137 Questions:
Criminal Risk Assessment Algorithms as a Case Study for Data Ethics

Xinlan Emily Hu

4 June 2019

Suppose that I ask you 137 questions.

How many times have you moved?

Did you ever skip classes?

Do you ever become bored?

Then, having fed your answers through a sophisticated algorithm, I assign you a score from 1 to 10—and if the result is too high, a judge will make you ineligible for release on bail or recommend a harsher criminal sentence.

This is the reality of the Correctional Offender Management Profiling for Alternative Sanctions (COMPAS), a common American criminal risk assessment tool[1]. COMPAS's 137 questions include details about the individual's personal and criminal history—including the seemingly mundane questions above[2]. Its use in criminal sentencing has recently sparked controversy; in 2016, the civic organization ProPublica published a report that criticized COMPAS for bias against black defendants. Northpointe, the algorithm's producer, rebutted the report with its own analysis of fairness[3]. At the heart of the COMPAS controversy is an ethical question surrounding the use of information technology in determining the course of lives.

My paper centers the COMPAS case within current debates in the field of data ethics. Broadly speaking, data ethics deals with moral questions related to using data—its generation, curation, and dissemination, as well as its role in algorithms and artificial intelligence[4]. COMPAS is a particularly salient case study in this field because it sits at the crossroads of industry (for it is a proprietary piece of software), government (where it is purchased and used), and civil society (which is concerned about its legitimacy).

Unlike some scholars and advocacy groups in the field, I argue that there is a place for algorithmic tools in the criminal justice system. However, their use should be redefined under a human rights-based framework. The law loses transparency and accountability otherwise, becoming subject to the arbitrary design choices of Northpointe and other profit-seeking corporations. I make three recommendations for governments in future implementations of risk-assessment tools. Drawing from an international body of literature, I argue that these tools

---

[1] Huq, Aziz Z. "Racial Equity in Algorithmic Criminal Justice." (*Duke Law Journal* 68, no. 6 (March 2019): 1043-134. Accessed May 27, 2019. https://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=3972&context=dlj), 1047.

[2] Angwin, Julia. "Sample-COMPAS-Risk-Assessment-COMPAS-"CORE"." DocumentCloud. Accessed June 01, 2019. https://www.documentcloud.org/documents/2702103-Sample-Risk-Assessment-COMPAS-CORE.html.

[3] Corbett-Davies, Sam, Emma Pierson, Avi Feller, and Sharad Goel. "A Computer Program Used for Bail and Sentencing Decisions Was Labeled Biased against Blacks. It's Actually Not That Clear." The Washington Post. October 17, 2016. Accessed June 01, 2019. https://www.washingtonpost.com/news/monkey-cage/wp/2016/10/17/can-an-algorithm-be-racist-our-analysis-is-more-cautious-than-propublicas/?utm_term=.43d0246f0685.

[4] Floridi, Luciano, and Mariarosaria Taddeo. "What Is Data Ethics?" (*Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*374, no. 2083 (2016): 20160360. Accessed October 17, 2018. doi:10.1098/rsta.2016.0360), 1.

should emphasize context rather than a raw score, that they should be decoupled from corporations in order to increase transparency, and that they should be paired with initiatives that generate oversight and raise awareness about bias.

The central tension in data ethics is one between *opportunity* and *opportunism*: exabytes of collected information can better the lives of citizens in unprecedented ways, but also leave room for breaches of ethics and privacy at an enormous scale. In other words, the line between usefulness and invasion is incredibly fine. Consider the case of personal financial and residential data, which play a large role in streamlining government services. Massive collections of such data enable the identification of otherwise indiscernible patterns: the American Internal Revenue Service uses data analytics to determine tax evasion and fraud; cities use big data to prevent unnecessary residential vacancies and create smarter housing[5].

Simultaneously, the same data produces personal risk, particularly to marginalized individuals and minorities who may be identified with the data. In a world where websites can guess a user's identity with just three basic data points—date of birth, ZIP code, and gender[6]—data re-identification is a legitimate problem. When more pieces of information exist about an individual, the *k-anonymity* (number of individuals who share the same trait) of the data decreases[7]. This risk is further compounded when large data sets are merged, linked, or re-used, making it even more likely that information about the same individual could be connected[8]. Once this occurs, the data may be used for the purpose of "group discrimination (e.g., ageism, ethnicism, sexism)" and even "group-targeted forms of violence[9]."

