

THE BIASED ALGORITHM: EVIDENCE OF DISPARATE IMPACT ON HISPANICS

Melissa Hamilton*

INTRODUCTION

Automated risk assessment is all the rage in the criminal justice system. Proponents view risk assessment as an objective way to reduce mass incarceration without sacrificing public safety. Officials thus are becoming heavily invested in risk assessment tools—with their reliance upon big data and algorithmic processing—to inform decisions on managing offenders according to their risk profiles.

While the rise in algorithmic risk assessment tools has earned praise, a group of over 100 legal organizations, government watch groups, and minority rights associations (including the ACLU and NAACP) recently signed onto “A Shared Statement of Civil Rights Concerns” expressing unease with whether the algorithms are fair.¹ In 2016, the investigative journalist group ProPublica kickstarted a public debate on the topic when it proclaimed that a popular risk tool called COMPAS was biased against Blacks.² Prominent news sites highlighted ProPublica’s message that this proved yet again an area in which criminal justice consequences were racist.³ Yet the potential that risk algorithms are unfair to another minority group has received far less attention in the media or amongst risk assessment scholars and statisticians: Hispanics.⁴ The general disregard here exists despite the fact that Hispanics represent an important cultural group in the American population with recent estimates revealing that they are the largest minority with almost fifty-eight million members, and that number is rising quickly.⁵

* Senior Lecturer of Law & Criminal Justice, University of Surrey School of Law; J.D., The University of Texas at Austin School of Law; Ph.D, The University of Texas at Austin (criminology/criminal justice). © 2019, Melissa Hamilton.

1. Ted Gest, *Civil Rights Advocates Say Risk Assessment May “Worsen Racial Disparities” in Bail Decisions*, THE CRIME REPORT (July 31, 2018), <https://thecrimereport.org/2018/07/31/civil-rights-advocates-say-risk-assessment-may-worsen-racial-disparities/>.

2. See *infra* Section I.C.

3. See, e.g., Max Ehrenfreund, *The Machines that Could Rid Courtrooms of Racism*, THE WASHINGTON POST: WONKBLOG (Aug. 18, 2016), https://www.washingtonpost.com/news/wonk/wp/2016/08/18/why-a-computer-program-that-judges-rely-on-around-the-country-was-accused-of-racism/?noredirect=on&utm_term=.ce854f237cfe; NPR, *The Hidden Discrimination in Criminal Risk-Assessment Scores* (May 24, 2016, 4:32 PM), <https://www.npr.org/2016/05/24/479349654/the-hidden-discrimination-in-criminal-risk-assessment-scores>; Ed Yong, *A Popular Algorithm Is No Better at Predicting Crime than Random People*, THE ATLANTIC (Jan. 18, 2018), <https://www.theatlantic.com/technology/archive/2018/01/equivant-compas-algorithm/550646/>; Li Zhou, *Is Your Software Racist?*, POLITICO (Feb. 7, 2018, 5:56 AM), <https://www.politico.com/agenda/story/2018/02/07/algorithmic-bias-software-recommendations-000631>.

4. Stephane M. Shepherd & Roberto Lewis-Fernandez, *Forensic Risk Assessment and Cultural Diversity: Contemporary Challenges and Future Directions*, 22 PSYCHOL. PUB. POL’Y & L. 427, 428 (2016).

5. Antonio Flores, *How the U.S. Hispanic Population Is Changing*, PEW RESEARCH CENTER (Sept. 18, 2017), <http://www.pewresearch.org/fact-tank/2017/09/18/how-the-u-s-hispanic-population-is-changing/>.

This Article intends to partly remedy this gap in interest by reporting on an empirical study about risk assessment with Hispanics at the center. The study uses a large dataset of pretrial defendants who were scored on a widely-used algorithmic risk assessment tool soon after their arrests. The report proceeds as follows. Section II briefly reviews the rise in algorithmic risk assessment in criminal justice generally, and then in pretrial contexts more specifically. The discussion summarizes the ProPublica findings regarding the risk tool COMPAS after it analyzed COMPAS scores comparing Blacks and Whites.

Section III discusses further concerns that algorithmic-based risk tools may not be as transparent and neutral as many presume them to be. Insights from behavioral sciences literature suggest that risk tools may not necessarily incorporate factors that are universal or culturally-neutral. Hence, risk tools developed mainly on Whites may not perform as well on heterogeneous minority groups. As a result of these suspicions, experts are calling on third parties to independently audit the accuracy and fairness of risk algorithms. The study reported in Section IV responds to this invitation. Using the same dataset as ProPublica, we offer a range of statistical measures testing COMPAS' accuracy and comparing outcomes for Hispanics versus non-Hispanics. Such measures address questions about the tool's validity, predictive ability, and the potential for algorithmic unfairness and disparate impact upon Hispanics. Conclusions follow.

I. ALGORITHMIC RISK ASSESSMENT

Risk assessment in criminal justice is about predicting an individual's potential for recidivism in the future.⁶ Predictions have long been a part of criminal justice decision-making because of legitimate goals of protecting the public from those who have already been identified as offenders.⁷ Historically, risk predictions were generally based on gut instinct or the personal experience of the official responsible for making the relevant decision.⁸ Yet advances in behavioral sciences, the availability of big data, and improvements in statistical modeling have ushered in a wave of more empirically-informed risk assessment tools.

A. *The Rise of Algorithmic Risk Assessment in Criminal Justice*

The “evidence-based practices movement” is the now popular term to describe the turn towards using behavioral sciences data to improve offender

6. See Melissa Hamilton, *Risk-Needs Assessment: Constitutional and Ethical Challenges*, 52 AM. CRIM. L. REV. 231, 232 (2015) (“The evidence-oriented model utilizes the best data available from the empirical sciences to identify and classify individuals based on their potential future risk of reoffending, and then to manage offender populations accordingly”).

7. Jordan M. Hyatt et al., *Reform in Motion: The Promise and Perils of Incorporating Risk Assessments and Cost-Benefit Analysis into Pennsylvania Sentencing*, 49 DUQ. L. REV. 707, 724–25 (2011).

8. Cecelia Klingele, *The Promises and Perils of Evidence-Based Corrections*, 91 NOTRE DAME L. REV. 537, 556 (2015).

classifications.⁹ Scientific studies on recidivism outcomes are benefiting from the availability of large datasets (i.e., big data) tracking offenders post-release to statistically test for factors that correlate with recidivism.¹⁰ Risk assessment tool developers use computer modeling to combine factors of sufficiently high correlation and to weigh them accordingly with increasingly complex algorithms.¹¹ Broadly speaking, “[d]ata-driven algorithmic decision making may enhance overall government efficiency and public service delivery, by optimizing bureaucratic processes, providing real-time feedback and predicting outcomes.”¹² With such a tool in hand, criminal justice officials can more consistently input relevant data and receive software-produced risk classifications.¹³ Dozens of automated risk assessment tools to predict recidivism are now available and are quite popular.¹⁴

The utility of risk instruments has attracted energetic support from reputable policy centers, namely the Council of State Governments Justice Center,¹⁵ the Justice Management Institute,¹⁶ the Center for Effective Public Policy,¹⁷ the Vera Institute of Justice,¹⁸ and the Center for Court Innovation.¹⁹ News headlines and academic literature have also been expounding upon the potential harms and benefits generated by the government’s use of big data to predict the future risk posed by individuals.²⁰ Algorithmic risk assessment tools offer the ability to reduce mass

9. Faye S. Taxman, *The Partially Clothed Emperor: Evidence-Based Practices*, 34 J. CONTEMP. CRIM. JUST. 97, 97–98 (2018).

10. Kelly Hannah-Moffat, *Algorithmic Risk Governance: Big Data Analytics, Race and Information Activism in Criminal Justice Debates*, THEORETICAL CRIMINOLOGY (forthcoming 2019).

11. An algorithm refers to “computational procedures (which can be more or less complex) drawing on some type of digital data (‘big’ or not) that provide some kind of quantitative output (be it a single score or multiple metrics) through a software program.” Angèle Christin, *Algorithms in Practice: Comparing Web Journalism and Criminal Justice*, BIG DATA & SOC’Y 1, 2 (2017).

12. Bruno Lepri et al., *Fair, Transparent, and Accountable Algorithmic Decision-Making Processes*, 31 PHIL. & TECH. 611, 612 (2018).

13. J. Stephen Wormith, *Automated Offender Risk Assessment: The Next Generation or a Black Hole?*, 16 CRIMINOLOGY & PUB. POL’Y 281, 285 (2017).

