

Artificial Intelligence and the Racial Gap in Mortgage Lending

Yongqiang Chu ^{*} Jing He[†] Leiju Qiu[‡] Daxuan Zhao[§]

Abstract

We examine whether banks' adoption of artificial intelligence (AI) widens racial disparities in mortgage lending. Using 2022–2024 HMDA data merged with a bank AI adoption measure derived from job postings, we find that a one-standard-deviation increase in AI adoption raises the denial-rate gap between minority and non-minority applicants by 0.75 percentage points and the interest-rate gap by 4 basis points. Instrumental-variable estimates exploiting banks' historical exposure to the supply of AI graduates support a causal interpretation. The effects concentrate in purchase loans and are absent for GSE-backed and refinance loans. In contrast, AI adoption does not amplify disparities among FinTech lenders, where algorithmic underwriting has been embedded from inception.

Keywords: Artificial Intelligence; Algorithmic Lending; Racial Gap; Mortgage Banking; Credit Allocation

JEL Code: G21; J15; D81

^{*}Department of Finance, Belk College of Business, University of North Carolina at Charlotte. 9201 University City Blvd, Charlotte, NC 28223. Email: yongqiang.chu@uncc.edu.

[†]Department of Finance, School of Business, Renmin University of China. 59 Zhongguancun Street, Haidian District, Beijing, China 100872. Email: hejing0116@ruc.edu.cn.

[‡]China Center for Internet Economy Research, Central University of Finance and Economics, 39 South College Road, Beijing 100081, China. Email: leijuqiu@cufe.edu.cn.

[§]Department of Finance, School of Business, Renmin University of China. 59 Zhongguancun Street, Haidian District, Beijing, China 100872. Email: zhaodaxuan@ruc.edu.cn.

1 Introduction

Mortgage credit is the largest liability for most U.S. households and the primary vehicle for wealth accumulation through homeownership. Persistent racial disparities in mortgage approval and pricing have first-order implications for wealth inequality: gaps in access to credit and loan terms translate directly into gaps in housing equity and intergenerational wealth transfer (Munnell et al., 1996; Ambrose, Conklin, and Lopez, 2021; Frame et al., 2025). At the same time, mortgage origination is undergoing a technological transformation. Lenders increasingly deploy automated underwriting systems, digital application channels, and algorithmic screening to reduce costs and scale originations (Fuster et al., 2019), and the rapid diffusion of artificial intelligence (AI) and machine learning (ML) promises further gains in predictive accuracy and operational efficiency. Whether these technologies narrow or widen racial disparities in credit access and pricing is both economically consequential and empirically open, with direct bearing on bank governance, fair-lending enforcement, and the design of financial regulation.

Theory yields sharply opposing predictions. On one hand, AI and machine learning (ML) may *reduce* racial disparities by displacing human discretion with objective, data-driven risk assessment, correcting systematic misclassification or bias in traditional underwriting (Berg et al., 2020; Kleinberg et al., 2018; Jagtiani and Lemieux, 2018; Howell et al., 2024). On the other hand, Fuster et al. (2022) identify two channels through which ML may *amplify* disparities. The *flexibility effect* arises because ML models exploit high-dimensional, nonlinear interactions among observables, generating finer individual-level risk estimates and greater dispersion in credit outcomes that can widen group-level gaps when the dispersion differen-

tially affects minority borrowers. The *triangulation effect* arises because richer feature spaces let models proxy for excluded characteristics, including race, through correlated observables. Both channels operate even when lenders comply with fair-lending rules by excluding race from model inputs.

Existing quantitative and experimental evidence tilts toward the amplification side. Calibrating their model to U.S. mortgage markets, [Fuster et al. \(2022\)](#) find that Black and Hispanic borrowers benefit less from ML adoption than White and Asian borrowers, with the former reclassified as lower risk at roughly half the rate of the latter. [Bowen et al. \(2025\)](#) provide experimental evidence that large language models reproduce racial biases in underwriting decisions. Together, these studies motivate a specific, testable hypothesis: banks' AI adoption, on net, *widens* racial disparities in lending outcomes. Whether this hypothesis holds in practice remains an open question, because existing evidence relies on calibration exercises or controlled experiments rather than on banks' actual AI adoption decisions and the outcomes they generate for millions of real applicants.

We test these predictions using nationwide mortgage application data from the Home Mortgage Disclosure Act (HMDA) for 2022–2024. We merge these applications with a bank-year measure of AI adoption derived from AI-skill intensity in job postings ([Babina et al., 2024](#)). This proxy captures banks' organizational investment in AI capabilities and varies both across banks and within banks over time. Our empirical design exploits within-bank variation while controlling for rich borrower and loan characteristics and a demanding set of fixed effects, including bank $\times$ year, MSA, and census-tract $\times$ year. The bank $\times$ year fixed effects absorb all time-varying bank-level confounders, including changes in bank strategy, risk appetite, and lending standards, so that identification comes from the differential treat-

ment of minority versus non-minority applicants within the same bank-year. The census-tract $\times$ year fixed effects control for local demand-side shocks and neighborhood compositional changes.

We find that higher AI adoption is associated with a statistically and economically significant widening of racial disparities in both underwriting and pricing. A one-standard-deviation increase in a bank’s AI job-posting share (2.02 percentage points) is associated with a 0.75-percentage-point increase in the denial-rate gap between minority and non-minority applicants, a meaningful effect relative to the unconditional minority denial-rate gap. Among approved loans, the same increase in AI adoption is associated with roughly a 4-basis-point increase in the interest-rate gap. These patterns persist after controlling for detailed observables and location-by-time shocks and are robust to alternative proxies for AI adoption.