Crucially, then, data ethics is charged with the task of balancing the potential for social innovation with its risks. The world must grapple with a future where once well-established institutions will experience a radical transformation. In the United States, for instance, the incarceration rate is the highest in the world, and well over twice that of Israel—the distant second place. Against a backdrop of overcrowded prisons and rising incarceration rates, criminal justice reform turned digital: to the criminal risk assessment algorithm[10]. Those optimistic about tools like COMPAS hope that, just as big data had boosted the efficiency of the IRS and helped diagnose climate change problems in the UK[11], so it could rescue a crippled U.S. criminal justice system.

The resulting 137-question tool, intended to assess the rate of reoffending for a given defendant, soon became mired in controversy. Although race is not among the factors assessed in the COMPAS survey, the tool has suffered criticism for assigning black defendants higher risk

---

[5] Kim, Gang-Hoon, Silvana Trimi, and Ji-Hyong Chung. "Big-data Applications in the Government Sector." (*Communications of the ACM* 57, no. 3 (2014): 78-85. Accessed June 1, 2019. doi:10.1145/2500873), 83

[6] Dale, Brady, and Brady Dale. "Websites Can Probably Guess Your Identity With Three Basic Data Points." Observer. June 16, 2016. Accessed June 01, 2019. https://observer.com/2015/09/http-injection-browser-fingerprinting-w3c-access-mark-nottingham/.

[7] "Taking Charge of Your Data: Understanding Re-identification Risk and Quasi-identifiers with Cloud DLP | Google Cloud Blog." Google. Accessed June 01, 2019. https://cloud.google.com/blog/products/identity-security/taking-charge-of-your-data-understanding-re-identification-risk-and-quasi-identifiers-with-cloud-dlp.

[8] Floridi and Taddeo, What Is Data Ethics?", 3.

[9] Ibid.

[10] The Partnership on AI. *Report on Algorithmic Risk Assessment Tools in the U.S. Criminal Justice System*. (Report. San Francisco, California: Partnership on AI (PAI), 2019), 7-8.

[11] Kim, et. al., "Big-data Applications in the Government Sector," 83.

scores. As ProPublica points out, among those who ultimately did not reoffend, blacks had a 42% chance of being assigned a high risk score—whereas whites were incorrectly labeled high-risk only 22% of the time[12]. Meanwhile, Northpointe rebuts that its tool has fair accuracy rates: among those who were assigned a 7 out of 10, 60% of white defendants reoffended, as did 61% of black defendants[13].

This case is representative of the data ethics debate at its core: is the potentially transformative benefit of an algorithmic tool (in this case, streamlining an overcrowded system and adding a sense of impartiality) worth the cost of revealing and making skewed judgments on the basis of race? And what, if anything, is the way forward for governments, industry, and civil society stakeholders?

Presently, civil society groups such as the Partnership on AI are calling for "jurisdictions to cease using the tools in decisions" altogether, arguing that the technology is simply not ready in its current form[14]. However, this conclusion is too naive. It does not at all grapple with the fact the government remains badly in need of reform, and ignores untapped potential in algorithmic tools. Discrediting big data in criminal justice altogether throws the baby out with the bathwater.

I instead propose a reframing of the COMPAS debate. The question is not so much a dichotomy—use COMPAS or entirely discard it—as it is one of how governments can appropriately and equitably integrate algorithmic tools into criminal justice. The COMPAS debate, as it currently stands, is irresolvable because it relies on two mathematically contradictory definitions of fairness. It is impossible to both be "fair" in Northpointe's sense (to have the score of 7 correspond to 60% regardless of race) and to be "fair" in ProPublica's sense (to have the same rate of false positives across races). Due to existing biases in the criminal justice system, the recidivism rate of black defendants is higher than that of whites (52% versus 39%). This means that "a greater share of black defendants will be classified as high risk. And if a greater share of black defendants is classified as high risk, then…a greater share of black defendants who do not reoffend will also be classified as high risk[15]"—causing the systematic racial bias that ProPublica points out.