14. See DANIELLE KEHL ET AL., ALGORITHMS IN THE CRIMINAL JUSTICE SYSTEM: ASSESSING THE USE OF RISK ASSESSMENTS IN SENTENCING 9 (Responsive Cmtys. Initiative, Harvard Law Sch. 2017).

15. COUNCIL OF STATE GOV’T, JUSTICE CTR., LESSONS FROM THE STATES: REDUCING RECIDIVISM AND CURVING CORRECTIONS COSTS THROUGH JUSTICE REINVESTMENT 6–7 (2013), http://csgjusticecenter.org/wp-content/uploads/2013/04/FINAL_State_Lessons_mbedit.pdf.

16. MAREA BEEMAN & AIMEE WICKMAN, THE JUSTICE MGMT. INST., RISK AND NEEDS ASSESSMENT 2 (2013).

17. CTR. FOR EFFECTIVE PUB. POLICY, BEHAVIOR MANAGEMENT OF JUSTICE INVOLVED INDIVIDUALS: CONTEMPORARY RESEARCH AND STATE-OF-THE-ART POLICY AND PRACTICE 19 (2015).

18. Memorandum from Vera Inst. of Justice to Del. Justice Reinvestment Task Force 1–2 (Oct. 12, 2011).

19. MICHAEL REMPEL, CTR. FOR COURT INNOVATION, EVIDENCE-BASED STRATEGIES FOR WORKING WITH OFFENDERS 1–2 (2014).

20. E.g., Andrew Guthrie Ferguson, *Big Data and Predictive Reasonable Suspicion*, 163 U. PA. L. REV. 327, 406–07 (2015); Crysta Jentile & Michelle Lawrence, *How Government Use of Big Data Can Harm Communities*, FORD FOUNDATION: EQUALS CHANGE BLOG (Aug. 30, 2016), <https://www.fordfoundation.org/ideas/equals-change-blog/posts/how-government-use-of-big-data-can-harm-communities/>; Sony Kassam, *Legality of Using Predictive Data to Determine Sentences Challenged in Wisconsin Supreme Court Case*, A.B.A. J. (June 27, 2016, 1:07 PM), http://www.abajournal.com/news/article/legality_of_using_predictive_data_to_determine_sentences_challenged_in_wisc.

incarceration by diverting low risk defendants from prison, while targeting greater supervision and services to those at higher risk.²¹

Many parties presume that algorithmic risk assessment tools developed on big data represent a transparent, consistent, and logical method for classifying offenders.²² The mathematical character of risk assessment suggests the ability to quantify the future and transport it into the present.²³ Evidence-based practices thereby present a welcome displacement of human instinct.²⁴ Risk assessment practices have been heavily oriented to back-end decisions, such as sentencing, early release decisions, and post-incarceration supervision.²⁵ More recent attention considers the potential benefits that automated risk assessment practices provide in pretrial settings.²⁶

B. Pretrial Risk Assessment

Algorithmic risk assessment informs such pretrial decisions as deferred adjudication and bail.²⁷ The basic idea of risk-informed decisions for pretrial purposes has a longer trajectory, being first approved by the United States Supreme Court in a 1987 opinion. In *United States v. Salerno*, the Court held that the practice of ordering pretrial detention based on an estimate of the individual defendant's future dangerousness is not facially unconstitutional.²⁸ Still, these predictions must be taken with care because of the potential consequences to individual rights. "Pretrial decision-making involves a fundamental tension between the court's desire to protect citizens from dangerous criminals, ensure that accused individuals are judged before the law, and minimize the amount of pretrial punishment meted out to legally innocent defendants."²⁹

The terminology has changed since *Salerno* from the vagueness of "future dangerousness" to a more refined perspective of "risk assessment."³⁰ Risk assessment

21. See Sarah L. Desmarais et al., *Performance of Recidivism Risk Assessment Instruments in U.S. Correctional Settings*, 13 PSYCHOL. SERVS. 206, 206–07 (2016) (arguing that risk assessment tools are producing more reliable results than unstructured risk assessments even with limits on such tools' effectiveness).

22. Hyatt et al., *supra* note 7, at 725.

23. See M. Roffey & S.Z. Kaliski, *To Predict or Not to Predict—That Is the Question*, 15 AFR. J. PSYCHIATRY 227, 227 (2012) (suggesting that risk assessment is "supposedly more objective" as a predictor of future risk as compared to psychiatric "predictions" due to its mathematical character).

24. Alfred Blumstein, *Some Perspectives on Quantitative Criminology Pre-JQC: And then Some*, 26 J. QUANTITATIVE CRIMINOLOGY 549, 554 (2010).

25. See Jessica M. Eaglin, *Constructing Recidivism Risk*, 67 EMORY L.J. 59, 67–68 (2017) (stating that risk assessment tools rely on "factors not necessarily connected to the criminal justice system").

26. See Kristin Bechtel et al., *A Meta-Analytic Review of Pretrial Research: Risk Assessment, Bond Type, and Interventions*, 42 AM. J. CRIM. JUST. 443, 444 (2017).

27. BEEMAN & WICKMAN, *supra* note 16, at 3.

28. *United States v. Salerno*, 481 U.S. 739, 751–52 (1987).

29. THOMAS BLOMBERG ET AL., *VALIDATION OF THE COMPAS RISK ASSESSMENT CLASSIFICATION INSTRUMENT 4* (2010).

30. Cf. KIRK HEILBRUN, *EVALUATION FOR RISK OF VIOLENCE IN ADULTS 18–19* (2009).

has evolved over the last three decades from unstructured clinical decisions to the actuarial form involving algorithmic processing.³¹

To improve the fairness and effectiveness of pretrial decisions, behavioral science experts encourage officials to use more objective criteria, such as those offered by evidence-based risk tools.³² Legal reformers generally welcome this practice as well. Risk assessment has become a foundation for the bail reform movement by offering a substitute to a long-standing dependence upon monetary bail.³³ Releasing more defendants who do not pose a substantial risk can alleviate the harms that money bail systems disproportionately wreak on poor and minority defendants.³⁴ At the same time, reducing the rate of pretrial detention prevents other negative consequences to individual offenders, as studies consistently show that pretrial detention is correlated with a greater likelihood of a guilty plea, a longer sentence, job loss, family disruption, and violent victimization in jail.³⁵

Notwithstanding the broad support and high hopes for algorithmic risk assessment practices in criminal justice, an investigative report publicized in 2016 called into question their objectivity and fairness.

C. *The ProPublica Study*

News journalists at ProPublica reported on statistical analyses the group had conducted involving a real dataset and a popular risk tool named COMPAS – the acronym for Correctional Offender Management Profiling for Alternative Sanctions. ProPublica investigators obtained through Freedom of Information Act requests the data on over 7,000 arrestees who were scored on COMPAS in a pretrial setting in a southern county of Florida.³⁶ ProPublica concluded that COMPAS discriminates against Blacks because its algorithm produced a much higher false positive rate for Blacks than Whites, meaning that it overpredicted high risk for Blacks.³⁷

COMPAS's corporate owner, Northpointe, quickly rejected such characterization.³⁸ After running their own statistical analyses on the same dataset ProPublica

31. See Klingele, *supra* note 8, at 551–67.

32. See Arthur Rizer & Caleb Watney, *Artificial Intelligence Can Make Our Jail System More Efficient, Equitable and Just*, 23 TEX. REV. L. & POL. 181, 184 (2018) (noting the potential benefits that artificial intelligence capacities could bring to the criminal justice system).

33. Megan Stevenson, *Assessing Risk Assessment in Action*, 103 MINN. L. REV. 303, 319–20 (2018).

34. John Logan Koepke & David G. Robinson, *Danger Ahead: Risk Assessment and the Future of Bail Reform*, 93 WASH. L. REV. 1725, 1769 (2018).

35. Laura I. Appleman, *Justice in the Shadowlands: Pretrial Detention, Punishment, & the Sixth Amendment*, 69 WASH. & LEE L. REV. 1297, 1320–21 (2012).

36. Julia Angwin et al., *Machine Bias*, PROPUBLICA (May 23, 2016), <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.

