefits.<sup>35</sup>
- by the U.S. Department of Agriculture in conjunction with state governments to detect fraud in SNAP.<sup>36</sup>
- in Illinois to detect benefits fraud.<sup>37</sup>
- in Arkansas to determine how much Medicaid assistance patients would receive.<sup>38</sup>
- in Los Angeles to rank the distribution of housing for the homeless.<sup>39</sup>
- in Pennsylvania to detect child abuse and neglect.<sup>40</sup>

## Which other domains are ADS used in?

Artificial intelligence systems are also used in national security, law enforcement, education (e.g. instructor evaluations), and criminal justice systems (e.g. pretrial detention risk scores).<sup>41</sup> While these systems are not generally called “automated decision systems” in other domains, the underlying technology is much the same.

---

<sup>33</sup> Strandburg 2019.

<sup>34</sup> Engstrom and Ho 2020.

<sup>35</sup> Engstrom, Ho, Sharkey, and Cuéllar 2020.

<sup>36</sup> Office, U.S. Government Accountability. 2018. “Supplemental Nutrition Assistance Program: Disseminating Information on Successful Use of Data Analytics Could Help States Manage Fraud Risks.” U.S. Government Accountability Office (U.S. GAO), October 2, 2018. <https://theintercept.com/2020/03/02/ice-algorithm-bias-detention-aclu-lawsuit/>.

<sup>37</sup> Gilman, Michele. 2020. “AI Algorithms Intended to Root out Welfare Fraud Often End up Punishing the Poor Instead.” Chicago Reporter, February 21, 2020. <https://www.chicagoreporter.com/ai-algorithms-intended-to-root-out-welfare-fraud-often-end-up-punishing-the-poor-instead/>.

<sup>38</sup> Lecher 2018.

<sup>39</sup> Eubanks, Virginia. 2018. “High-Tech Homelessness.” American Scientist, 2018. <https://www.americanscientist.org/article/high-tech-homelessness>.

<sup>40</sup> Eubanks 2020.

<sup>41</sup> For a discussion on algorithmic decision making in the criminal justice system, see: Larson, Jeff, and Julia Angwin. 2016. “Machine Bias.” ProPublica, May 23, 2016. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>. See also: Metz, Cade, and Adam Satariano. 2020. “An Algorithm That Grants Freedom, or Takes It Away.” The New York Times, February 6, 2020. <https://www.nytimes.com/2020/02/06/technology/predictive-algorithms-crime.html>.

One example of such a system in law enforcement is an algorithm used by Immigrations and Customs Enforcement (ICE) to recommend detention or release for arrested immigrants.<sup>42</sup> Facial recognition is another prominent example of these systems in use.

ADS are also used in many commercial applications, e.g. in consumer reporting, hiring, tenant screening, and more.<sup>43</sup>

## **What are some examples of how agencies have successfully managed to build technical capacity?**

Most federal agencies are still working to build capacity and data cultures. However, one agency that may serve as a partial role model is the Securities and Exchange Commission (SEC), which has developed a number of algorithms to track financial and securities fraud. While the system at the SEC is as yet unperfected, laudable steps taken by the agency include<sup>44</sup>:

- Testing algorithms on historical data to evaluate performance (“back-testing”).
- Assigning technical staff to work with subject-matter experts in the agency to evaluate historical algorithm performance.<sup>45</sup>
- The collection of highly granular securities trading activity data that could be used to train future ADS.
- The rapid expansion of the agency’s information technology workforce; from 2017-2019, the SEC’s IT headcount grew by 15%, and its information security headcount grew by 75%.<sup>46</sup>

## **What are some examples of ADS transparency or notice-and-comment procedures that have been implemented successfully?**

One instance of a process for ADS regulation similar to notice-and-comment was implemented in Canada in 2019 as the “Directive on Automated Decision-Making.” Among other standards, the Canadian federal government now requires<sup>47</sup>:

---

<sup>42</sup> Biddle, Sam. 2020. “ACLU Sues to End ICE’s Rigged Algorithm for Detaining Immigrants.” *The Intercept*, March 2, 2020. <https://theintercept.com/2020/03/02/ice-algorithm-bias-detention-aclu-lawsuit/>.

<sup>43</sup> For a list of automated decision systems in these contexts, see: Gilman, Michele. 2020. “Poverty Lawgorithms: A Poverty Lawyer’s Guide to Fighting Automated Decision-Making Harms on Low-Income Communities.” *Data & Society*, September 15, 2020. <https://datasociety.net/library/poverty-lawgorithms/>. For a specific discussion of ADS in credit reporting, see: Citron, Daniele Keats, and Frank A. Pasquale. 2014. “The Scored Society: Due Process for Automated Decisions.” *Washington Law Review* 89 (January 7, 2014). <https://cyberlaw.stanford.edu/publications/scored-society-due-process-automated-predictions>.

<sup>44</sup> Engstrom, Ho, Sharkey, & Cuéllar 2020. 22-29.

<sup>45</sup> *Ibid* 26.