In other words, it is not possible to design an algorithm that satisfies all definitions of fairness. The United States must therefore reconsider its definition of fair, ethical policy positions in the context of algorithmic tools.

Here, I draw from two recently proposed frameworks for the ethical use of artificial intelligence. Vidushi Marda describes a three-stage framework for AI— Data, Model, and Application[16]—through a case study of India. Marda's argument is applicable to AI in general, calling for an evaluation of factors such as training, data parity, and security within each stage. More recently, ARTICLE19 proposed the notion of a "human rights-based approach to AI,"

---

[12] Corbett-Davies, et. al., "A Computer Program Used for Bail and Sentencing Decisions Was Labeled Biased against Blacks."

[13] Ibid.

[14] The Partnership on AI, *Report on Algorithmic Risk Assessment Tools,* 33.

[15] Ibid.

[16] Marda, Vidushi. "Artificial Intelligence Policy in India: A Framework for Engaging the Limits of Data-driven Decision-making." (*Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*376, no. 2133 (2018). Accessed June 3, 2019. doi:10.1098/rsta.2018.0087), 2.

drawing from an international ethical and legal standard[17]. Both of these frameworks ground a better approach to AI in the criminal justice system.

As the direct implementers of algorithmic tools, it is incumbent upon states to make the ethical use of artificial intelligence a priority. ARTICLE19's recommendation calls upon states to "[h]old AI systems to accountability, responsibility, and constitutional standards without dilution or exception"—a duty that is as important as it is incredibly expansive. Three specific standards from Marda's framework help to narrow the concerns relevant to criminal justice reform: System and Historical Bias, Feature Selection, and Fairness. (Of course, there is also plenty to be said of other issues, but I have chosen to focus on these three particularly central standards.)

The first standard from Marda's framework asks governments to consider whether data can "cement, formalize and imbibe" biases in the algorithm's greater environment[18]. The second standard mandates oversight for data reidentification. Careless algorithmic design can inadvertently reveal protected attributes (such as race) through the feature selection process[19]. Many of the features used are often a proxy for race, even if race is never explicitly mentioned. Finally, governments should pay close attention to tradeoffs inherent in attempting to achieve a fair algorithm. The third standard makes clear that fairness often erodes competing values. For instance, "[r]emoval of discrimination has been shown to reduce overall accuracy in a model[20]."

Held up against these standards, COMPAS fails in a few notable ways. Its one-size-fits-all numerical score often erases nuances behind the data. The collated risk score conflates correlated factors that in fact have very different causes. "The reasons for someone not appearing in court, getting re-arrested, and/or getting convicted of a future crime are all very distinct, so a high score…would group together people who are likely to have a less dangerous outcome…with [those likely to have] more dangerous outcomes[21]." COMPAS thus produces the very risk that Marda warns of in the second standard. Even though race is not explicitly used as a feature, many of the features disproportionately relate to race—for instance, black defendants are likely to have more prior arrests, causing the algorithm to implicitly create a racial bias even when race is removed from the training data[22].

These considerations point to the fact that, while Northpointe's definition of fairness (that 60% of white defendants and 61% of black defendants reoffend) is legitimate on face, it is not worth the obvious tradeoffs to having fair outcomes. Evaluating COMPAS under a human rights-centric framework should therefore lead us to prioritize the quality of treatment of minority groups over the sanctity of a numerical fairness standard.

Note here that, unlike the Partnership on AI, I am not concluding that algorithmic tools like COMPAS should be entirely abandoned. I advocate instead for a paradigm shift on the part of governments as they implement algorithmic decision-making tools in criminal justice.

---

[17] ARTICLE 19. *Governance with Teeth: How Human Rights Can Strengthen FAT and Ethics Initiatives on Artificial Intelligence*. (PDF. London, April 2019), 17.

[18] Marda, Vidushi. "Artificial Intelligence Policy in India," 7.

[19] Ibid., 8.

[20] Ibid., 9.

[21] The Partnership on AI, *Report on Algorithmic Risk Assessment Tools,* 22.

[22] Corbett-Davies, et. al., A Computer Program Used for Bail and Sentencing Decisions Was Labeled Biased against Blacks."