37. *Id.*

38. Northpointe re-branded with the trade name “equivant” (lower case intended) in January 2017. Press Release, equivant, Courtview, Constellation & Northpointe Re-brand to equivant (Jan. 9, 2017).

had compiled, Northpointe statisticians asserted that their results demonstrated COMPAS outcomes achieved predictive parity for Blacks and Whites.³⁹

It turns out a rather simple explanation accounts for the dispute: contrasting measures of algorithmic fairness. ProPublica touted the false positive rate, while Northpointe preferred an alternative measure called the positive predictive value.⁴⁰ As will be addressed later below, these measures are not synonymous and offer distinct, sometimes conflicting, impressions of a tool's accuracy.⁴¹

II. AUDITING THE BLACK BOX

For purposes here, a “black-box” tool refers to an algorithmic risk instrument which is not transparent about what is input into the software program and/or how the outputs are generated and quantified.⁴² This characterization is more probably appropriate in the case of an algorithmic instrument that is proprietary and its owner protects information based on a claim of trade secrets.⁴³ COMPAS, for example, is proprietary and its corporate owner declines to reveal its algorithm, which likely is a reason for ProPublica's interest in auditing it.⁴⁴

A. *Calls for Third Party Audits*

ProPublica's study certainly brought the issue of algorithmic fairness to the forefront in the popular media.⁴⁵ Questions are being raised in the scientific and policy communities that otherwise “reasonable algorithms” may fail to result in fair and equitable treatment of diverse populations.⁴⁶ There is even a new scientific literature on the topic named FATML, or “fairness, accountability and transparency in machine learning.”⁴⁷ There is some overlap among the FATML goals:

Fairness can . . . be related to the notion of transparency— the question of how much we are entitled to know about any automated system that is used to make or inform a decision that affects us. Hiding the inner workings of an algorithm from public view might seem preferable, to avoid anyone gaming the system. But without transparency, how can decisions be probed and challenged?⁴⁸

39. WILLIAM DIETERICH ET AL., NORTHPOINTE INC., COMPAS RISK SCALES: DEMONSTRATING ACCURACY EQUITY AND PREDICTIVE PARITY 1 (July 8, 2016).

40. Tafari Mbadiwe, *Algorithmic Injustice*, THE NEW ATLANTICS, Winter 2018, at 18.

41. See *infra* Section IV.B.3.d.

42. ROBYN CAPLAN ET AL., DATA & SOCIETY, ALGORITHMIC ACCOUNTABILITY: A PRIMER 2–3 (2018).

43. See SARAH TAN ET AL., DISTILL-AND-COMPARE: AUDITING BLACK-BOX MODELS USING TRANSPARENT MODEL DISTILLATION 1 (2017), <https://arxiv.org/pdf/1710.06169> (stating that black-box models are “typically proprietary or opaque”).

44. See *State v. Loomis*, 881 N.W.2d 749, 761 (Wisc. 2016).

45. See, e.g., Zhou, *supra* note 3.

46. OSONDE OSABA & WILLIAM WELSER IV, AN INTELLIGENCE IN OUR IMAGE: THE RISKS OF BIAS AND ERRORS IN ARTIFICIAL INTELLIGENCE 19 (Rand Corporation 2017).

47. HARSH GUPTA, CONSTITUTIONAL PERSPECTIVES ON MACHINE LEARNING 4 (2017), <https://osf.io/preprints/socarxiv/9v8js/>.

48. Sofia Olhede & Patrick J. Wolfe, *When Algorithms go Wrong, Who is Liable?*, 14 SIGNIFICANCE 8, 9 (2017).

Importantly, no formal mechanism in the law or in the sciences exists to consistently enforce any form of algorithmic accountability.⁴⁹ Thus, observers call for third party auditing, such as that exemplified by ProPublica's efforts, to engage in any form of scientific inquiry that may reveal information about the empirical validity and fairness of black-box tools.⁵⁰ Such data will be useful to legal practitioners and policymakers in considering or reevaluating the use of automated risk assessment to inform criminal justice decisions which carry significant consequences for individuals.⁵¹

Moreover, despite the many advantages of algorithmic assessment, risk profiling may fail to alleviate all of the harms of mass incarceration as some "scholars are suspicious that contemporary extensions of risk assessment and risk reduction will likely only reproduce, or may even exacerbate, the injustices of contemporary criminal justice policy under a more 'objective' guise."⁵² For instance, a White House Report on Big Data from the Obama administration lauded the public benefits of big data in criminal justice, but also promoted academic research into big data systems "to ensure that people are treated fairly."⁵³

A particularly acute focus is the concern that potential unfairness will likely fall predominately on already beleaguered minority groups. "The use of big data analytics might have disproportionate adverse impacts on protected classes, even when organizations do not intend to discriminate and do not use sensitive classifiers like race and gender."⁵⁴ Moreover, cross-disciplinary sharing is necessary because of difficulties in translation where data scientists are often not trained in law and policy, while civil rights experts in turn may not have statistical expertise.⁵⁵ In sum, many are just realizing that big data analytics can create civil rights problems,⁵⁶ such that the "patina of fairness" that otherwise seems to attach to big data algorithms may be unjustified.⁵⁷ Thus, to the extent that justice decisions may bring negative consequences upon defendants, it is particularly advisable to study whether and how an algorithmic tool disparately impacts protected groups.⁵⁸

49. See ROBYN CAPLAN ET AL., *supra* note 42, at 10.

50. *See Id.* at 10–11.

51. Jennifer Skeem et al., *Gender, Risk Assessment, and Sanctioning: The Cost of Treating Women Like Men*, 40 LAW & HUM. BEHAV. 580, 590 (2016).

52. Seth J. Prins & Adam Reich, *Can we Avoid Reductionism in Risk Reduction?*, 22 THEORETICAL CRIMINOLOGY 258, 259 (2018).

53. EXECUTIVE OFFICE OF THE PRESIDENT, BIG DATA: A REPORT ON ALGORITHMIC SYSTEMS, OPPORTUNITY, AND CIVIL RIGHTS 22–23 (2016), https://obamaWhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf.

54. Mark MacCarthy, *Standards of Fairness for Disparate Impact Assessment of Big Data Algorithms*, 48 CUMB. L. REV. 67, 80 (2017).

55. *Cf.* SOLON BAROCAS ET AL., BIG DATA, DATA SCIENCE, AND CIVIL RIGHTS 6–7 (2017), <https://arxiv.org/pdf/1706.03102> (arguing cross-training of policymakers and data scientists is necessary to ensure ethical use of "big data" in fields that implicate civil rights).

56. *Id.* at 1.

57. Anupam Chander, *The Racist Algorithm?*, 115 MICH. L. REV. 1023, 1034 (2017).

58. MacCarthy, *supra* note 54, at 67.

B. Black-Box Tools and Ethnic Minorities

The ProPublica study was concerned with Black minorities. This paper focuses on the ethnic minority group of Hispanics. This Article follows the tradition of the United States Census Bureau in classifying Hispanics as an ethnicity rather than a race.⁵⁹

Importantly, Hispanics comprise a significant proportion of the American population, making them a reasonable population to analyze. Plus, reasons exist to suspect that an algorithm may not assess an ethnic minority group very well.

“[A] transparent, facially neutral algorithm can still produce discriminatory results.”⁶⁰ Even if a particular tool is shown to perform well on its training sample(s), it is not advisable to simply transport that tool across to new populations and settings because of the potential for risk-relevant differences in offenders and the availability of rehabilitation-oriented services that can undermine the tool’s performance.⁶¹ Unfortunately, officials often disregard this advisory. Recent experience informs that “the application of risk knowledge is often haphazard: jurisdictions frequently deploy pre-existing screening tools in settings for which they were neither designed nor calibrated.”⁶²

Another reason for suspicion is that, despite their disproportionate presence in criminal justice statistics, minorities tend to be underrepresented in testing or validation samples for most risk assessment tools.⁶³ It is unreasonable, though, to assume that risk assessment tools will perform as well for culturally diverse minority groups.⁶⁴ A risk assessment process that presumes that risk tools are somehow universal, generic, or culturally-neutral may well be flawed on the following grounds:

The over or under estimation of risk that can ensue from this process is entirely plausible given (a) the potential omission of meaningful risk items specific to minority populations, (b) the inclusion of risk factors that are more relevant to White offenders, and (c) variation in the cross-cultural manifestation and expression of existing risk items.⁶⁵

For example, a risk tool may yield unequal results for cultural minorities if it fails to incorporate or otherwise consider their unique “behavioral practices and expectations, health beliefs, social/environmental experiences, phenomenology, illness

59. U.S. CENSUS BUREAU, RACE AND ETHNICITY (2017), <https://www.census.gov/mso/www/training/pdf/race-ethnicity-onepager.pdf>.

60. Chander, *supra* note 57, at 1024.

61. See Desmarais et al., *supra* note 21, at 207 (noting that the “performance of risk assessments may differ as a function of the characteristics of the research samples and study designs”).