To address endogeneity, we follow [Babina et al. \(2024\)](#) and instrument bank AI adoption with banks’ *ex ante* exposure to the supply of AI graduates from universities. Using Revelio Lab data, we identify all employees at sample banks in 2008–2010 and their alma maters, then apply [Babina et al. \(2024\)](#)’s method to classify universities as AI strong. Our instrument, *AI Strong*, equals one if at least one employee in 2008–2010 graduated from an AI strong university, and zero otherwise.

This instrument leverages banks’ historical ties to universities with subsequent AI prominence. Conditional on our rich fixed effects, this supply-driven variation is highly unlikely to correlate with contemporaneous shocks to minority mortgage outcomes, thereby plausibly satisfying the exclusion restriction. A placebo test using pre-2008 mortgage data confirms that the instrument does not predict racial gaps in denial before the relevant AI adoption

period, further supporting its validity. The IV estimates are more than double the OLS magnitudes, consistent with downward bias in OLS, plausibly because banks that adopt AI are simultaneously pursuing other strategies to mitigate racial disparities, attenuating the observed relationship.

A battery of tests rules out alternative explanations. AI adoption does not significantly affect the share of minority loan officers or the racial composition of applicants, ruling out compositional channels. Minority applicants at high-AI banks are, if anything, observably *less* risky along key dimensions (combined loan-to-value and loan-to-income), so the disparity cannot be attributed to risk-based sorting. The effects concentrate in underwriting decisions rather than at the pre-application stage (withdrawals and incomplete applications), where AI involvement is minimal. Robustness tests using McDash data, which contain credit scores and precise loan-to-value ratios not available in HMDA, confirm that our findings survive the inclusion of these critical underwriting variables.

The cross-sectional pattern of effects further illuminates the mechanism. The widening of the denial gap is concentrated in loans held on balance sheet and loans sold to non-GSE investors and is absent for loans sold to GSEs, consistent with GSEs’ standardized underwriting guidelines mitigating AI-driven amplification. The effect is also concentrated in purchase loans, where borrower risk is less directly observable and screening intensity is highest, and is absent for refinance loans, where existing payment histories limit the flexibility and triangulation channels emphasized by [Fuster et al. \(2022\)](#). These patterns map closely onto the theoretical mechanism: AI amplifies disparities precisely where its feature space has the most room to operate.

As a contrast case, we examine FinTech lenders, where fully automated underwriting

with limited loan-officer discretion has been embedded from inception (Buchak et al., 2018; Fuster et al., 2019). AI adoption does not amplify racial disparities in this segment. The divergence indicates that AI technology per se does not widen racial gaps; rather, its equity implications depend on the institutional environment in which it is deployed. At traditional banks, AI interacts with legacy processes, historical data, and residual human discretion in ways that amplify historical disparities, whereas technology-native FinTech operations appear to insulate underwriting from such amplification.

This paper contributes to three strands of the literature. First, and most directly, we provide the first large-sample empirical evidence that banks’ actual AI adoption amplifies racial disparities in mortgage lending. Foundational work documents systematic racial disparities in mortgage approval and pricing (Munnell et al., 1996; Hunter, 1996; Yinger, 1995), and more recent studies relate persistent gaps to intermediary characteristics and organizational practices (Ambrose, Conklin, and Lopez, 2021; Frame et al., 2025; Chu, 2025). Evidence from FinTech settings shows that algorithmic lending can coexist with disparate outcomes (Bartlett et al., 2022). Fuster et al. (2022) provide a theoretical framework predicting that ML widens racial disparities through flexibility and triangulation channels, and Bowen et al. (2025) offer experimental evidence that LLMs reproduce racial biases. Our contribution is to move from theory and simulation to field evidence: using data on millions of actual mortgage applications and a measure of banks’ revealed investment in AI, we show that the disparity-widening predictions of Fuster et al. (2022) materialize in practice. The causal evidence from our IV estimates, the concentration of effects in purchase loans and their absence in GSE-backed and refinance loans, and the sharp contrast with FinTech lenders collectively establish that the institutional context of AI deployment is a first-order determinant of its

equity implications.

Second, we contribute to research on the role of AI and technology in financial markets. Prior work examines how automation reshapes origination processes and borrower–lender matching (Fuster et al., 2019; Buchak et al., 2018), but has largely focused on efficiency and aggregate access to credit. We show that AI adoption has meaningful *distributional* consequences that differ across borrower groups and across institutional contexts. The divergence between traditional banks and FinTech lenders is particularly informative: it suggests that the equity implications of AI depend not on the technology per se but on the organizational environment (legacy processes, human discretion, and historical data) in which it is embedded. This finding has direct implications for how regulators and bank managers should approach AI governance.

Third, our analysis connects to the growing literature on algorithmic fairness and the distributional effects of predictive technologies (Dwork et al., 2012; Hardt, Price, and Srebro, 2016; Kleinberg et al., 2018). That literature emphasizes that improvements in predictive accuracy need not translate into equitable outcomes when protected characteristics are correlated with model inputs through historical and structural channels. We provide field evidence from a high-stakes financial market consistent with this concern: AI-assisted underwriting changes the mapping from applicant information to credit outcomes in ways that systematically disadvantage minority borrowers, even within narrowly defined geographic and time cells. These findings underscore the importance of model-risk management, disparate-impact testing, and transparency in AI deployment, and highlight the need to evaluate not only average predictive performance but also distributional impacts across demographic groups.