Algorithms are "more than simple mathematical problems: they are socio-technical systems that depend on the contextual setting in which they function[23]." As noted in the introduction, I make three broad recommendations to this end: first, I recommend that criminal risk assessment tools prioritize giving statistical insights over delivering a consolidated risk score. Second, I recommend decoupling these tools from corporations, as copyrighted tools are difficult to scrutinize. And finally, I recommend that the corrections departments introducing these tools undertake initiatives to mitigate the risk of bias—for instance, by establishing an ethical advisory board, conducting regular audits, and providing appropriate bias training for police officers.

My first recommendation stems from a need to shift away from a single numerical score. In addition to the mathematical impossibility of creating an unbiased score, providing police officers a number with no context causes *automation bias*—in which "information presented by a machine is viewed as inherently trustworthy and above skepticism[24]." COMPAS is, in fact, an example of this effect. Its current use suggests a tendency to trust the algorithm even in cases where it is contextually inappropriate. For example, COMPAS is typically used on a local or regional basis, yet it was trained on a nationwide sample (which may not be representative of local trends). Moreover, although COMPAS was not originally developed for use during sentencing—and was meant to merely assess recidivism risk—judges frequently do incorporate COMPAS scores in sentencing decisions[25].

In the future, numerical scores should not be presented in isolation. They should, at minimum, include details about the source of data, the potential bias, and the original intended use of the tool. Officers who use the tools should be empowered to make decisions given a full understanding, rather than be blindly led to trust the computer. Individuals whose lives are impacted by an algorithm's judgement also have a right to understand how and why the decision was found. Even though current technology is far from a truly "explainable AI," the algorithmic system should disclose the specific measures that were used in the decision-making process. These explanation facilities would enable "individuals and users assess [to] whether a given output is justified, and whether they should seek a remedy through the courts[26]."

Furthermore, tools should prioritize delivering information about trends at the systemic level rather than attempting to over-generalize group data for individual defendants. One yet-unexplored use of big data is to question longstanding assumptions in criminal justice: whether mental health leads to violence; whether incarceration effectively decreases recidivism;

---

[23] Marda, Vidushi. "Artificial Intelligence Policy in India," 2.

[24]The Partnership on AI, *Report on Algorithmic Risk Assessment Tools,* 23.

[25] Kehl, Danielle, Priscilla Guo, and Samuel Kessler. "Algorithms in the Criminal Justice System: Assessing the Use of Risk Assessments in Sentencing." (*Responsive Communities Initiative, Berkman Klein Center for Internet & Society*, July 2017. Accessed June 3, 2019. http://nrs.harvard.edu/urn-3:HUL.InstRepos:33746041), 20.

[26] The Law Society Commission on the Use of Algorithms in the Justice System, and The Law Society of England and Wales. "Algorithm Use in the Criminal Justice System Report." (The Law Society. Accessed June 04, 2019. https://www.lawsociety.org.uk/support-services/research-trends/algorithm-use-in-the-criminal-justice-system-report/), 66.

whether community policing reduces crime[27]. In this way, a new opportunity for dialogue opens.

My second recommendation emerges from comparing COMPAS to its UK counterpart, the Harm Assessment Risk Tool (HART). Notably, the UK has a far less commercial approach to developing such tools. Unlike COMPAS, which is owned by Northpointe and licensed to individual jurisdictions, HART was created in a collaborative effort between the Durham Constabulary (a UK police department) and the University of Cambridge[28]. This collaborative approach increases the transparency of the development process. In contrast, despite widespread criticism of bias, it is difficult to fully investigate COMPAS because it is a piece of proprietary software. As one group of computer science researchers explain, "Northpointe has refused to disclose the details of its proprietary algorithm. …That's understandable: Northpointe needs to protect its bottom line. But it raises questions about relying on for-profit companies to develop risk assessment tools[29]." I recommend that corrections departments avoid using algorithmic tools that, for copyright reasons, cannot be scrutinized. Governments should not outsource the execution of their laws to a corporation.

Finally, I recommend establishing institutional oversight for algorithmic tools. One generally successful example of such oversight is the West Midlands Ethics Committee in the United Kingdom. Following the introduction of algorithmic decision-making for criminal cases, the West Midlands Police and Crime Commissioner established an ethics board to oversee algorithms' use and evaluation. The board is an excellent example of engaging civil society stakeholders. It is comprised mostly of residents in the West Midlands community, who have a "diverse range of relevant skills and experiences" and an equal gender balance[30].