62. Prins & Reich, *supra* note 52, at 260 (internal citations omitted).

63. Shepherd & Lewis-Fernandez, *supra* note 4, at 428.

64. *Id.*

65. *Id.* at 429.

narratives, deviant conduct, and worldview.”⁶⁶ It is not surprising, then, when risk tools are originally normed on largely White samples that at least some studies show risk tools provide more accurate predictions for Whites than other groups.⁶⁷

Regrettably, relatively few validation studies on ethnic minorities exist.⁶⁸ Two recent validation studies that have included an ethnic minority group may be telling. The data reported with revalidation studies of the federal Pretrial Risk Assessment and the Post-Conviction Risk Assessment tools show that each tool overpredicted or underpredicted recidivism for Hispanics for at least a few of the recidivism outcomes tested.⁶⁹ Despite the seeming importance of these findings, the authors do not expound upon these results.

With respect specifically to Hispanic Americans, academics have suggested that risk assessment tools may not perform well if they fail to “consider the centrality of family, acculturation strain, religiosity, gender role expectations, and culturally stoic responses to adversity” unique to this particular cultural group.⁷⁰

C. Legal Challenges

The legal and big data communities are beginning to realize that there may be legal implications in terms of disparate impact if algorithmic bias exists.⁷¹ One commentator, while acknowledging the potential for unfairness in risk assessment-led decisions in criminal justice, has suggested that the practice at least is better than uninformed judgments about dangerousness and that the adversarial process will encourage attorneys to challenge the tools, which may yield improvements to the science underlying them.⁷² So far there is very little evidence of such confrontations in the courts. An exception is a 2016 case, in which a defendant protested the black-box nature of COMPAS (the same tool in ProPublica’s research and used in the study presented herein).

In *State v. Loomis*, the defendant claimed the use of COMPAS in his sentencing proceeding constituted a due process violation because the software owner’s claims of trade secrets prevented him from challenging the tool’s validity or its

66. *Id.*

67. Jay P. Singh et al., *A Comparative Study of Violence Risk Assessment Tools: A Systematic Review and Metaregression Analysis of 68 Studies Involving 25,980 Participants*, 31 *CLINICAL PSYCHOL. REV.* 499, 510 (2011); Jay P. Singh & Seena Fazel, *Forensic Risk Assessment: A Metareview*, 37 *CRIM. JUST. & BEHAV.* 965, 978 (2010).

68. T. Douglas et al., *Risk Assessment Tools in Criminal Justice and Forensic Psychiatry: The Need for Better Data*, 42 *EUR. PSYCHIATRY* 134, 135 (2017).

69. Thomas A. Cohen & Christopher T. Lowenkamp, *Revalidation of the Federal PTRR: Testing the PTRR for Predictive Biases*, 46 *CRIM. JUST. & BEHAV.* 234, 255 (2019); Christopher T. Lowenkamp, *PCRA Revisited: Testing the Validity of the Federal Post Conviction Risk Assessment (PCRA)*, 12 *PSYCHOL. SERV.* 149, 155 tbl.5 (2015).

70. Shepherd & Lewis-Fernandez, *supra* note 4, at 429.

71. Osaba & Welsler, *supra* note 46, at 19.

72. David E. Patton, *Guns, Crime Control, and a Systemic Approach to Federal Sentencing*, 32 *CARDOZO L. REV.* 1427, 1456–58 (2011).

algorithmic scoring.⁷³ The Wisconsin Supreme Court found against Loomis, but with some caveats.⁷⁴ The majority ruled that any use of COMPAS must come with certain written cautions, including warnings regarding the proprietary (hence, secretive) nature of the tool and the potential that it may disproportionately classify minorities as high risk.⁷⁵ The court seemed marginally concerned with the fact that COMPAS had never been validated on a Wisconsin-based population. The court suggested that more or less weight might be put on a COMPAS score if such a localized validation study were to be conducted.⁷⁶

In contrast, a very recent decision by the Canadian Supreme Court found the lack of a relevant validation study dispositive. In *Ewert v. Canada*, the defendant, an indigenous minority, successfully protested the use of a different risk assessment tool (i.e., not COMPAS) because of the nonexistence of any validation studies of the tool specifically on indigenous groups.⁷⁷ The Canadian Supreme Court, ruling in Ewert's favor, determined that without evidence that the particular tool was free of cultural bias, it was unjust to use it on indigenous inmates.⁷⁸ The justices observed that "substantive equality requires more than simply equal treatment" as treating groups identically may itself produce inequalities.⁷⁹

Despite the divergence of the rulings in *Loomis* and *Ewert*, both courts supported the relevance of validation studies appropriate to the underlying population(s) on which the tool is to be used.⁸⁰ According to the 100 plus groups of civil rights and defense counsel organizations in their "Shared Statement of Civil Rights Concerns," the stakes are significant.⁸¹ "If the use of a particular pretrial risk assessment instrument by itself does not result in an independently audited, measurable decrease in the number of people detained pretrial, the tool should be pulled from use until it is recalibrated to cause demonstrably decarceral results."⁸² Moreover, this Shared Statement encourages defense attorneys to inspect any risk tool to review how it was created, the input factors used, weights assigned to the factors, thresholds for risk bins, and outcome data used in its development.⁸³

The next Section responds to the call for third-party auditing by academics to evaluate how a tool may perform on a minority group.

73. *State v. Loomis*, 881 N.W.2d 749, 753 (Wisc. 2016).

74. *Id.* at 769.

75. *Id.*

76. *Id.*

77. *Ewert v. Canada*, 2018 S.C.R. 30, para. 66 (S.C.C. June 13, 2018).

78. *Id.* at para. 63.

79. *Id.* at para. 54.

80. See *State v. Loomis*, 881 N.W.2d 749, 769 (Wisc. 2016); *Ewert*, 2018 S.C.R., at para. 66.

81. See AFRICAN AMERICAN MINISTERS IN ACTION ET AL., THE USE OF PRETRIAL "RISK ASSESSMENT" INSTRUMENTS: A SHARED STATEMENT OF CIVIL RIGHTS CONCERNS 10 (2018), <http://civilrightsdocs.info/pdf/criminal-justice/Pretrial-Risk-Assessment-Full.pdf> (showing over 100 signees).

82. *Id.* at 5.

83. *Id.* at 7.

IV. EVALUATING ALGORITHMIC FAIRNESS WITH ETHNICITY

This section presents the results of an original study of the fairness of COMPAS, a popular algorithmic risk assessment tool, employing the same real-world dataset that ProPublica compiled for its evaluation of outcomes for Blacks versus Whites.⁸⁴ The data and the tool will be briefly addressed. Then, a variety of empirical methods will assess the accuracy, validity, and predictive ability of the tool and determine whether algorithmic unfairness or disparate impact appear present.

A. *The Samples and the Test*

The primary dataset includes individuals arrested in Broward County, Florida who were scored on a COMPAS risk scale soon after their arrests in 2013 and 2014.⁸⁵ Notably, Broward County is among the top twenty largest American counties by population, thus improving the potential for a large and diverse sample set. The pretrial services division of the Broward County Sheriff's Department has been using COMPAS since 2008 to inform judicial determinations concerning pre-trial release.⁸⁶ This study uses two subsets of data, one of which tracks general recidivism ($n=6,172$) and the other violent recidivism ($n=4,020$). The follow-up recidivism period is two years.

COMPAS is a software application widely used across correctional institutions, offering a general recidivism risk scale and a violent recidivism risk scale.⁸⁷ The general recidivism risk scale contains about two dozen items related to age at first arrest, age at intake, criminal history, drug problems, and vocational/educational problems (e.g., employment, possessing a skill or trade, high school grades, suspensions).⁸⁸ The violent recidivism scale differs in that in lieu of criminal history generally it uses a history of violence, and instead of drug problems it incorporates factors related to a history of non-compliance (e.g., previous parole violations, arrested while on probation).⁸⁹

The COMPAS algorithms produce outcomes as decile scores of 1-10 with higher deciles representing greater predicted risk. COMPAS then subdivides decile

84. COMPAS is the "most successful and popular application of machine learning to offender risk assessment." Wormith, *supra* note 13, at 285 (citing developers using a decision tree model to educate the tool).

85. See generally Jeff Larson et al., *How We Analyzed the COMPAS Recidivism Algorithm*, PROPUBLICA (May 23, 2016), <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm> (explaining procedure used to generate the dataset). ProPublica has generously made the data available for other researchers to access at: <https://github.com/propublica/compas-analysis>.

86. BLOMBERG ET AL., *supra* note 29, at 15–16.

87. See NORTHPOINTE, PRACTITIONER'S GUIDE TO COMPAS CORE 2 (2015), http://www.northpointeinc.com/downloads/compas/Practitioners-Guide-COMPAS-Core-_031915.pdf.

