

Recommendations to ensure AI systems help close the racial wealth gap.

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. 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 – housing, lending, and jobs – that contribute to the racial wealth gap.

Tenant-Screening Recommendations

- **Maintain internal systems that promptly identify, document, and correct inaccurate or outdated data and biased model behavior.** Developers and deployers should maintain clear, public channels for receiving error and bias reports from end users and affected communities, and incorporate validated corrections into data pipelines, model updates, and documentation to prevent recurring inaccuracies in future iterations. These channels should feed into systems that enable deployers to flag and correct suspected errors or patterns of bias and ensure that such updates have auditable trails. Developers and deployers should also publish regular public reports summarizing, in aggregate, the errors and bias patterns identified, corrections made, and outcomes, so that advocates, regulators, and affected communities can assess whether internal accountability mechanisms are working.
- **Ensure that tenant-screening tools are built and configured to predict rental performance specifically.** Developers and deployers should regularly validate that the variables used in screening models predict outcomes like on-time rent payment and lease completion, and remove those that do not. Metrics adapted from credit or mortgage contexts, where different outcomes are at stake, should be reassessed and replaced with indicators that bear directly on a prospective tenant's current ability to pay rent.

- **Design screening tools that do not rely on minimum credit score thresholds or rigid income-to-rent ratios as default qualifying criteria.** Because traditional credit and income data have uncertain predictive value for rental performance, these tools should be validated against actual rental outcomes and adjusted accordingly. Tools should also be capable of incorporating alternative indicators of current ability to pay, such as bank account transaction (cashflow) data, for applicants with limited or nontraditional credit profiles. The use of cashflow and bank transaction data, however, raises privacy concerns, including ensuring that applicants provide meaningful informed consent, that data collection is limited to what is necessary for rental assessment, and that such data is not retained, shared, or repurposed beyond the screening context.
- **Integrate debiasing techniques into tenant-screening systems.** Developers and deployers should integrate debiasing techniques into their screening tools and decision-making processes to reduce the influence of biased or low-quality data, including eviction records, criminal histories, and credit reports, and incorporate training techniques specifically designed to mitigate disparate impacts. Developers and deployers should conduct rigorous testing of existing and more inclusive screening models using civil rights-centered testing frameworks, in coordination with fair housing advocates and other subject-matter experts, to evaluate which approaches produce the most reliable, accurate, and equitable tenant recommendations. As noted above, these technical interventions are necessary but not sufficient alone, and should be understood as a floor, rather than a ceiling, for responsible practice. These debiasing techniques must be enacted in concert with the other safeguards and backstops outlined in this document.
- **Compile and maintain data in a manner that ensures reliability and facilitates timely correction.** Developers and deployers should refresh and validate databases on a regular, predictable schedule, so the data reflects people’s current circumstances, rather than relying on ad hoc or discretionary updates.
- **Provide internal systems that enable prompt identification and correction of inaccurate or outdated records.** End users should be given a clear and accessible way to flag and correct potential errors, and those reports should be sent through a clear process for review, correction, and confirmation so that inaccuracies are resolved and do not recur. In addition, developers and deployers should equip operational staff with appropriate tools or dashboards to review and update data.

- **Include explicit good-faith and fair-dealing expectations in deployment contracts.** Service agreements should confirm that housing providers remain responsible for housing decisions made with automated systems, even when those systems come from third-party vendors. Contracts should also state that deployers will maintain accurate and current data, keep systems properly updated, and offer clear processes to identify, challenge, and correct errors so that these technologies support fair and lawful decision-making.
- **Push for contractual transparency and actionable explainability in tenant screening.** Provisions requiring developers and deployers to provide meaningful transparency into data inputs, scoring factors, performance metrics, and potential disparate impacts, along with audit rights and notice of material model changes, should be included in their service agreements. Contracts should also require developers and deployers to supply clear, plain-language explanations and actionable guidance that landlords can share with applicants so individuals can understand adverse decisions and take concrete steps to improve future housing outcomes.