The remainder of the paper proceeds as follows. Section 2 describes the data and the

construction of the bank-level AI adoption measure and mortgage outcomes. Section 3 presents the baseline results, instrumental-variables analysis, and heterogeneity tests. Section 4 rules out several alternative explanations. Cross-sectional analyses, robustness checks, and additional tests are reported in Sections 5 and 6. Section 7 concludes.

2 Data and Sample

2.1 Bank AI Adoption

Following the human-capital-based approach in [Babina et al. \(2024\)](#), we construct a bank-level measure of artificial intelligence (AI) adoption using job postings data. We obtain the job posting data from Revelio Labs, a global workforce intelligence provider that aggregates and harmonizes job postings and employee records from a wide range of public and proprietary sources, including online job boards, company career pages, and professional networking platforms. Revelio Labs standardizes employer identities, job titles, and skill taxonomies, enabling consistent measurement of skill demand across firms and over time. Such job-postings data are widely used in research to measure skill demand, technology diffusion, and firm behavior ([Babina et al., 2024, 2023, 2025](#); [Babina and Fedyk, 2025](#); [Deming and Kahn, 2018](#); [Alekseeva et al., 2021](#); [Azar et al., 2018](#); [Goldfarb, Taska, and Teodoridis, 2023](#); [Blair and Deming, 2020](#)).

The Revelio job-posting data are from two sources, LinkedIn and Indeed. The coverage of Indeed begins in 2016 and the coverage for LinkedIn begins in 2021. LinkedIn has better coverage of higher-skilled white-collar jobs, which is critical for measuring AI-related job

postings. We therefore use the Revelio job postings from 2021–2023 to construct banks’ AI adoption.

Specifically, for each bank-year, we compute the share of job postings that require AI-related skills relative to the bank’s total number of job postings in that year. This share serves as our primary proxy for bank AI adoption. It captures the variation in banks’ strategic emphasis on AI while mitigating concerns related to bank size or overall hiring activity. This AI adoption measure exhibits substantial cross-sectional heterogeneity. A large fraction of banks report zero AI-related job postings in a given year, while banks in the upper tail display meaningful and persistent AI hiring intensity. This distributional pattern is consistent with uneven and gradual adoption of AI technologies across banks, and mirrors the stylized facts documented for non-financial firms in [Babina et al. \(2024\)](#).

The Revelio job-posting data do not directly identify whether the AI-related postings are for mortgage-related business. Because we measure firm-wide AI investment rather than mortgage-specific AI use, our coefficients capture the bank-level treatment effect of AI adoption on mortgage outcomes. To the extent that firm-wide AI investments are concentrated outside the mortgage business, this would attenuate our estimates relative to the effect of mortgage-area AI per se. At the same time, even AI investments outside of mortgage origination (for example, in enterprise data analytics, risk management, or fraud detection) can still affect mortgage outcomes by changing the underlying risk-assessment and decision-making infrastructure on which mortgage underwriting relies.

2.2 Mortgage Lending

We obtain mortgage application level data from the Home Mortgage Disclosure Act (HMDA) database covering the period 2022–2024 because our bank AI adoption measure is available for 2021–2023. Under HMDA reporting requirements, regulated financial institutions with assets exceeding \$30 million, including commercial banks, credit unions, and mortgage companies, are required to disclose detailed information on their mortgage lending activities. The HMDA data include the lender’s identity, the geographic location of the property, loan amount, application year, and application outcomes, indicating whether a loan is approved or denied. For approved applications, the dataset further reports pricing information, such as the contractual interest rate and the rate spread, the difference between the annual percentage rate (APR) and the average prime offering rate. In addition, HMDA provides detailed borrower characteristics, including reported income, race, ethnicity, and gender.

We restrict the sample to mortgage applications with complete information on borrower characteristics, loan attributes, and lender identifiers. The denial analysis is conducted on the full set of applications with non-missing denial outcomes. Analyses of interest rates and rate spreads are based on subsamples of approved loans with available pricing information, resulting in smaller sample sizes due to missing rate disclosures.

For our baseline testing sample, we focus on non-FinTech lenders because FinTech lenders have embedded algorithmic underwriting from inception, and their AI adoption may not reflect the same dynamics as traditional banks. We identify FinTech lenders using the list constructed by [Buchak et al. \(2018\)](#). The main sample thus consists of applications to

traditional banks, while we conduct separate analyses for the FinTech subsample to examine heterogeneity.

Using the HMDA data, we construct a comprehensive set of loan-level outcome variables, borrower characteristics, and loan attributes. Mortgage lending outcomes include an indicator for loan denial, the contractual interest rate for approved loans, and the rate spread, defined as the difference between a loan’s annual percentage rate (APR) and the average prime offering rate on similar mortgages. Borrower demographic characteristics, including race, ethnicity, gender, and reported income, are obtained directly from HMDA filings. Loan attributes include the loan amount, loan maturity, combined loan-to-value (CLTV) ratio, the presence of a co-applicant, and debt-to-income information.

2.3 Sample Construction

We merge the mortgage lending data with the bank-level AI adoption measure. The Revelio Lab data provide the link between the Revelio company identifier (RCID) and the LEI, which is the HMDA lender identifier. We then merge the bank-year AI adoption measure with the HMDA data using the LEI and year identifiers. We restrict the sample to applications with non-missing values for key variables, including denial outcomes, interest rates, and borrower characteristics. We also exclude applications from lenders that cannot be matched to the Revelio Lab data, resulting in a final sample of 13,915,800 loan applications for the denial regressions and 11,086,894 approved loans for the interest rate analyses.¹

¹HMDA data are publicly available from the Consumer Financial Protection Bureau (<https://ffiec.cfpb.gov/data-browser/>). Revelio Labs and ICE McDash data are proprietary and can be obtained from the respective vendors under standard licensing terms. Code to construct all variables and reproduce the tables in the paper is available from the authors upon request and will be posted in a public replication archive upon publication.