88. See NORTHPOINTE, COMPAS CORE NORMS FOR ADULT INSTITUTIONS 80 tbl.3.41 (Feb. 11, 2014), https://epic.org/algorithmic-transparency/crim-justice/EPIC-16-06-23-WI-FOIA-201600805-WIDOC_DAL_norm_report021114.pdf (listing risk factors including employment, possessing a skill or trade, high school grades, and suspensions).

89. See NORTHPOINTE, *supra* note 87, at 38. PRACTITIONER'S GUIDE TO COMPAS CORE 28 (2015).

scores into three, ordinal risk bins: low (deciles 1-4), medium (deciles 5-7), and high (deciles 8-10).

These analyses generally followed the methodology of the ProPublica researchers in terms of defining what acts comprise general or violent recidivism, excluding cases with missing data, and excluding cases where the individuals were not scored on COMPAS in a timely manner.⁹⁰

B. Differential Validity Measures

Validity, while a technical term, simply means the extent to which a test properly reflects the concept it is designed to reflect.⁹¹ For purposes here, validity asks whether COMPAS adequately measures recidivism risk. The term *differential validity* applies when test validity varies between groups.⁹²

Multiple measures related to validity are available to judge the diagnostic and prognostic capabilities of an assessment tool. *Discrimination* reflects how well the tool distinguishes higher risk from lower risk (i.e., relative accuracy).⁹³ *Calibration* concerns how accurate the tool statistically estimates the outcome of interest (i.e., absolute predictive accuracy).⁹⁴

A few descriptive statistics of the samples are provided in Figure 1.

Figure 1. Descriptive Statistics (means or proportions)

Factor	General Recidivism Sample		Violent Recidivism Sample	
	Hispanic	Non-Hispanic	Hispanic	Non-Hispanic
Recidivism Rate	.37	.46***	.10	.17***
Decile Score	3.38	4.51***	2.64	3.33***
No. Prior Offenses	2.10	3.35***	1.70	2.52***
Age	35.02	34.49	36.14	35.70
Females	.16	.19	.17	.21
<i>n</i>	509	5663	355	3665

Notes: * $p < .05$, ** $p < .01$, *** $p < .001$.

90. See generally Larson et al., *supra* note 85.

91. MICHAEL G. MAXFIELD & EARL BABBIE, RESEARCH METHODS IN CRIMINAL JUSTICE AND CRIMINOLOGY 127 (2nd ed. 1998).

92. Christopher M. Berry et al., *Can Racial/Ethnic Subgroup Criterion-to-Test Standard Deviant Ratios Account for Conflicting Differential Validity and Differential Prediction Evidence for Cognitive Ability Tests?*, 87 J. OCCUPATIONAL & ORG. PSYCHOL. 208, 209 (2014).

93. L. Maaiké Helmus & Kelly M. Babchishin, *Primer on Risk Assessment and the Statistics Used to Evaluate its Accuracy*, 44 CRIM. JUST. & BEHAV. 8, 11 (2017).

94. *Id.*

Figure 1 shows that general and violent recidivism rates are substantially lower for Hispanics, twenty and forty-one percent lower than non-Hispanics respectively. Decile score outcomes on both scales for Hispanics were also substantially lower. Lower mean decile scores for Hispanics can partly be explained by their also having lower mean numbers of prior offences as each scale includes factors that count criminal history.

The table also includes a comparison on age and gender as these factors have been shown in previous studies to be strong predictors of recidivism.⁹⁵ Here, test comparison results between groups of the average age of offenders and their gender makeup are not statistically significant. Thus, these results generally indicate that any differential validity between these groups is not obviously a result of age or gender disparities.

Figure 2 includes statistics to explore the degree (i.e., the strength) of the relationship between COMPAS scores and recidivism for each scale and compares Hispanics with non-Hispanics.

Figure 2. COMPAS Failure Rates

Risk Level	General Recidivism Sample		Violent Recidivism Sample	
	Hispanic	Non-Hispanic	Hispanic	Non-Hispanic
--				
Low	.30	.32	.09	.11
Medium	.55	.55	.08	.27***
High	.57	.75***	.46	.49
AUC	.637	.714***	.641	.722**
<i>r</i>	.238	.372**	.148	.319***

Notes: * $p < .05$, ** $p < .01$, *** $p < .001$.

The tabular statistics in Figure 2 contain a host of relevant information about COMPAS' performance.

1. Validity of Risk Bins

Three rows display actual rates of recidivism by offenders classified in each of the COMPAS risk bins of low, medium, and high. The results show both intergroup and intragroup disparities. Significant intergroup differences in recidivism rates occur in the high risk classification for the general recidivism risk scale (57% for Hispanics versus 75% for non-Hispanics) and in the medium risk bin for the

95. See generally David E. Olson et al., *Comparing Male and Female Prison Releasees across Risk Factors and Postprison Recidivism*, 26 WOMEN & CRIM. JUST. 122, 123, 126 (2016) (listing studies).

violent recidivism scale (8% for Hispanics versus 27% for non-Hispanics). Thus, these two risk bins perform disparately regarding reoffending based on ethnicity.

For the non-Hispanics group, actual recidivism rates operate in a monotonic fashion as risk bins rise (low-medium-high) with substantively increasing rates of recidivism on each scale, thereby exemplifying favorable performance. Contrastingly, this is not the case for Hispanics. Highlighted within Table 2, one can observe that with the general recidivism scale, the recidivism rate for Hispanics is only slightly higher in the high risk bin than the medium bin, with the difference not statistically significant. Then with the violent recidivism scale, the rate for Hispanics unexpectedly decreases from the low to the medium bins. Hence, these measures confirm differential validity in that COMPAS fails to perform equivalently for these groups. Indeed, the three risk bin strategy collapses for Hispanics. (A two risk bin scheme appears more appropriate for Hispanics whereby the medium and high risk bins could be combined for general recidivism while a low and medium risk bin combination might better fit the data for violent recidivism).

2. *The Area Under the Curve*

The next row in Figure 2 involves a metric called the area under the curve (AUC)—it is derived from a statistical plotting of true positives and false positives across a risk tool's rating system.⁹⁶ More specifically, an AUC is a discrimination index that represents the probability that a randomly selected recidivist received a higher risk classification than a randomly selected non-recidivist.⁹⁷ AUCs range from 0–1.0, with .5 indicating no better accuracy than chance, and a 1.0, meaning perfect discrimination.⁹⁸ For this statistic, COMPAS decile scores are utilized, rather than the ordinal bins.

Risk assessment scholars often refer to AUCs of .56, .64, and .71 as the thresholds for small, medium, and large effect sizes, respectively.⁹⁹ Using these suggestions, it would mean that COMPAS operates effectively for both groups, albeit at disparate magnitudes. The AUCs would represent moderate effect sizes for Hispanics and large effect sizes for non-Hispanics. Notwithstanding, agreement on the strength of AUCs is not universal.¹⁰⁰ A more conservative conceptualization is that AUCs between .60 and .69 are poor, .70 to .79 are fair, .80 to .89 are good, and

96. Jay P. Singh, *Predictive Validity Performance Indicators in Violent Risk Assessment*, 31 BEHAV. SCI. & L. 8, 15 (2013).

97. Jay P. Singh et al., *Measurement of Predictive Validity in Violence Risk Assessment Studies*, 31 BEHAV. SCI. & L. 55, 64 (2013) (stating that 34 percent of “all articles” that use AUC analysis define AUC in this way).

98. Martin Rettenberger et al., *Prospective Actuarial Risk Assessment: A Comparison of Five Risk Assessment Instruments in Different Sexual Offender Subtypes*, 54 INT'L J. OFFENDER THERAPY & COMP. CRIMINOLOGY 169, 176 (2010).

99. Helmus & Babchishin, *supra* note 93, at 12.

100. Jay P. Singh, *Five Opportunities for Innovation in Violence Risk Assessment Research*, 1 J. THREAT ASSESSMENT & MGMT. 179, 181 (2014).

over .90 are excellent.¹⁰¹ Based on these standards, the AUCs for Hispanics would be judged as poor and for non-Hispanics as fair.

Still, the AUC has serious limitations as discussed elsewhere and thus cannot present a holistic portrait of a tool's abilities.¹⁰² Unfortunately, the AUC is too commonly misinterpreted as measuring calibration accuracy; but a higher AUC does not mean more accurate prospective prediction.¹⁰³ Further, the AUC cannot calculate how well an instrument selects those at high risk.¹⁰⁴ For example, the AUC could be very high even if no recidivists were ranked as high risk. To use a hypothetical, the AUC for COMPAS would actually reflect perfect accuracy (AUC=1.0) where all recidivists were classified Decile 2 and all non-recidivists as Decile 1 (i.e., all were classified as low risk), with very little distinction considering the decile scale ranges from 1 to 10.