Homeownership and Lending Recommendations

- **Prioritize cashflow data over rent- and utility-payment reporting in alternative credit models.** Credit scoring systems seeking to reach borrowers excluded from traditional credit markets should favor bank account transaction (“cashflow”) data, which requires consumer opt-in and does not risk introducing negative payment information into permanent credit bureau files. If rent-reporting is considered, it must be strictly limited to positive-only, opt-in programs as full-file reporting risks causing more harm than good, particularly for the most at-risk renters, who may be denied housing based on any negative payment entry.
- **Fairness-aware model design for developers and deployers requires affirmative design, testing, and correction obligations.** Design, train, and validate systems using fairness-aware methods, including routine testing for disparate impacts across protected classes and evaluation against defined equity benchmarks and accuracy metrics. Where testing reveals unjustified disparities, developers and deployers should adjust data pipelines, model architectures, and post-processing techniques to mitigate discriminatory outcomes, and should document both corrective measures and residual risks so deployers and regulators can assess compliance and real-world impact.

- **Develop inclusive AI tools.** To address longstanding structural barriers, such as fewer physical banks in majority-Black counties and higher unbanked rates, developers and deployers should produce AI tools that expand access to banking, enable personalized financial planning, and improve compliance with fair lending laws. Fintech companies should create alternative credit metrics, offer mobile banks and mission-driven institutions, offer low-cost services, and provide wealth-building platforms.
- **Promote transparent, consumer-centered AI tools by providing clear disclosures about algorithmic decision-making in lending.** Strengthen safeguards against biased credit algorithms through mandatory pre-deployment testing and ongoing monitoring, and ensure consumers have meaningful rights to explanation and contestation, not merely financial education. As scholars have noted, framing financial “literacy” as a solution to systemic inequity can reinforce harmful narratives about individual blame while deflecting from the structural reforms and corporate accountability that actually drive change.
- **Design AVMs to mitigate neighborhood-level racial bias.** AVMs should be trained and validated to avoid systematically undervaluing homes in majority-Black and majority-Brown neighborhoods. This requires auditing comparable selection algorithms for racially correlated proxies, testing valuation accuracy disaggregated by neighborhood demographics, and correcting models where disparities are identified.
- **Large lenders and AVM providers should undergo regular third-party audits for racial bias and publish disaggregated outcome data.** These audits should test models against clear fairness metrics across protected classes, with results made publicly available in accessible formats to enable community and regulatory oversight.

Jobs Recommendations

- **Design screening tools to evaluate substantive qualifications rather than surface-level formatting or stylistic choices.** Train and test résumé and application models to ensure they do not disproportionately exclude qualified applicants based on immaterial formatting, linguistic variation, or presentation differences that correlate with race or class but not with job performance. Use counterfactual testing (e.g., same qualifications, different names or styles) to verify that candidates of color are not screened out at higher rates.

- **Conduct pre-deployment and ongoing bias audits across race, ethnicity, and other protected characteristics and publish standardized impact documentation.** Provide employers and regulators with clear, accessible “model cards” or impact statements that explain data sources, known limitations, and documented disparate-impact test results, enabling informed decisions about deployment and oversight.
- **Build contestability and human-in-the-loop controls directly into system architecture by default.** Make it easy for employers to override automated outcomes, fix obvious mistakes, and send close or high-stakes cases (like rejections, deactivations, or pay cuts) to a human for review. Give employers built-in tools to record why a decision was made and to handle appeals so that workers can meaningfully challenge decisions that appear biased.
- **Limit surveillance and wage-personalization features that enable exploitative labor practices.** Avoid designing bossware defaults that facilitate pervasive, minute-by-minute tracking or dynamic wage adjustments that workers cannot predict or verify. Where monitoring is necessary for legitimate purposes such as safety, collect only the minimum data required and use high-level summaries instead of detailed, person-by-person tracking that can be used to control or punish workers, especially workers of color.