Table 1 reports the summary statistics for the main variables used in the analysis. The sample includes 13.9 million mortgage applications spanning 2022–2024. The mean denial rate is 0.201, indicating that approximately 20.1% of applications are rejected, with a standard deviation of 0.400. Among the 11.1 million approved loans with available pricing information, the average contractual interest rate is 6.16 percent (standard deviation: 2.02%), ranging from 2.99% at the fifth percentile to 9.50% at the ninety-fifth percentile. The average rate spread, defined as the difference between the loan’s annual percentage rate and the average prime offering rate, is 0.465 percentage points (standard deviation: 1.374 percentage points), with substantial heterogeneity across borrowers (fifth percentile: -1.325% ; ninety-fifth percentile: 2.750%).

Demographic characteristics reveal meaningful sample composition. Minority applicants (African American or Hispanic) account for 26.0% of the sample, and 37.0% of applicants are female. The borrower income distribution shows a mean log income of 4.627 with a standard deviation of 0.680, translating to a median income of approximately \$99,000. Loan-to-income ratios average 2.236 (standard deviation: 1.562), indicating substantial variability in leverage across borrowers. Co-applicants are present in 41.3% of applications. The combined loan-to-value ratio averages around 0.73 (73.27%), reflecting typical down payment and equity positions.

The bank AI adoption measure exhibits considerable heterogeneity. The mean AI job-posting share is 1.14% (standard deviation: 2.02%), with a median of zero, reflecting that the majority of banks report minimal or no AI-related hiring in a given year. However, the distribution is highly right-skewed, with banks in the upper tail showing substantially higher AI intensity: the ninety-fifth percentile reaches 6.03%. This skewed distribution accurately

captures the reality of uneven AI adoption across the banking sector, with most traditional lenders having limited AI investment while a subset of larger, technology-focused institutions report meaningful AI hiring intensity. The presence of substantial cross-sectional and time-series variation in both AI adoption and lending outcomes provides rich identifying variation for our empirical design.

INSERT TABLE 1 HERE

3 Bank AI Adoption and the Racial Gap in Mortgage Lending

This section presents our baseline results on bank AI adoption and racial disparities in mortgage lending and then reports instrumental-variables results.

3.1 Research Design

We estimate the following specification:

$$Y_{i,b,m,c,t} = \beta_1 (Minority_i \times Bank AI_{b,t-1}) + \beta_2 Minority_i + \Gamma' X_i + \delta_{b \times t} + \delta_m + \delta_{c \times t} + \boldsymbol{\delta}' \mathbf{Z}_i + \varepsilon_{i,b,m,c,t}, \quad (1)$$

where $Y_{i,b,m,c,t}$ denotes lending outcomes for application i at bank b in year t for a borrower in MSA m and census tract c . We study two outcomes: an indicator for denial (*Denial*) and the contractual interest rate for approved loans (*Interest Rate*).

The independent variable of interest is $Minority_i \times Bank AI_{b,t-1}$, where $Minority_i$ equals

one for African American or Hispanic applicants and zero otherwise, and $Bank\ AI_{b,t-1}$ is the bank’s AI job-posting share in $t - 1$. The vector X_i captures application and borrower characteristics, including co-applicant, loan-to-income, log income, loan term, CLTV, and a female indicator. $\mathbf{Z}_i$ collects fixed effects for debt-to-income bins, loan type, and applicant age. In particular, the debt-to-income ratios are sorted into less than 20%, 20%–29%, 30%–35%, 36%–40%, 41%–45%, 46%–49%, 50%–60%, and greater than 60%.²

All regressions include bank $\times$ year fixed effects $\delta_{b\times t}$, which absorb all time-varying bank-level confounders, including changes in bank strategy, risk appetite, and overall lending standards, so that identification comes from the differential treatment of minority versus non-minority applicants within the same bank-year. MSA fixed effects δ_m and census-tract $\times$ year fixed effects $\delta_{c\times t}$ control for demand-side factors, including local economic conditions, neighborhood composition, and house price dynamics, that might be correlated with both minority applicant characteristics and lending outcomes. Standard errors are clustered at the bank level to account for within-bank correlation in AI adoption and lending decisions.

3.2 Application Denial

We present the results for application denial in Table 2. The coefficient estimate on $Bank\ AI\times Minority$ is positive and statistically significant at the 1% level, suggesting that AI adoption increases the denial rates of minority applications relative to non-minority applications. A one standard deviation increase in $AI\ Job\ Share$ (2.02%) is associated with about 0.75 percentage-point increase in the denial gap between minority and non-minority applicants.

²The HMDA data report actual debt-to-income ratios when the ratio is between 36% and 49%. Outside that range, we use the same bins as reported by HMDA. The results remain robust if we use alternative bins.

The coefficient estimate on *Minority* remains positive and statistically significant, consistent with the vast literature documenting persistent racial gaps in mortgage approval (Munnell et al., 1996; Ambrose, Conklin, and Lopez, 2021; Bartlett et al., 2022; Lewis-Faupel and Tenev, 2024; Federal Reserve Bank of Minneapolis, 2024; Chu, Huang, and Zhang, 2025). In column (1), the coefficient estimate on *Bank AI* is negative and statistically significant, suggesting that higher AI adoption is associated with lower denial rates for non-minority applicants. This pattern is consistent with the theoretical predictions in Fuster et al. (2022) that AI adoption can lead to greater dispersion in credit outcomes, benefiting some borrowers while disadvantaging others.

