

May 2025

Exploring Mutable Characteristics and Discriminatory Perceptions in Justice Systems

Daniel L. Chen
Harvard University

Follow this and additional works at: <https://lawandinequality.org/>

Recommended Citation

Daniel L. Chen, *Exploring Mutable Characteristics and Discriminatory Perceptions in Justice Systems*, 43 L. & INEQUALITY 131 (2025).

Available at: <https://scholarship.law.umn.edu/lawineq/vol43/iss3/10>

Exploring Mutable Characteristics and Discriminatory Perceptions in Justice Systems

Daniel Chen[†]

Both sophisticated analyses of historical evidence and simple real-world observation indicate that judicial decisions demonstrate systematic racial and gender bias. For example, Republican-appointed federal judges sentence Black defendants more severely and female defendants more leniently.¹ Federal judges behave more politically before presidential elections,² especially for judges residing in states with close races in presidential elections.³ A judge's political party of appointment can be predicted by the citations they choose to motivate their decisions.⁴

[†]. Daniel L. Chen is an Evelyn Green Davis Fellow at the Radcliffe Institute for Advanced Study at Harvard, the Director of Research at the French National Centre for Scientific Research (CNRS), a Professor at the Toulouse School of Economics (TSE), a Senior Fellow at the Institute For Advanced Study in Toulouse (IAST), and a Co-Chair of Moral AI at Artificial and Natural Intelligence Toulouse Institute (ANITI). Chen received his B.A. (summa cum laude) and M.S. from Harvard University in Applied Mathematics and Economics, completed his Economics Ph.D from MIT, and obtained a J.D. from Harvard Law School. Chen is the founder of oTree Open Source Research Foundation. oTree is the leading open-source ecosystem for conducting strategic interaction and survey experiments in the lab, online, and in the field. He is the Lead Principal Investigator for AMICUS (Analytical Metrics for Informed Courtroom Understanding and Strategy), a team of computer scientists, economists, and lawyers aiming to revolutionize how justice systems' legitimacy and equality are measured, understood, and enhanced. They use cutting-edge research techniques to measure efficiency, quality, and access to justice and impacts on downstream outcomes such as economic growth, conflict and violence, and corruption. Chen's research is in law and economics. Many of his recent papers focus on judicial bias, AI, and court reform.

1. Alma Cohen & Crystal S. Yang, *Judicial Politics and Sentencing Decisions*, 11 AM. ECON. J. ECON. POL'Y 160 (2019).

2. Carlos Berdejó & Daniel L. Chen, *Electoral Cycles Among US Courts of Appeals Judges*, 60 J.L. & ECON. 479, 492 (2017).

3. Daniel L. Chen, *Priming Ideology: Why Presidential Elections Affect U.S. Judges* 18 (SSRN Working Paper No. 2816245, 2019), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2816245 [<https://perma.cc/34KE-84SW>].

4. Wei Lu & Daniel L. Chen, *Motivated Reasoning in the Field: Polarization of Prose, Precedent, and Policy in U.S. Circuit Courts, 1891–2013*, PLoS ONE, Mar. 3, 2025, at 1, 4, <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0318790> [<https://perma.cc/8WXB-2WFH>].

In their examination of these issues, Naomi Cahn, June Carbone, and Nancy Levit's *Fair Shake* probes similar patterns of bias across diverse sectors.⁵ They explore how deeply-ingrained societal prejudices shape outcomes for marginalized groups and argue that systems often fail to provide equitable treatment to all, even when impartiality is purportedly a guiding principle.⁶ Like the judges whose behaviors fluctuate with political pressures, institutions and individuals—as described in *Fair Shake*—often operate on biases rooted in history, culture, and deeply internalized norms.⁷ The book presents a tapestry of narratives demonstrating how systemic inequities manifest not only in legal contexts but also in everyday life, perpetuating cycles of disadvantage.⁸

The data on the justice system reveals sustained bias, both in arbitrary decisions that correspond with the identity of the assigned judge⁹ (e.g., a judge's racial identity is predictive of the disparities in their sentencing decisions) and also those that correspond with trivialities like whether or not a judge's hometown football team won or lost,¹⁰ whether a defendant shares with a judge the same first letter in their name,¹¹ whether it is the litigant's birthday,¹² etc. Furthermore, minoritized groups consistently bear the punitive brunt of these deviations from objectivity.

As we grapple with bias in our societies, one controversial yet crucial battleground has emerged: artificial intelligence (AI). The digital world mirrors the analog, and our systems of AI are susceptible to our inherent biases. However, there is hope; unlike humans, AI is fundamentally flexible. It can be reprogrammed and adjusted to mitigate biases, a process far more direct than attempting to reshape deeply-ingrained human prejudices.