<sup>46</sup> 2019. “Oversight of the Securities and Exchange Commission: Wall Street’s Cop on the Beat. Testimony before the U.S. House of Representatives Committee on Financial Services of the 116th Congress, 2019. <https://www.sec.gov/news/testimony/testimony09-24-2019>.

<sup>47</sup> Government of Canada. 2019. “Directive on Automated Decision-Making,” February 5, 2019. <https://www.tbs-sct.gc.ca/pol/doc-eng.aspx?id=32592>.

- Public notice to be provided for “moderate-impact” ADS and above<sup>48</sup>:
  - For “moderate-impact” systems, there must be a “Plain language notice posted on the program or service website.”
  - For “high-impact systems” and “very high-impact systems,” documentation must be provided concerning ADS functionality, its usage in decision making, audit results, and descriptions of datasets used to train the ADS.<sup>49</sup>
- If the ADS uses proprietary software components, “the Government of Canada retains the right to access and test the Automated Decision System, including all released versions of proprietary software components” in certain instances.<sup>50</sup>
- Completion of “an Algorithmic Impact Assessment prior to the production of any Automated Decision System”<sup>51</sup>:
  - If the ADS “functionality or the scope” is changed, the assessment must be updated.<sup>52</sup>
  - The final assessment must also be released to the public.<sup>53</sup>
- All ADS must be tested for bias prior to implementation<sup>54</sup>:
  - For “moderate-impact” ADS and above, the system must also undergo peer review by qualified experts.<sup>55</sup>

## What costs and challenges might be associated with a notice-and-comment procedure for ADS?

A notice-and-comment procedure for ADS would introduce additional administrative burdens on agencies. Specifically, agencies would be required to read and respond to every publicly submitted comment, which may require significant resources depending on how many comments are submitted.

Moreover, notice-and-comment would also delay the implementation of any ADS. Agencies’ comment periods generally last at least 30 days<sup>56</sup>, and additional time would be required to review and finalize the ADS after the close of the comment period. Thus, an ADS might not be implemented for months pending the conclusion of the notice-and-comment process.

---

<sup>48</sup> Ibid § 6.2.1. See Appendix B for definitions.

<sup>49</sup> Ibid Appendix C.

<sup>50</sup> Ibid § 6.2.5.2.

<sup>51</sup> Ibid § 6.1.1.

<sup>52</sup> Ibid § 6.1.3.

<sup>53</sup> Ibid § 6.1.4.

<sup>54</sup> Ibid § 6.3.1.

<sup>55</sup> Ibid § 6.3.4. See also Appendix C.

<sup>56</sup> Carey, Maeve P. 2013. “The Federal Rulemaking Process: An Overview.” Congressional Research Service, June 17, 2013. <https://www.everycrsreport.com/reports/RL32240.html>

## What steps have states and municipalities taken to regulate ADS?

Some states and municipalities have proposed regulatory standards for ADS:

- In California, AB-2269 would require that businesses using ADS continually test those systems for bias.<sup>57</sup>
- In Washington state, SB-5527 / HB-1655 would regulate the use of ADS by governmental agencies in a wide range of use cases, as well as empower individuals harmed by ADS to take legal action.<sup>58</sup>
- Task forces have been established in New York City and Vermont to assess the impact of ADS and artificial intelligence.<sup>59</sup>

For a full and updated list of artificial intelligence-related legislation introduced in Congress and in the California state legislatures, see the database maintained by the CITRIS Policy Lab at the University of California.<sup>60</sup>

---

<sup>57</sup> AB-2269 (2019-2020). [https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill\\_id=201920200AB2269](https://leginfo.legislature.ca.gov/faces/billTextClient.xhtml?bill_id=201920200AB2269)

<sup>58</sup> SB-5527 (2019). <http://lawfilesexternal.leg.wa.gov/biennium/2019-20/Pdf/Bills/Senate%20Bills/5527.pdf>

<sup>59</sup> While the New York City Automated Decision Systems Task Force failed to issue a final report, a shadow report was released by a coalition of NYC-based NGOs. See: Rashida, Richardson, ed. 2019. "Confronting Black Boxes: A Shadow Report of the New York City Automated Decision System Task Force." AI Now Institute, December 4, 2019. <https://ainowinstitute.org/ads-shadowreport-2019.html>. For the report of the Vermont Artificial Intelligence Task Force, see: Rep. Artificial Intelligence Task Force Final Report, January 15, 2020. <https://accd.vermont.gov/economic-development/artificial-intelligence-task-force>.

<sup>60</sup> "Federal & California AI Legislation." 2020. CITRIS Policy Lab. October 13, 2020. <https://citrispolicylab.org/ailegislation>.

# DAY ONE PROJECT

## About the Author



Kevin Wei is a graduate student in computer science at the Georgia Institute of Technology. His research interests include bias and fairness in machine learning and artificial intelligence, content regulation on the Internet, and the social impacts of technology. In addition to his academic work, he is an activist based in New York City.



## About the Day One Project

The Day One Project is dedicated to democratizing the policymaking process by working with new and expert voices across the science and technology community, helping to develop actionable policies that can improve the lives of all Americans, and readying them for Day One of a future presidential term. For more about the Day One Project, visit [dayoneproject.org](http://dayoneproject.org).