The coefficient estimates on most controls align with standard underwriting patterns. Applications with co-applicants are less likely to be denied and, conditional on approval, tend to obtain more favorable pricing (Fetter et al., 2023; Goodman and Zhu, 2023). The negative coefficient estimates on *Loan to Income* are driven by the inclusion of the debt-to-income bucket fixed effects. By contrast, higher loan-to-income is associated with worse outcomes. The results are similar when we proxy AI adoption using the number (or log number) of AI-related job postings.³

INSERT TABLE 2 HERE

3.3 Mortgage Interest Rate

In Table 3, we present the results for mortgage interest rate. The coefficient estimate on *Bank AI* $\times$ *Minority* is again positive and statistically significant. A one-standard-deviation increase in *AI Job Share* widens the racial gap in mortgage interest rates by almost 4 basis

³These results are available upon request.

points. Although 4 basis points appears small, the effect is economically meaningful: on a 30-year, \$300,000 mortgage, an additional 4 basis points translates into roughly \$2,500 in extra interest payments over the life of the loan. Scaled by the unconditional minority share of approved loans, a back-of-the-envelope calculation implies an aggregate transfer on the order of several hundred million dollars per year from minority borrowers at high-AI banks.

The coefficient on *Minority* itself is small and statistically insignificant in our 2022–2024 sample. This contrasts with the 2009–2015 evidence in [Bartlett et al. \(2022\)](#), which documents an unconditional minority pricing premium of roughly 8 basis points. The difference is consistent with secular compression of unconditional racial pricing gaps following the post-2015 expansion of CFPB pricing-disclosure rules and the diffusion of automated pricing engines across traditional lenders. Our identification does not rely on the level of the minority dummy; the parameter of interest is the interaction with bank AI adoption, which captures how AI deployment differentially affects minority pricing within bank-year and tract-year cells.

The coefficient estimates on most control variables are consistent with the literature. Mortgage interest rates increase with the combined loan-to-value ratio and decrease with the presence of co-applicants.

INSERT TABLE 3 HERE

3.4 Instrumental Variable Estimates

Although our baseline specification includes demanding fixed effects, endogeneity remains a concern. Bank×year fixed effects absorb all time-varying bank-level confounders, and

census-tract $\times$ year fixed effects absorb local demand-side shocks, so identification comes from the differential treatment of minority versus non-minority applicants within the same bank-year and location-year cell. Nevertheless, bank AI adoption could still be correlated with unobserved shifts in how banks differentially treat minority applicants. The direction of potential bias is ambiguous. If banks adopt AI partly to mitigate human-driven racial bias, pursuing AI as one element of a broader fairness strategy, then the OLS estimates would understate the true disparity-widening effect of AI. Conversely, if banks adopt AI to better capture the credit risk of minority borrowers who exhibit higher *ex post* default rates, the OLS estimates could overstate the causal effect.

To mitigate this concern, we follow Babina et al. (2024) and instrument bank AI adoption with banks’ *ex ante* exposure to the supply of AI graduates from universities. Using Revelio Lab data, we identify all employees at sample banks in 2008–2010 and their alma maters, then apply Babina et al. (2024)’s method to classify universities as AI strong. Our instrument, *AI Strong*, equals one if at least one employee in 2008–2010 graduated from an AI strong university, and zero otherwise.⁴ This instrument leverages banks’ historical ties to universities with subsequent AI prominence. Conditional on our rich fixed effects, this supply-driven variation is highly unlikely to correlate with contemporaneous shocks to minority mortgage outcomes, thereby plausibly satisfying the exclusion restriction.

To further support the instrument’s validity, we examine whether *AI Strong* predicts the racial gap in mortgage lending before 2008, when AI adoption had not yet begun. A pre-period correlation would indicate that the instrument captures other persistent bank

⁴We do not use the share of employees from AI strong universities as the instrument because the dummy measure is empirically stronger, both in the coefficient estimates and first-stage F-statistics, suggesting that it is a more powerful predictor of AI adoption.

characteristics that affect the racial gap, such as exposure to other technologies or systematic differences in executive composition, thereby violating the exclusion restriction.

To test this, we regress mortgage denial decisions on the interaction between the minority dummy and the instrument, using mortgage applications from 2005–2007. We cannot perform the same test for mortgage interest rates because the interest rate data are available only starting from 2018. We use a three-year sample because our baseline results are also based on a three-year sample. We first present the results for each year separately in columns (1)–(3), and for the full sample in column (4) of Table 4. The coefficient estimates on $AI \times Strong \times Minority$ are all small and statistically insignificant, suggesting that the instrument does not affect the racial gap in mortgage lending before 2008. This evidence supports the validity of our instrument.

INSERT TABLE 4 HERE

We then implement a two-stage least squares approach using this instrument to estimate the causal effect of bank AI adoption on the racial gap in mortgage lending. The first- and second-stage regression results are reported in Table 5. In the first stage, we regress the interaction term $Bank \times AI \times Minority$ on the interaction between the instrument $AI \times Strong$ and the minority dummy $Minority$, along with the other control variables. The coefficient estimate on $AI \times Strong \times Minority$ is positive and statistically significant, suggesting that the instrument satisfies the relevance condition. The first-stage Kleibergen-Paap F-statistic is 10.26, above the conventional rule-of-thumb threshold. To address residual weak-instrument concerns, we further report the [Anderson and Rubin \(1949\)](#) weak-instrument-robust 95% confidence intervals for $Bank \times AI \times Minority$; these intervals exclude zero for both the denial

and interest-rate specifications, suggesting that our results are not biased by weak instruments.