HART provides another example of oversight. Officers who use HART undergo substantial bias awareness training. The training materials make clear that HART is not meant to provide a comprehensive picture; "the custody officers… remain the decision makers and must ensure that the HART output is but one factor they consider alongside all of the many other factors they are statutorily obliged to consider[31]."

These examples constitute important steps in the right direction. Certainly the UK's approach is far from perfect: a June 2019 report by the Law Society of England and Wales found "a lack of explicit standards, best practice, and openness or transparency about the use of algorithmic systems in criminal justice across England and Wales[32]." Moreover, HART has raised its own set of controversies over bias and data privacy[33]. Though we should laud these

---

[27] Mears, Daniel P. "How Big Data Can save America's out of Control Criminal Justice Policies." LSE US Centre. January 15, 2018. Accessed June 01, 2019. https://blogs.lse.ac.uk/usappblog/2018/01/10/how-big-data-can-save-americas-out-of-control-criminal-justice-policies/.
[28] Oswald, Marion, Jamie Grace, Sheena Urwin, and Geoffrey C. Barnes. "Algorithmic Risk Assessment Policing Models: Lessons from the Durham HART Model and 'Experimental' Proportionality." (*Information & Communications Technology Law* 27, no. 2 (2018): 223-50. Accessed June 3, 2019. doi:10.1080/13600834.2018.1458455), 225.
[29] Qtd. in Kehl, et. al., "Algorithms in the Criminal Justice System", 33.
[30] "Ethics Committee." West Midlands Police and Crime Commissioner- Ethics Committee. Accessed June 03, 2019. https://www.westmidlands-pcc.gov.uk/transparency/ethics-committee.
[31] Oswald, Marion, et. al., "Algorithmic Risk Assessment Policing Models," 231.
[32] The Law Society Commission on the Use of Algorithms in the Justice System, "Algorithm Use," 4.
[33] Ibid., 46.

examples of involving civil society stakeholders and implementing anti-bias training, clearly much work remains.

In short, as governments decide how to incorporate big data into institutional reform, they should adopt a critical lens that acknowledges the risk that algorithms will link disparate features together, re-identifying sensitive demographic information and perpetuating oppression of minority groups. Especially because algorithms appear unbiased on the surface, "it will often not be clear to a human operator that an algorithmic criminal justice tool needs reconsideration[34]." States that hope to invest in AI must recognize that these tools can as easily solve their problems as they can re-intrench them. When the tools are implemented, it is governments' duty to enshrine the human rights-based approach to AI in policy. Per ARTICLE19's recommended paradigm, states should conduct "human rights impact assessments…[and] continuous auditing. These are not systems that can simply be rolled out. They should instead be tailored to the exact context and use for which they are intended[35]."

These ethical protections should also be understood in the context of equal protection and due process law within their respective jurisdictions. Although a full legal review of the scholarship on this matter is out of scope, it should be noted that many of the concerns raised in this paper also have a basis in legal literature. Differences in legal tradition also explain the distinct approaches to algorithmic decision-making seen in the U.S. and the UK. However, I have attempted in this paper to ground my recommendations in an international human rights paradigm, rather than appeal to the law of any particular jurisdiction.

Ultimately, racial equity in the criminal justice system should be motivated by a human-rights centered data ethics, in which issues such as reidentification risk and automation bias remain front and center. In a world where 137 questions can condemn a life, governments ought to orient themselves toward a more equitable future, rather than cling to a flawed past.

[34] Huq, Racial Equity in Algorithmic Criminal Justice," 1067.
[35] ARTICLE 19. *Governance with Teeth,* 21.

Bibliography

Angwin, Julia. "Sample-COMPAS-Risk-Assessment-COMPAS-"CORE"." DocumentCloud. Accessed June 01, 2019. https://www.documentcloud.org/documents/2702103-Sample-Risk-Assessment-COMPAS-CORE.html.

ARTICLE 19. *Governance with Teeth: How Human Rights Can Strengthen FAT and Ethics Initiatives on Artificial Intelligence*. PDF. London, April 2019.