5. NAOMI CAHN, JUNE CARBONE & NANCY LEVIT, *FAIR SHAKE: WOMEN AND THE FIGHT TO BUILD A JUST ECONOMY* (2024)

6. *Id.*

7. *Id.*

8. *Id.*

9. Abrams D., Bertrand M. & Mullainathan S., *Do Judges Vary in Their Treatment of Race?*, 41 J. LEGAL STUD. 347 (2012).

10. Özkan Eren & Naci Mocan, *Emotional Judges and Unlucky Juveniles*, 10 AM. ECON. J. APPL. ECON. 171, 199–200 (2018); Markus Loecher & Daniel L. Chen, *Mood and the Malleability of Moral Reasoning: The Impact of Irrelevant Factors on Judicial Decisions*, J. BEHAV. & EXPERIMENTAL ECON., June 2025, at 1, 10.

11. See Daniel L. Chen, *The Judicial Superego: Implicit Egoism, Internalized Racism, and Prejudice in Three Million Sentencing Decisions*, 77 KYKLOS 1004, 1016 (2024).

12. Arnaud Philippe & Daniel L. Chen, *Clash of Norms: Judicial Leniency on Defendant Birthdays*, 211 J. ECON. BEHAV. & ORG. 324, 335 (2023).

Using a study on gender attitudes in U.S. Circuit Courts as a case in point,¹³ I will illustrate how AI holds the potential to counter biases more efficiently than humans. It can do so by diagnosing bias in a manner that humans cannot.¹⁴

The study in question uses Natural Language Processing (NLP), a branch of AI, to detect judges' attitudes toward females.¹⁵ The researchers developed a "gender slant" measure to gauge how judges associate men with careers and women with families in their written opinions.¹⁶ This nuanced approach brought to light subtle gender biases not captured by direct ruling analysis.¹⁷

The unique strength of AI here is the ability to aggregate and analyze enormous amounts of data objectively, far beyond human capacity. NLP provided a quantitative, unambiguous measure of gender slant by examining 380,000 published opinions,¹⁸ something that would be impractical if not impossible for humans to achieve in a comparable timeframe. AI, therefore, holds an exceptional potential for bias detection that goes beyond human capabilities.

Gender slant is a convincing proxy for bias.¹⁹ Female and younger judges display lower gender slant.²⁰ Having a daughter reduces gender slant.²¹ Lower gender slant is also associated with more frequent use of gender-neutral pronoun constructions, such as "he or she" or "they."²² Finally, judges with higher slant tend to express less empathy toward women in their writing.²³

Judges not only systematically differed in the way they write about gender; these differences are also predictive of how they decide gender rights cases and how they treat their female colleagues.²⁴ The study examines how judges with different gender slant levels interact with female judges in three areas: reversals of lower court decisions, opinion assignments, and citations.²⁵ Results

13. Elliott Ash, Daniel L. Chen & Arianna Ornaghi, *Gender Attitudes in the Judiciary: Evidence from U.S. Circuit Courts*, 16 AM. ECON. J.: APPL. ECON. 314 (2024).

14. *See generally id.* at 316 (describing the "growing literature using word embeddings to analyze bias in text").

15. *Id.* at 315.

16. *Id.*

17. *Id.*

18. Ash et al., *supra* note 13, at 323.

19. *Id.* at 334.

20. *Id.* at 316.

21. *Id.*

22. *Id.*

23. *Id.*

24. Ash et al., *supra* note 13, at 317.

25. *Id.* at 338, 341, 344.

show that judges with higher gender slant are more likely to reverse decisions by female district judges, less likely to assign opinion authorship to female judges, and less likely to cite female judges' opinions.²⁶ These judges also tend to vote more conservatively in gender-related cases.²⁷ Findings suggest that gender bias could hinder female judges' career progression and reinforce gender disparity in the judiciary.²⁸

The underrepresentation of women at the top of the legal profession is an issue that has received considerable attention in the United States. It's troubling that although nearly 45% of law school graduates since the 1990s have been women, females still account for only 20% of equity partners in large law firms and 30% of state and federal judgeships.²⁹ The disparities in these numbers speak to a systemic problem: the differential treatment of female judges, possibly due to gender attitudes among their colleagues.