INSERT TABLE 5 HERE

In column (2) of Table 5, we present the second-stage regression results for denial. The coefficient estimate on *Bank AI* $\times$ *Minority* remains positive and statistically significant. The magnitude of the IV estimate is more than double that of the OLS estimate in Table 2, suggesting that the OLS estimate is biased downward. One possible interpretation is that banks pursue multiple strategies to mitigate human-driven racial disparities, and AI adoption is one such strategy whose disparity-widening effect is partially offset by other contemporaneous fairness-oriented investments. Conventional OLS, which absorbs both the positive AI effect and the offsetting fairness investments, will then understate the AI-only effect that the IV isolates.

In column (3) of Table 5, we present the second-stage regression results for interest rate. The coefficient estimate on *Bank AI* $\times$ *Minority* is again positive and statistically significant, and the magnitude is more than double that of the OLS estimate in Table 3. Overall, the 2SLS results support a causal interpretation of the relationship between bank AI adoption and the widening racial gap in mortgage lending.

4 Alternative Explanations

4.1 The Composition of Loan Officers

We argue that the impact of AI on the racial gap is driven by AI’s flexibility to uncover structural relationships or its triangulation of excluded characteristics, such as race. However, the effect documented above could also be driven by the changing racial composition of the mortgage lenders induced by the adoption of AI. For example, if banks’ AI adoption disproportionately replaces minority loan officers, the decreased presence of minority loan officers could then lead to an increased racial gap ([Frame et al. 2025](#)). To see whether this is the case, we regress the percentage of minority loan officers on *Bank AI* in a lender-year regression. We report the results in column (1) of Table 6. The coefficient estimate on *Bank AI* is small and statistically insignificant, suggesting that AI adoption does not affect the share of minority loan officers at the lender.

INSERT TABLE 6 HERE

4.2 The Composition of Loan Applications

Alternatively, the adoption of AI and the associated announcement could affect borrowers’ choices. For example, minority borrowers may think that the bank AI adoption could mitigate human bias and then choose to borrow from banks that use more AI. More applications in general or more applications from minority applicants may affect banks’ underwriting or pricing decisions. To see whether this is the case, we first examine how bank AI adoption affects the total number of applications using the fixed-effects Poisson model.

The result is presented in column (2) of Table 6. The coefficient estimate on *Bank AI* is again statistically insignificant. However, the estimate is positive, suggesting that AI adoption may increase the total number of applications. We then focus on the percentage of minority applications to see if AI adoption disproportionately attracts minority applicants. The result is presented in column (3) of Table 6. The coefficient estimate is again small and statistically insignificant, suggesting that AI adoption does not disproportionately increase minority applications.

4.3 Bank AI Adoption and Applicant Risk

An alternative concern is that banks adopt AI when the riskiness of their applicant pool changes, which could mechanically affect approval and pricing outcomes. Table 7 addresses this concern by testing whether bank AI adoption is associated with observable measures of applicant risk, focusing on the combined loan-to-value ratio (CLTV) and the loan-to-income ratio. Columns (1) and (2) examine the average relationship between lagged bank AI adoption and these two risk proxies. Columns (3) and (4) include the interaction term *Bank AI* × *Minority* to test whether AI adoption is associated with differential changes in these risk measures for minority applicants.

INSERT TABLE 7 HERE

The results suggest that AI adoption does not increase applicant risk, and in particular does not increase the risk of minority applications. In Columns (1) and (2), the coefficients on *Bank AI* are negative and statistically insignificant for both CLTV and loan-to-income, indicating no meaningful change in these risk proxies as banks adopt AI. In Columns (3)

and (4), the interaction term *Bank AI*×*Minority* is negative and statistically significant, implying that, if anything, minority applicants at banks with greater AI adoption appear observably less risky along these dimensions. Taken together, these findings alleviate concerns that our main results are driven by banks attracting riskier applicants (especially riskier minority applicants) when adopting AI, and instead support an interpretation in which AI adoption affects underwriting decisions and pricing rather than the risk composition of loan applications.

5 Cross-sectional Heterogeneity

5.1 Securitization Status

The impact of AI adoption on racial disparities in mortgage lending may vary across different segments of the market. One important dimension of heterogeneity is the securitization status of loans. Securitized loans are those that are bundled and sold to investors as mortgage-backed securities, while non-securitized loans are held on banks’ balance sheets. In particular, a substantial share of mortgages are sold to government-sponsored enterprises (GSEs), including Fannie Mae, Freddie Mac, and Ginnie Mae. These GSEs have their own underwriting and pricing algorithms, and hence the lenders’ adoption of AI may have limited effect on these loans. For originated loans, the HMDA data report whether a loan is sold during the reporting year, and if it is sold, to whom it is sold. We use this information to split the sample of approved loans into three samples: not sold during the reporting year, sold to GSEs during the reporting year, and sold to other entities during the reporting year.

