AI in the Racial Wealth Gap

the CENTER for Civil Rights and Technology

Executive Summary

The racial wealth gap in the United States is one of the most persistent measures of economic inequality. Without safegaurds, AI will widen this gap – not narrow it.

As of 2022, the median white household held approximately $285,000 in net wealth compared to just $45,000 for the median Black household—a staggering gap of over $240,000 that has continued to widen considerably since the early 1960s. Rooted in centuries of systemic exclusion, the racial wealth gap has been perpetuated across generations through three critical domains: housing, lending, and employment, and has deprived communities of investments and infrastructure across education, transportation, safety, broadband, and many other areas. Today, as artificial intelligence (AI) tools become embedded in the systems that govern access to housing, lending, and employment, as well as education, health and other critical areas, a critical question emerges: will AI disrupt this pattern of inequality and narrow the racial wealth gap, or will it further systematize and scale the pattern, widening that gap? This report addresses this question through an examination of AI's documented impacts across housing, lending, and employment domains, the three domains through which people in the U.S. historically build wealth.

The evidence presented in this report demonstrates that without deliberate and outcome-oriented interventions, AI will likely amplify the racial wealth gap rather than narrow it.

A consistent pattern currently emerges across housing, lending, and employment: AI systems trained on data that has been shaped by systemic racial disparities reproduce those patterns and can do so at massive scale, expanding and accelerating these trends across millions of transactions, judgments, and real-world interactions. Importantly, this report does not contend that automated technologies are inherently harmful; rather, the research confirms that when AI systems are trained on data shaped by longstanding structural discrimination, their outputs carry a substantial risk of reproducing and further entrenching those same inequities.

To be clear, this does not have to be a foregone conclusion. AI technologies can produce positive outcomes when developed, tested, and adopted with principled and civil-rights-protective guardrails. **But without safeguards, governance, and corrective measures, AI risks reinforcing existing disparities under the guise of neutrality and efficiency.**

Companies developing or deploying AI must do so responsibly. To address these concerns, transparency and accountability in algorithmic decision-making must be established across these three domains that contribute to the racial wealth gap.

Housing

In housing, landlords and tenant-screening companies should be required to disclose both that AI is being used and what their algorithmic criteria are before applicants apply, provide written tenant-selection plans, and offer rejected applicants actionable explanations for denial. Fair-chance housing laws should be expanded to limit consideration of arrests and convictions: Evidence shows that criminal histories are poor predictors of tenant success. In mortgage lending, third-party audits of algorithmic underwriting and credit scoring systems should be mandated for racial bias, and lenders should be required to publish disaggregated outcome data by race, ethnicity, and other demographic categories. Borrowers must have strengthened rights to explanation and contestation of algorithmic decisions.

Lending

Data governance and the use of alternative credit metrics represent another critical intervention point. Many borrowers—especially those from marginalized communities—are excluded from traditional credit markets because they lack formal credit histories, despite having consistent records of financial responsibility. Not all alternative data sources carry the same promise or risks, however. Incorporating rent and utility payment history into traditional credit bureau files raises serious concerns: negative payment data can block at-risk renters from finding housing altogether because most landlords use credit reports in their tenant-screening processes, which can treat any record of a late rent payment as disqualifying. A more consumer-protective approach would be the use of bank account transaction (or "cashflow") data, which typically requires consumer opt-in, does not embed data permanently into credit bureau files, and has been identified by the Consumer Financial Protection Bureau (CFPB) as a promising pathway to expand credit access for populations with historically low scores. Inaccurate data used in tenant screening, including in eviction records and criminal background checks, must be audited and corrected; clear look-back periods must be established; and charges that have been dismissed, sealed, or expunged must be excluded from consideration.

Jobs

Algorithmic hiring and performance evaluation systems should be regulated to prevent discriminatory screening, and companies should be required to inform workers about how algorithms assess them not only at hiring, but also in post-hire use-cases that shape workplace conditions, promotions, discipline, and outcomes. Worker protections and labor rights must be strengthened to address AI's role in workplace exploitation. Enforcement against worker misclassification should be enhanced to ensure gig workers and contingent workers are not deprived of legal protections through algorithmic management. Just-cause and fair-scheduling standards should be implemented to limit arbitrary termination and exploitative scheduling practices enabled by algorithms.

Finally, community-centered governance must underpin all interventions. This means that people most affected by algorithmic systems, particularly communities of color, have a real say in how those systems are built, tested, and monitored—not merely the opportunity to comment after decisions have already been made. In practice, community members should help select auditors who review algorithmic systems, should sit on oversight boards with actual decision-making authority, and should be empowered to enforce agreements requiring companies to demonstrate fair outcomes before and during deployment.

This report is grounded in a mixed-methods approach that combines multiple research streams. The literature review comprises academic studies, policy reports, and regulatory documents examining AI's effect on the racial wealth gap. This includes empirical studies on algorithmic impacts in housing, lending, and employment, policy evaluations of regulatory frameworks governing AI in these domains, qualitative research documenting lived experiences of automated system use, and data from federal agencies including the Federal Reserve, Census Bureau and Federal Trade Commission (FTC). Beyond the literature, the report draws on qualitative research that centers the voices of those directly affected by AI-driven decision-making systems. Semi-structured interviews were conducted with renters, prospective homebuyers, and employees from various backgrounds. These qualitative data provide crucial context and insight into how algorithmic systems function in practice and suggest meaningful remedies and protections.

Written by:
Yvette N.A. Pappoe
Nadiyah J. Humber
Darlis Pantoja-Benavides


AI
in the
Racial
Wealth
Gap