Importantly, substantial group differences indicated by corresponding *z*-tests for AUCs will reveal the presence of differential discrimination accuracy in terms of degree.¹⁰⁵ Here, the AUCs indicate that COMPAS is a weaker discriminatory tool for Hispanics than non-Hispanics as the AUC measures on both the general and violent recidivism tools are lower for Hispanics, and with statistical significance (*z*-tests of the differences in proportions indicate levels of $p=.000$ and $p=.001$ for general and violent recidivism, respectively). These results on the AUCs are another indicator that COMPAS exhibits differential accuracy in its degree of discrimination utility on both scales relative to Hispanic ethnicity.

3. Correlations

The final row in Figure 1 contains correlation coefficients representing another barometer of the strength of COMPAS decile scores relative to recidivism. Like the AUC, correlations test discrimination ability.¹⁰⁶ Separate statistical runs (not shown in Figure 2) find that each of the four individual correlation coefficients are statistically significant at $p=.005$ or below. These individual correlations thus signify that the scales achieve some statistical strength relative to recidivism in both groups. Notwithstanding, correlation coefficients can signify differential validity as well.¹⁰⁷ Here, comparing correlation statistics shows that COMPAS scales are more weakly associated with general and violent recidivism for Hispanics, and that these differences are statistically significant. In sum, this result adds to the cumulative evidence of differential validity for COMPAS regarding Hispanics.

101. Helmus & Babchishin, *supra* note 93, at 12.

102. Singh, *supra* note 96, at 16–18.

103. *Id.* at 16.

104. *Id.* at 17.

105. Anthony W. Flores et al., *False Positives, False Negatives, and False Analyses: A Rejoinder to "Machine Bias: There's Software Used Across the Country to Predict Future Criminals. And It's Biased Against Blacks,"* 80 FED. PROBATION 38, 41 (2016).

106. Helmus & Babchishin, *supra* note 93, at 11.

107. Berry et al., *supra* note 92, at 209.

a. Test Bias

A well calibrated instrument is one that is also “free from *predictive bias*.”¹⁰⁸ Differential prediction demonstrates group disparities in predictive ability.

Researchers examining group bias in psychological testing in education have standardized a methodology to empirically study it, with the endorsement of the American Psychological Association.¹⁰⁹ Group bias represents test bias, which refers to the existence of systematic errors in how a test measures members of any group.¹¹⁰ This gold standard for evaluating test bias involves a series of nested models of regression equations involving the test, the group(s) of interest, and an interaction term (test*group) as predictors of test outcomes.¹¹¹ Basically, a regression analysis is a statistical method to evaluate the relationship between one or more predictors with a response (outcome) variable.¹¹² Then, an interaction term refers to the product of two predictor variables (here, test and group) to determine whether the effect on the outcome of either predictor is moderated by the presence of the other (i.e., changes its strength or direction of influence).¹¹³

This nested models method detects group differences in the form of the relationship between the test and the outcome in terms of intercepts and slopes¹¹⁴ in order to reveal differential prediction.¹¹⁵ The rule of thumb in the psychological assessment field is that a significant group difference in either the intercept or the slope represents that a single regression equation for the groups combined will predict inaccurately for one or both groups, and therefore a separate equation for each group must be considered.¹¹⁶ Unequal intercepts or slopes also signify *disparate impact*.¹¹⁷ A system that treats persons unfairly suggests disparate impact may be present without requiring evidence of any discriminatory intent.¹¹⁸ Selected

108. Alexandria Chouldechova, *Fair Prediction with Disparate Impact: A Study of Bias in Recidivism Prediction Instruments*, 5 *BIG DATA* 153, 154 (2017) (emphasis in original).

109. Nathan R. Kuncel & Davide M. Kleiger, *Predictive Bias in Work and Educational Settings*, in *THE OXFORD HANDBOOK OF PERSONNEL ASSESSMENT* 462, 463 (Neil Schmitt ed., 2012) (confirming endorsements also from the National Council on Measurement in Education and the American Educational Research Association).

110. Adam W. Meade & Michael Fetzer, *Test Bias, Differential Prediction, and a Revised Approach for Determining the Suitability of a Predictor in a Selection Context*, 12 *ORG. RES. METHODS* 738, 738 (2009).

111. Jeanne A. Teresi & Richard N. Jones, *Bias in Psychological Assessment and Other Measures*, in *APA HANDBOOK OF TESTING AND ASSESSMENT IN PSYCHOLOGY* 139, 144 (vol. 1 2013).

112. RONET D. BACHMAN & RAYMOND PATERNOSTER, *STATISTICS FOR CRIMINOLOGY AND CRIMINAL JUSTICE* 675 (1997).

113. JAMES JACCARD, *INTERACTION EFFECTS IN LOGISTIC REGRESSION* 12 (2001).

114. Jennifer L. Skeem & Christopher T. Lowenkamp, *Race, Risk, and Recidivism: Predictive Bias and Disparate Impact*, 54 *CRIMINOLOGY* 680, 692 (2016).

115. Meade & Fetzer, *supra* note 110, at 740.

116. Cecil R. Reynolds & Lisa A. Suzuki, *Bias in Psychological Assessment: An Empirical Review and Recommendations*, in *HANDBOOK OF PSYCHOLOGY* 82, 101 (Irving B. Weiner ed., 2003).

117. Meade & Fetzer, *supra* note 110, at 741.

118. See Andrew D. Selbst, *Disparate Impact in Big Data Policing*, 52 *GA. L. REV.* 109, 121–22 (2017) (explaining how different reports have shown unintentional bias).

researchers in criminal justice have recently begun to apply this methodological practice of nested models to evaluate group bias in recidivism risk tools.¹¹⁹

The nested model structure here utilized variables labeled as Hispanic (coded as Hispanic=1, non-Hispanic=0), the COMPAS decile score, and an interaction between them as Hispanic*COMPAS decile score. A four model structure is employed with the outcome variable being recidivism. Model 1 tests only the Hispanic variable; Model 2 tests just COMPAS Decile score; Model 3 includes both the Hispanic and COMPAS decile score variables; and Model 4 retains Hispanic and COMPAS decile score while adding the interaction term. The general recidivism model utilizes the COMPAS general recidivism scale while the violent recidivism model relies upon the COMPAS violent recidivism scale.

Figures 3 and 4 provide the relevant logistic regression equation results for general recidivism and violent recidivism, respectively. *Logistic* regression applies when investigating an association between one or more predictor variables with an outcome of interest that is dichotomous in nature (e.g., recidivist=yes/no).¹²⁰ The logistic coefficients have been translated in the tables as odds ratios for interpretive purposes.

Figure 3. Logistic Regression Predicting the Odds of General Recidivism

Predictor	Model 1	Model 2	Model 3	Model 4
Hispanic	.686***	---	.914	1.284
Decile	---	1.323***	1.322***	1.331***
Hispanic*Decile Interaction	---	---	---	.911*
Constant	.861	.239	.242	.235
-2LL	8490.49	7650.88	7650.11	7644.43
χ^2	15.93	855.53	856.30	861.99
<i>n</i> =6,172				

Notes: * $p < .05$, ** $p < .01$, *** $p < .001$. Coefficients represent odds ratios.

Model 1 indicates that Hispanics are significantly less likely to engage in any act of recidivism. Model 2 supports the overall utility of COMPAS for the groups combined in that the odds of general recidivism increases 32% for every one unit increase in COMPAS decile score. In Models 3 and 4, when adding the decile score, Hispanic ethnicity is no longer statistically significant. While the odds are

119. See Skeem et al., *supra* note 51, at 585 (applying the aforementioned methodological practice).

120. See generally FRED C. PAMPEL, LOGISTIC REGRESSION 1 (2000) (describing the advantages of using dichotomous indicators).

less for Hispanics in Model 3, such that there is evidence of overprediction for them, it is not statistically significant ($p=.380$). However, the interaction term in Model 4 is statistically significant ($p=.015$). This means that Hispanic ethnicity significantly moderates the relationship between COMPAS decile score and general recidivism. As the interaction term is below one, the regression slope is less steep for Hispanics, revealing that as the COMPAS score increases, it has weaker influence on predicting recidivism for Hispanics. In other words, COMPAS does not predict as strongly for Hispanics and thus “is not as valid a predictor for that subgroup.”¹²¹ Further, unequal slopes reflect differential test prediction and thus present as disparate impact.