We then examine the effect of Bank AI adoption on racial disparities in interest rate in each of the three loan segments. The results are presented in Table 8. In column (1), for loans that are not sold during the reporting year, the coefficient estimate on *Bank AI*×*Minority* is positive and statistically significant. In contrast, in column (2) for loans sold to GSEs, the coefficient estimate is much smaller and statistically insignificant, suggesting that AI adoption does not significantly affect racial disparities in this segment. In column (3) for loans sold to other entities, the coefficient estimate is again positive and statistically significant, although the magnitude is smaller than that for non-securitized loans. These patterns are consistent with the institutional features of these segments: non-securitized loans are more likely to be influenced by lenders’ underwriting and pricing decisions, while GSE-sold loans are subject to standardized underwriting criteria that may limit the scope for AI-driven disparities. Loans sold to other entities may fall somewhere in between, depending on the nature of the buyers and their own underwriting practices.

INSERT TABLE 8 HERE

5.2 Purchase Versus Refinance Loans

Table 9 reports results separately for purchase and refinance loans. For purchase mortgages, the key interaction term, *Bank AI*×*Minority*, is positive and statistically significant in both the denial and interest-rate regressions, indicating that AI adoption is associated with a larger racial gap in underwriting and pricing for home-purchase credit. By contrast, the corresponding estimates for refinance loans are much smaller and statistically insignificant.

INSERT TABLE 9 HERE

These differences are consistent with institutional features of the two markets. Purchase loans are often originated to households facing tighter cash-flow constraints and greater information frictions, and the underwriting decision must incorporate property-level details and down-payment constraints. As a result, screening intensity is high and model-based decision tools can have a large marginal effect on approval and pricing. If AI models exploit variables that proxy for protected characteristics through historical and structural channels, the resulting screening rules can widen minority gaps most strongly in the purchase segment.

Refinance loans, in contrast, are typically made to borrowers with an existing mortgage and an observable payment history, implying positive selection on unobserved credit quality and richer lender-observable information. In addition, refinancing is heavily driven by interest-rate incentives and often processed through more standardized, eligibility-based workflows. These features reduce the scope for incremental AI-driven screening to differentially affect minority versus non-minority applicants. Consistent with this interpretation, Table 9 shows that the differential effect of AI adoption is concentrated in purchase loans.

5.3 FinTech Lenders

FinTech lenders represent a distinct segment of the mortgage market characterized by technology-intensive business models, algorithmic underwriting platforms, and digital-first customer acquisition (Buchak et al., 2018; Fuster et al., 2019). These lenders have embedded automated decision-making into their core operations from inception, rather than retrofitting AI into legacy processes. This structural difference raises the question of whether AI adoption has similar effects on racial disparities in FinTech versus traditional banking contexts.

Table 10 presents results for the FinTech subsample. In stark contrast to our main findings for traditional banks, the coefficient on *Bank AI* $\times$ *Minority* is negative and statistically insignificant across all specifications for both denial and interest rate outcomes. For denial, the point estimates range from -0.274 to -0.351 and are never significant at conventional levels. Similarly, for interest rates, the estimates are imprecise and do not indicate that AI adoption widens racial gaps among FinTech lenders.

INSERT TABLE 10 HERE

Several institutional features may explain why AI adoption does not amplify racial disparities in the FinTech segment. First, FinTech lenders typically rely on fully automated underwriting systems with limited loan officer discretion throughout the lending process. Because human judgment plays a minimal role, the marginal impact of adopting additional AI capabilities may be smaller. These firms already operate in a high-automation environment where incremental AI investments affect operational scale rather than fundamentally altering decision rules. Second, FinTech lenders often employ newer data sources and alternative credit scoring models that were developed in the post-fair-lending regulatory environment and may incorporate fairness constraints or bias-mitigation techniques more systematically than legacy models at traditional banks. Third, FinTech business models emphasize standardization and scale, which may reduce the scope for heterogeneous treatment effects across borrower subgroups that can emerge when AI tools are layered onto discretionary human processes.

The divergence between FinTech and traditional bank results suggests that the institutional context in which AI is deployed matters critically for its distributional consequences.

At traditional banks, AI adoption appears to interact with existing organizational structures, legacy data, and human decision-making in ways that disadvantage minority applicants. In contrast, FinTech lenders' technology-native operations may better insulate underwriting decisions from the amplification of historical disparities. These findings underscore that the equity implications of AI in lending depend not only on the technology itself but also on the organizational and operational environment in which it is embedded.

6 Robustness and Additional Tests

6.1 Pre-Application Versus Underwriting Stage

To better understand where differential treatment arises in the lending process, we decompose mortgage outcomes by stage. The HMDA mortgage application data not only include completed applications that are submitted for underwriting, but also incomplete applications that are either withdrawn or closed because of incomplete information. In the analysis above, we also treat these applications as being denied. However, these outcomes occur before underwriting, and distinguishing these outcomes from underwriting decisions could provide more information on the source of the effect of AI on mortgage lending. In particular, application withdrawn and closed for incomplete information happen in the pre-application information collection stage, in which AI is less likely to be involved. In contrast, AI is more likely to be used in the underwriting stage to conduct in-depth analysis of the risk of the applications. We therefore expect the effect to concentrate in the underwriting stage if the effect is truly driven by AI adoption.

We therefore first examine how bank AI adoption affects whether an application is more or less likely to be withdrawn or closed for incomplete information. The results are presented in column (1) of Table 11. The coefficient estimate on *Bank AI* $\times$ *Minority* is positive but statistically insignificant; we cannot reject a zero differential effect at the pre-application stage. We then focus on completed applications that are submitted for underwriting. The results are presented in column (2) of Table 11. The coefficient estimate on *Bank AI* $\times$ *Minority* is positive and statistically significant, indicating that the differential impact of AI on minority applicants is concentrated at the underwriting stage.