Corbett-Davies, Sam, Emma Pierson, Avi Feller, and Sharad Goel. "A Computer Program Used for Bail and Sentencing Decisions Was Labeled Biased against Blacks. It's Actually Not That Clear." The Washington Post. October 17, 2016. Accessed June 01, 2019. https://www.washingtonpost.com/news/monkey-cage/wp/2016/10/17/can-an-algorithm-be-racist-our-analysis-is-more-cautious-than-propublicas/?utm_term=.43d0246f0685.

Dale, Brady, and Brady Dale. "Websites Can Probably Guess Your Identity With Three Basic Data Points." Observer. June 16, 2016. Accessed June 01, 2019. https://observer.com/2015/09/http-injection-browser-fingerprinting-w3c-access-mark-nottingham/.

"Ethics Committee." West Midlands Police and Crime Commissioner- Ethics Committee. Accessed June 03, 2019. https://www.westmidlands-pcc.gov.uk/transparency/ethics-committee.

Floridi, Luciano, and Mariarosaria Taddeo. "What Is Data Ethics?" *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*374, no. 2083 (2016): 20160360. Accessed October 17, 2018. doi:10.1098/rsta.2016.0360.

Floridi, Luciano. "Soft Ethics, the Governance of the Digital and the General Data Protection Regulation." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*376, no. 2133 (2018): 20180081. Accessed October 17, 2018. doi:10.1098/rsta.2018.0081.

Hassani, Hossein, Xu Huang, and Emmanuel Silva. "Big Data and Climate Change." *Big Data and Cognitive Computing*3, no. 1 (2019): 12. doi:10.3390/bdcc3010012.

Huq, Aziz Z. "Racial Equity in Algorithmic Criminal Justice." *Duke Law Journal*68, no. 6 (March 2019): 1043-134. Accessed May 27, 2019. https://scholarship.law.duke.edu/cgi/viewcontent.cgi?article=3972&context=dlj.

Kehl, Danielle, Priscilla Guo, and Samuel Kessler. "Algorithms in the Criminal Justice System: Assessing the Use of Risk Assessments in Sentencing." *Responsive Communities*

Initiative, Berkman Klein Center for Internet & Society, July 2017. Accessed June 3, 2019. http://nrs.harvard.edu/urn-3:HUL.InstRepos:33746041.

Kim, Gang-Hoon, Silvana Trimi, and Ji-Hyong Chung. "Big-data Applications in the Government Sector." *Communications of the ACM*57, no. 3 (2014): 78-85. Accessed June 1, 2019. doi:10.1145/2500873.

Marda, Vidushi. "Artificial Intelligence Policy in India: A Framework for Engaging the Limits of Data-driven Decision-making." *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*376, no. 2133 (2018). Accessed June 3, 2019. doi:10.1098/rsta.2018.0087.

Mears, Daniel P. "How Big Data Can save America's out of Control Criminal Justice Policies." LSE US Centre. January 15, 2018. Accessed June 01, 2019. https://blogs.lse.ac.uk/usappblog/2018/01/10/how-big-data-can-save-americas-out-of-control-criminal-justice-policies/.

Oswald, Marion, Jamie Grace, Sheena Urwin, and Geoffrey C. Barnes. "Algorithmic Risk Assessment Policing Models: Lessons from the Durham HART Model and 'Experimental' Proportionality." *Information & Communications Technology Law*27, no. 2 (2018): 223-50. Accessed June 3, 2019. doi:10.1080/13600834.2018.1458455.

"Taking Charge of Your Data: Understanding Re-identification Risk and Quasi-identifiers with Cloud DLP | Google Cloud Blog." Google. Accessed June 01, 2019. https://cloud.google.com/blog/products/identity-security/taking-charge-of-your-data-understanding-re-identification-risk-and-quasi-identifiers-with-cloud-dlp.

The Law Society Commission on the Use of Algorithms in the Justice System, and The Law Society of England and Wales. "Algorithm Use in the Criminal Justice System Report." The Law Society. Accessed June 04, 2019. https://www.lawsociety.org.uk/support-services/research-trends/algorithm-use-in-the-criminal-justice-system-report/.

The Partnership on AI. *Report on Algorithmic Risk Assessment Tools in the U.S. Criminal Justice System*. Report. San Francisco, California: Partnership on AI (PAI), 2019.