Figure 4. Logistic Regression Predicting the Odds of Violent Recidivism

Predictor	Model 1	Model 2	Model 3	Model 4
Hispanic	.540***	---	.678*	1.033
Decile	---	1.379***	1.375***	1.384***
Hispanic*Decile Interaction	---	---	---	.891
Constant	.202	.056	.058	.057
-2LL	3550.99	3207.91	3203.36	3201.18
χ^2	13.00	356.08	360.63	362.80
<i>n</i> =4,020				

Notes: * $p<.05$, ** $p<.01$, *** $p<.001$. Coefficients represent odds ratios.

With respect to the COMPAS violent risk scale, there are some consistencies with the general risk scale just discussed. In Figure 4, Model 1 shows that Hispanics are significantly less likely to commit a violent offense. The utility of the COMPAS violent recidivism scale generally is supported in Model 2. For every one unit increase in the violent decile score, the odds of violent recidivism increases by 38%.

The results here diverge from the regression models for general recidivism in two ways. The Hispanic variable in Model 3 remains statistically significant ($p=.04$), meaning that there are discrepancies in the intercepts of the regression lines for violent recidivism. The lower intercept means overprediction of Hispanics on the violent recidivism scale.¹²² But as the Model 4 interaction is not significant ($p=.143$), there is no detectable variation in the slopes.

121. See Christopher M. Berry, *Differential Validity and Differential Prediction in Cognitive Ability Tests*, 2 ANN. REV. ORG. PSYCHOL. & ORG. BEHAV. 435, 443 (2015).

122. See *id.*

In sum, the nested regression models reveal differential predictive validity in COMPAS based on Hispanic ethnicity, though the form varies. For general recidivism, there is asymmetry in the slopes, while for the violent recidivism scale the problem is unequal intercepts. In both cases, though, differential predictive validity is shown and COMPAS can be challenged as presenting unfair and biased algorithmic results based on Hispanic ethnicity. Test bias also symbolizes that while COMPAS scores may have some meaning within groups, comparisons across groups are problematic.¹²³

One may wonder if these regression models are too simplistic. For example, other researchers have found that gender, age, and criminal history are statistically significant factors in recidivism risk.¹²⁴ Because of this, separate regression models (not offered herein) were run to include controls for gender, age, and number of prior counts. The results were substantially the same: unequal slopes for the general recidivism scale and unequal intercepts for the violent recidivism scale.

b. Differential Prediction via E/O Measures

An alternative perspective on differential prediction looks to whether expected rates of recidivism are the same as observed recidivism rates in action, and that this is the case irrespective of group membership.¹²⁵ Expected rates are generated from logistic regression models whereas observed rates are the actual recidivism statistics. We can show the differences by plotting expected rates of recidivism against the observed rates at every COMPAS decile score (with expected versus observed referred to herein as “E/O”).

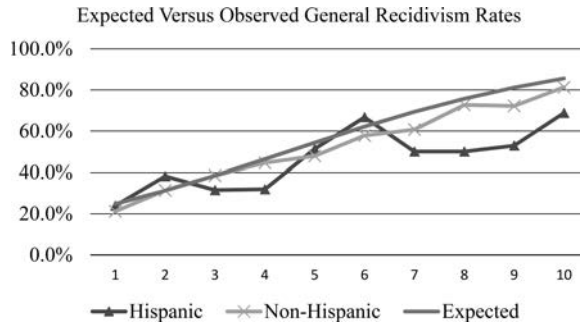
For each COMPAS scale, a single expected rate line was computed using Model 2 statistics for general and violent recidivism from Figures 3 and 4, respectively. The single expected rate line represents the expected recidivism rate at each decile score, regardless of group. Two lines of observed rates of recidivism are presented (one for Hispanics and the other for non-Hispanics). Figures 5 and 6 provide visual bar graphs of these E/O plots for general recidivism and violent recidivism, respectively.

123. See Meade & Fetzer, *supra* note 110, at 738.

124. See Nicholas Scurich & John Monahan, *Evidence-Based Sentencing: Public Openness and Opposition to Using Gender, Age, and Race as Risk Factors for Recidivism*, 40 *LAW & HUM. BEHAV.* 36, 37 (2016) (explaining that criticism of using these factors has not been based on their failure to statistically correlate with recidivism); Jennifer Skeem et al., *Gender, Risk, Assessment, and Sanctioning: The Cost of Treating Women Like Men*, 40 *LAW & HUM. BEHAV.* 580, 587 (2016) (finding that the difference between genders in PCRA scores “is small, but possibly meaningful”). See generally John Monahan et al., *Age, Risk Assessment, and Sanctioning*, 41 *LAW & HUM. BEHAV.* 191 (2017) (noting that such models might not treat the aforementioned factors accurately enough).

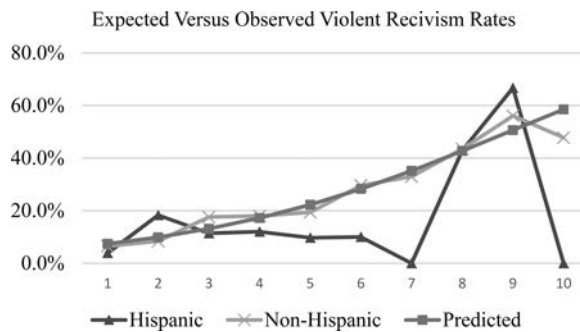
125. Richard Berk, *Accuracy and Fairness for Juvenile Justice Risks Assessments*, 16 *J. EMPIRICAL LEG. STUD.* 175, 181 (2019).

Figure 5. E/O General Recidivism Plots by Ethnicity



The solid line represents expected recidivism rates by decile regardless of grouping. The observed plotting for general recidivism outcomes for non-Hispanics moderately tracks the single, expected rates slope. The observed plotting for Hispanics, though, does not present a similar shape and, indeed, fails to reflect a consistent linear increase in actual recidivism rates as decile scores rise. This plotting visually represents differential prediction for COMPAS in the general recidivism scale. Figure 5 also graphically depicts accuracy errors for Hispanics.

Figure 6. E/O Violent Recidivism Plots by Ethnicity



The plotting in Figure 6 follows the same patterns just discussed for general recidivism rates. Notably, the discrepancies are more dramatic. Actual recidivism rates of violent reoffending for Hispanics oddly decline from decile 2 to a 0% recidivism rate at decile 7, then increase substantially in deciles 8 and 9, before plummeting back to 0% in decile 10. Hence, COMPAS violent decile score does not have a linear relationship to violent recidivism for Hispanics, thus undermining its performance for this group. This visual plotting also confirms differential prediction for COMPAS in its violent recidivism scale regarding Hispanic ethnicity.

So far, comparisons of risk bin outcomes, AUCs, and correlation coefficients reflect differential validity, while logistic regressions and E/O plots present differential prediction for COMPAS. Some argue that these are not the only appropriate measures of a test's ability to classify individuals, to predict risk, or to reflect unfairness.¹²⁶

c. Mean Score Differences

It was shown earlier that average COMPAS decile scores on the general and violent recidivism scales significantly varied by Hispanic ethnicity. Yet it is too simplistic to assert that mean score differences signify test bias (even though it violates the algorithmic fairness notion of statistical parity¹²⁷) as there may be some risk-relevant differences between groups.¹²⁸ Still, closely-related fairness indices are indicative of *differential calibration*—group variances in calibration—and potentially disparate impact. The algorithmic fairness concept of “balance for the positive class” refers to requiring that the mean test score for those in the positive class—here, meaning recidivists—be the same across groups.¹²⁹ Correspondingly, “balance for the negative class” requires equal mean test score for those in the negative class, i.e., non-recidivists. A test is well-calibrated with respect to these definitions of algorithmic fairness if a particular score is associated with the same likelihood of the outcome occurring, regardless of group membership.¹³⁰ Figure 7 presents mean COMPAS scores comparing recidivists and non-recidivists on both the general and violent recidivism scales.

Figure 7. Mean Decile Scores

	General Recidivism Scale		Violence Recidivism Scale	
	Recidivists	Non-Recidivists	Recidivists	Non-Recidivists
Hispanic	4.17	2.92***	3.54	2.54**
Non-Hispanic	5.65	3.53***	5.03	2.98***

Notes: * $p < .05$, ** $p < .01$, *** $p < .001$.

All four comparisons of mean decile score differences between groups are statistically significant at the level of $p = .005$ or below. Mean scores for recidivist and

126. Richard W. Elwood, *Calculating Probability in Sex Offender Risk Assessment*, 62 INT'L J. OFFENDER THERAPY & COMP. CRIMINOLOGY 1262, 1264 (2018).