Gender attitudes, or the biases and preconceptions one holds about social groups, notably women and racial minorities, are known to significantly influence judgments and choices. These biases affect decisions in a range of contexts, from physician treatments and hiring decisions to employer-employee interactions and even the effectiveness of teachers. If these attitudes imply differential treatment of female judges, they could be a contributing factor to the underrepresentation of women in the judiciary.

It's challenging to examine these issues among justice actors due to the lack of traditional measures of gender attitudes for judges. However, researchers have innovatively used recent developments in natural language processing (NLP) to propose a novel measure of gender attitudes.³⁰ By analyzing a large corpus of written text from appellate judges, researchers have developed a measure of gender bias based on how strongly judges associate men with careers and women with families in their writing.³¹ Using a technological tool called word embeddings, the researchers calculated a judge-specific gender bias measure.³²

26. *Id.* at 347.

27. *Id.*

28. *Id.* at 317.

29. *Women in the Legal Profession*, AM. BAR ASSOC. (2024), <https://www.americanbar.org/news/profile-legal-profession/women> [<https://perma.cc/Z79T-E9HN>].

30. Ash et al., *supra* note 13, at 315.

31. *Id.*

32. *Id.*

Unlike humans, where bias can become deeply entrenched over time, AI can be swiftly and effectively adjusted once a bias is detected. There's no need for lengthy educational or societal campaigns to alter AI behavior—a programmer can do it with a few lines of code.

Besides the text, there is voice. Voice, long considered a neutral channel of expression, is revealed by Chen, Halberstam, and Yu's studies³³ to be a surprisingly potent driver of biased decision-making within the U.S. Supreme Court bar. Their work, which delves into the influence of perceived masculinity and femininity in advocates' voices on case outcomes, raises critical questions about implicit bias, systemic inequality, and mutable characteristics that Cahn, Carbone, and Levit explore with equal vigor in *Fair Shake*.³⁴ In *Fair Shake*, Cahn, Carbone, and Levit argue for a reimagined vision of justice and equality—one that moves beyond traditional markers of equity and embraces deeper, more nuanced understandings of how individuals navigate the world through mutable and immutable aspects of their identities.

This essay reflects on the ways vocal traits serve as battlegrounds for broader social biases and how such subtle but pervasive discrimination often denies a “fair shake” to those who diverge from traditional norms. Chen, Halberstam, and Yu's studies provide an examination of how voice-based judgments influence Supreme Court outcomes. Their analysis of 1,901 oral arguments from 1998 to 2012 finds that male petitioners perceived as having less masculine voices are significantly more likely to succeed.³⁵ In contrast, women with more feminine-sounding voices fare better.³⁶ These patterns indicate that voice, a mutable characteristic, carries weighty consequences in high-stakes legal contexts—ones that ostensibly should be free from prejudicial influence.

One of the most poignant intersections between Chen, Halberstam, and Yu's studies and *Fair Shake* lies in their shared exploration of “covering”—the demand that individuals downplay

33. Daniel Chen, Yosh Halberstam & Alan C. L. Yu, *Perceived Masculinity Predicts U.S. Supreme Court Outcomes*, 11 PLOS ONE (2016) [hereinafter Chen et al., *Perceived Masculinity*], <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0164324> [https://perma.cc/AYX9-VZ7V]; Daniel Chen, Yosh Halberstam & Alan Yu, *Covering: Mutable Characteristics and Perceptions of Voice in the U.S. Supreme Court*, J.L. & EMPIRICAL ANALYSIS (forthcoming 2025) [hereinafter Chen et al., *Covering*] (on file with authors).

34. See Chen et al., *Covering*, *supra* note 33 (manuscript at 20).

35. Chen et al., *Perceived Masculinity*, *supra* note 33, at 7.

36. Chen et al., *Covering*, *supra* note 33 (manuscript at 20).

aspects of their identity to gain social acceptance.³⁷ Chen and colleagues note that lawyers may adopt vocal strategies to cover their natural speaking patterns to conform to expectations. Cahn, Carbone, and Levit emphasize that covering extends far beyond the legal arena, shaping how individuals navigate every facet of their lives. The authors recount examples of professionals who consciously alter their speech, behavior, or appearance to fit within rigid social norms, often at great personal cost.

The concept of covering underscores a tension at the heart of both works: assimilation often serves as a prerequisite for success in systems purportedly based on merit. Yet, as Chen, Halberstam, and Yu show, lawyers who don't conform to the masculine norm can be punished for being different: individuals misperceive more masculine-sounding lawyers as winners or have a taste for being around masculine-sounding lawyers. This pressure to conform can become dehumanizing, stripping individuals of their authenticity and perpetuating a cycle of marginalization.

Workplace practices that prioritize "cultural fit," an often-coded term used to enforce homogeneity and limit true diversity, can limit the opportunities for a fair shake. Just as law firms and judicial actors may inadvertently (or purposefully) penalize lawyers with voices that deviate from perceived norms, so too do workplaces often judge individuals harshly for failing to assimilate fully into dominant workplace cultures. By drawing attention to the stakes involved—whether in the form of a legal ruling or job security—both works highlight how mutable characteristics become barriers to equitable outcomes.

The correlation between voice-based perceptions and Supreme Court outcomes mirrors the discriminatory practices Cahn, Carbone, and Levit critique in *Fair Shake*. This subtle form of prejudice serves as a reminder that discrimination has evolved into more insidious and less overt forms. In both cases, individuals are judged based on attributes that deviate from the perceived "norm." In the courtroom, as on other societal stages, such judgments create a new class of outsiders and insiders—those who fit and those who do not.

A key reflection emerging from this integration is that bias does not operate solely at the level of conscious prejudice. It can manifest as statistical discrimination—where judgments are made

37. See *id.* (manuscript at 2) ("[W]hen dominant groups or courts enforce 'covering', or assimilation, on subordinate groups, it perpetuates a form of second-class citizenship."); CAHN ET AL., *supra* note 5, at 14.

based on assumed characteristics or misperceptions. Chen and colleagues' experiment using incentives and information illustrates that voice-based perceptions can be mitigated when individuals are made more aware of their biases or incentivized to make accurate judgments. Cahn, Carbone, and Levit similarly call for educational and structural reforms to promote empathy, challenge stereotypes, and create spaces where marginalized voices are heard on their own terms.

Chen, Halberstam, and Yu's experimental efforts to debias perceptions through information and incentives offer a glimmer of hope. By making biases explicit and offering incentives to counteract prejudice, they demonstrate that it is possible to attenuate the influence of voice-based judgments. This finding resonates with Cahn, Carbone, and Levit's call for systemic change through awareness-building, education, and policy reforms aimed at creating more inclusive institutions.

However, both works caution against over-reliance on individual efforts to address structural issues. As Cahn, Carbone, and Levit emphasize, achieving a "fair shake" requires not just debiasing individuals but transforming the structures that enable discrimination to persist. The studies by Chen and colleagues provide a starting point, revealing how mutable characteristics shape outcomes in high-stakes environments. But they also highlight the limits of individual adaptation in the face of entrenched systemic biases.

The intersection of Chen, Halberstam, and Yu's studies and *Fair Shake* reveals a complex tapestry of identity, perception, and justice. Mutable characteristics like voice may seem superficial, but they carry profound implications for how individuals are perceived and treated in society's most powerful institutions. For lawyers arguing before the Supreme Court, as for professionals in every field, success often hinges on the ability to navigate a web of biases that privilege conformity over authenticity.

The call to action from both works is clear: true justice requires dismantling the systems that perpetuate inequality, challenging the biases that shape perception, and creating environments where all individuals—regardless of their voices, gender, or other mutable traits—can receive a fair shake. By integrating empirical evidence with personal narratives, Chen, Halberstam, and Yu's studies and Cahn, Carbone, and Levit's book invite us to reflect deeply on the kind of society we wish to build—one where equity is not contingent on conformity but is grounded in respect for each individual's unique voice.

This vision of equity is particularly relevant in the context of emerging technologies like AI, which at once mirror, magnify, and mute human intent. The dichotomy of AI and bias highlights a crucial truth: AI is neither inherently good nor bad, but a tool whose impact depends on how we choose to wield it. While it has the potential to amplify existing prejudices, it also offers an unprecedented opportunity to reduce biases and foster greater fairness. The challenge lies in ensuring that AI aligns with the equitable society we aspire to create.

For instance, as we advance from the analysis phase to the application phase in the study on gender attitudes in U.S. circuit courts, AI could be utilized to counter the detected gender biases. Systems can be programmed to prompt human judges to self-reflect and reconsider potential implicit biases.

Ultimately, the goal of leveraging AI isn't to replace human judgment but to enhance it, to make us more aware of our inherent biases, and aid us in countering them. Just as a spellchecker alerts us to a misspelled word, AI could alert us to potential bias, pushing us toward a more fair and just society.

AI is only as good or as bad as we allow it to be. As researchers, programmers, and users, we must remain vigilant to the biases we could be unknowingly encoding into our AI systems. Only then can we truly unlock the potential of AI to reduce bias and discrimination, making strides toward a more equitable world and a more inclusive and unbiased environment.