127. JON KLEINBERG ET AL., INHERENT TRADE-OFFS IN THE FAIR DETERMINATION OF RISK SCORES 11–12 (2017), <https://arxiv.org/abs/1609.05807>.

128. Russell T. Warne et al., *Exploring the Various Interpretations of “Test Bias,”* 20 CULTURAL DIVERSITY & ETHNIC MINORITY PSYCHOL. 570, 572 (2014).

129. Kleinberg et al., *supra* note 127, at 2 (emphasis in original).

130. Chouldechova, *supra* note 108, at 154.

non-recidivist Hispanics are substantially lower than non-Hispanics. Practical differences exist as well. Mean scores for Hispanics exhibit much less variance between the recidivists and non-recidivists in both scales, being approximately one decile apart. The means for the non-Hispanics are at least two deciles apart, thus signifying greater mean score differences that distinguish non-Hispanic recidivists from non-recidivists in scoring to a greater degree. Overall, the metrics in Figure 7 signify a failure in balances for the positive and negative classes that again indicate algorithmic unfairness and disparate impact of COMPAS scoring based on Hispanic ethnicity.

d. Classification Errors

Additional computations measuring algorithmic accuracy and fairness are popular in the behavioral sciences and machine learning literatures. A few of these can provide further insight into possible disparate impact, which can occur when accuracy and error rates between groups are unequal.¹³¹

The first two are the true positive rate (TPR) and the true negative rate (TNR), representing a high risk and a low risk discrimination metric, respectively.¹³² The TPR is alternatively titled “sensitivity in the field of statistics” and represents the non-error rate for the recidivists.¹³³ The TNR is alternatively titled “specificity” and represents the non-error rate for the non-recidivists.¹³⁴

The TPR and TNR are *retrospective* in nature. Two metrics which more appropriately measure *prospective* predictive accuracy, and thereby are more important to practitioners who are interested in the predictive validity of risk tools, are the positive predictive value (PPV) and negative predictive value (NPV).¹³⁵ The PPV represents the probability that a high risk score will be correct, i.e., the proportion of high risk predictions who were recidivists.¹³⁶ The NPV then is the proportion of those classified as low risk who did not recidivate. The PPV is a high-risk calibration measure while the NPV is a low risk calibration measure.¹³⁷

These calculations require that the sample be divided into two groupings: one representing individuals predicted to be recidivists and the other non-recidivists. Regarding COMPAS, researchers typically opt to lump together COMPAS’ low and medium risk bins into one group, merge the medium and high-risk bins together, or, for better measure, offer both.¹³⁸ Figures 8 and 9 include the two different cutpoints for dichotomizing the COMPAS general and violent recidivism scales, respectively.

131. GEOFF PLEISS ET AL., ON FAIRNESS AND CALIBRATION 2, arxiv.org/pdf/1709.02012.pdf.

132. Singh, *supra* note 96, at 9.

133. Kristian Linnet et al., *Quantifying the Accuracy of a Diagnostic Test or Marker*, 58 CLINICAL CHEMISTRY 1292, 1296 (2012).

134. *Id.*

135. Singh, *supra* note 96, at 12.

136. ROBERT H. RIFFENBURGH, STATISTICS IN MEDICINE 254 Table 15.13 (2013).

137. Singh, *supra* note 96, at 11 Figure 1.

138. *Cf.* Flores et al., *supra* note 105, at 42.

Figure 8. Measures for General Recidivism

Measure	Low v. Medium/High		Low/Medium v. High	
	Hispanic	Non-Hispanic	Hispanic	Non-Hispanic
Discrimination				
TPR	.42	.63***	.14	.31***
TNR	.81	.69***	.94	.91
Calibration				
PPV	.56	.63**	.57	.75***
NPV	.70	.68	.65	.61

Notes: * $p < .05$, ** $p < .01$, *** $p < .001$.

In Figure 8, five of the eight comparisons are statistically significant, again confirming differential validity and differential prediction. With respect to the True Positive Rate, the tool is significantly better at discriminating the recidivists from the non-recidivists for non-Hispanics at both cutpoints. Indeed, the general recidivism TPRs for non-Hispanics are approximately 50% and 120% better than for Hispanics. The low TPRs for Hispanics are actually quite poor from quantitative and qualitative perspectives. For example, at the higher cutpoint (low/medium versus high) the TPR is only 14%, meaning the test is wrong at least eight out of ten times for Hispanics.

Likewise, the PPV differences show that COMPAS is a better predictor of recidivism for non-Hispanics at both cutpoints. For example, at the higher cutpoint, of those that were predicted at high risk of general recidivism, 57% of Hispanics were reoffenders while 75% of non-Hispanics were correctly classified. The PPVs for Hispanics at both cutpoints are also considered poor, indicating overprediction.¹³⁹

Interestingly one need not here choose between the preference of ProPublica for the TPR or Northpointe (the owner of COMPAS) regarding the PPV, or which cutpoint to use. COMPAS on any of those measures performs far worse at predicting recidivists for Hispanics. The results thus suggest algorithmic unfairness.

In terms of identifying non-recidivists, one of the four measures is statistically significant. The tool performs modestly better at identifying non-recidivists among Hispanics.

139. See James C. DiPerna et al., *Broadband Screening of Academic and Social Behavior, in UNIVERSAL SCREENING IN EDUCATIONAL SETTINGS* 223, 235, 239 (R.J. Kettler et al. eds., 2014).

Figure 9. Measures for Violent Recidivism

Measure	Low v. Medium/High		Low/Medium v. High	
	Hispanic	Non-Hispanic	Hispanic	Non-Hispanic
<i>Discrimination</i>				
TPR	.29	.54***	.14	.21**
TNR	.81	.77*	.98	.96
<i>Calibration</i>				
PPV	.14	.32***	.45	.49
NPV	.91	.89	.91	.86**

Notes: * $p < .05$, ** $p < .01$, *** $p < .001$.

Similar results are shown for violent recidivism in Figure 9. COMPAS generally performs much better from discrimination and calibration perspectives on classifying violent recidivists for non-Hispanics, with one caveat of the PPV at the higher cutoff where the PPV was still higher for non-Hispanics but not to a statistically significant degree. The violent recidivism TPRs for non-Hispanics are 86% and 50% better at the lower and higher cutpoints, respectively. The low PPVs for Hispanics again support the tendency of the tool to overclassify them. At the lower cutpoint, the predictions of violent recidivism for Hispanics were wrong at least 8 out of 10 times. Then, like the general recidivism scale, the TPRs and PPVs at both cutpoints indicate COMPAS achieves weaker accuracy at selecting recidivists for Hispanics. One need not choose between the TPR or PPV here: either definition of algorithmic fairness indicates unfairness to Hispanics.

Like the general recidivism scale, the tool performs modestly better at selecting violent non-recidivists among Hispanics.

e. Limitations

Several limitations of this study should be mentioned. The single site (i.e. Broward County, FL) limits generalization of results. This study relied upon archival data, and it is thereby possible for there to have been systematic errors in data collection that are not observable on secondary data analysis. Recidivism outcomes were from official records and thus will not include undetected crimes. The dataset did not include interrater reliability scores that would confirm the dependability of COMPAS scoring across evaluators and over time. It would have been preferable to control for aspects of supervision as pretrial services/conditions may moderate reoffending rates, but such an option is also not available in this secondary data analysis.

CONCLUSION

Algorithmic risk assessment holds promise in informing decisions that can reduce mass incarceration by releasing more prisoners through risk-based

selections that consider public safety. Yet caution is in order where presumptions of transparency, objectivity, and fairness of the algorithmic process may be unwarranted. Calls from those who heed such caution for third party audits of risk tools led to the study presented herein. This study rather uniquely focused on the potential of unfairness for Hispanics.

Using multiple definitions of algorithmic unfairness, results consistently showed that COMPAS, a popular risk tool, is not well calibrated for Hispanics. The statistics presented evidence of differential validity and differential predictive ability based on Hispanic ethnicity. The tool fails to accurately predict actual outcomes in a linear manner and overpredicts risk for Hispanics. Overall, there is cumulative evidence of disparate impact. It appears quite likely that factors extraneous to those scored by the COMPAS risk scales related to cultural differences account for these results.¹⁴⁰ This information should inform officials that greater care should be taken to ensure that proper validation studies are undertaken to confirm that any algorithmic risk tool used is fair for its intended population and subpopulations.

140. See Meade & Fetzer, *supra* note 110, at 740–41.