INSERT TABLE 11 HERE

6.2 Origination Fees

We also examine whether AI adoption affects the racial gap in origination fees, which are upfront fees charged to borrowers at loan closing. Origination fees are an important component of the total cost of credit and can reflect differences in underwriting risk assessment and pricing strategies. Bhutta and Hizmo (2021) show that ignoring origination fees could bias the estimates of the racial gap in mortgage pricing. The HMDA data provide the following information on origination fees: the total loan cost, the sum of borrower paid origination charges and borrower paid third party charges, such as appraisal fees and title insurance; the origination charges, the total fees charged by lenders for originating the loan, including discount points paid by the borrower; and the discount points, the total amount spent on discount points, including those paid by the borrower and those paid by the others. We scale these fees by the loan amount to construct the origination fee rate. We then regress

the origination fee rate on the same set of controls as in the interest rate regression, and focus on the coefficient estimate on *Bank AI×Minority* to see whether AI adoption is associated with a larger racial gap in origination fees.

The results are presented in columns (1)–(3) in Table 12. The coefficient estimates on *Bank AI×Minority* are all small and statistically insignificant, suggesting that Bank AI adoption does not affect the racial gap in origination fees. This pattern contrasts with the significant effect of AI adoption on the racial gap in interest rates, and suggests that the impact of AI on mortgage pricing is primarily reflected in the contractual interest rate rather than upfront fees. This finding is consistent with the interpretation that AI adoption affects risk assessment and pricing decisions embedded in the interest rate setting process, rather than fee structures that may be more standardized or subject to regulatory constraints.

INSERT TABLE 12 HERE

The HMDA data also report the rate spread, the annual percentage rate (APR) minus the average prime offer rate (APOR) for a comparable mortgage. The rate spread captures the total cost of credit relative to market benchmarks and can reflect both interest rate and fee components. We also regress the rate spread on the same set of controls, and focus on the coefficient estimate on *Bank AI×Minority* to see whether AI adoption is associated with a larger racial gap in overall pricing. The results are presented in column (4) of Table 12. The coefficient estimate on *Bank AI×Minority* is positive and statistically significant, suggesting that AI adoption is associated with a larger racial gap in overall mortgage pricing as measured by the rate spread. This pattern reinforces our findings for contractual interest rates.

6.3 Additional Evidence from McDash Data

The HMDA data miss one important underwriting variable, the credit score. Additionally, the HMDA data provide the combined loan-to-value ratio rather than the loan-to-value ratio. Both of these could raise concerns about omitted variable bias. To mitigate this concern, we augment the HMDA data with the McDash data. McDash provides comprehensive loan-level records with detailed credit characteristics not available in HMDA, including borrower credit scores and precise loan-to-value ratios. This additional information allows us to more thoroughly control for applicant creditworthiness and test whether our results are robust to the inclusion of these key underwriting variables. We follow the literature and merge the two datasets using variables common to both, including year of origination, location (census tract in HMDA and zip code in McDash), loan amount, loan term, property value, interest rate, the debt-to-income ratio, loan type (FHA, VA, or Conventional), loan purpose (refinance or purchase). We retain only unique matches and achieve a match rate of about 20%.⁵

Table 13 presents results for mortgage interest rates using the McDash sample. Column (1) replicates our baseline interest-rate specification from Table 3 in the matched data. The coefficient estimate on *Bank AI × Minority* is 0.845 and statistically significant at the 1% level, indicating that AI adoption is associated with higher interest rates for minority applicants relative to non-minority applicants. This estimate is smaller in magnitude than the corresponding estimates in Table 3.

INSERT TABLE 13 HERE

⁵Our match rate is lower than the match rate between confidential HMDA and McDash because confidential HMDA reports origination date instead of just origination year.

Column (2) adds two critical underwriting variables: the loan-to-value ratio and the borrower’s credit score. Both variables enter with expected signs: higher LTV increases the mortgage interest rate, while higher credit scores reduce the interest rate. Both coefficient estimates are statistically significant at the 1% level. Importantly, the inclusion of these additional risk controls leaves the coefficient on *Bank AI*×*Minority* virtually unchanged at 0.885, suggesting that our main results are not driven by omitted measures of applicant credit quality that happen to correlate with both AI adoption and minority status.

Columns (3) and (4) address concerns about time-varying compositional shifts across different mortgage market segments. We construct securitization buckets that classify loans by their conforming status and government guarantee eligibility, as these categories face different underwriting standards and competitive dynamics (Bartlett et al. 2022). Column (3) includes bucket×year fixed effects to control for annual shifts in lending standards within each market segment, while column (4) employs bucket×year-month fixed effects to capture even more granular temporal variation. The coefficient on *Bank AI*×*Minority* remains positive and statistically significant in both specifications (0.848 and 0.632, respectively), though the magnitude attenuates somewhat with the most demanding fixed effects structure.

The robustness of our findings in the McDash data is particularly reassuring given this dataset’s advantages for mortgage research. The availability of credit scores, arguably the single most important predictor in automated underwriting systems, allows us to verify that AI’s differential impact on minority borrowers persists even after accounting for this key variable. Moreover, the granular fixed effects in columns (3) and (4) address concerns that our results might be driven by banks adopting AI precisely when they are shifting their portfolio composition toward or away from particular loan segments that differ in minority representa-

tion. The stability of our core finding across these demanding specifications strengthens the interpretation that AI adoption itself contributes to widening racial disparities in mortgage lending.

6.4 Other Technology Investment

Banks' adoption of AI may be correlated with other technology investments that could affect mortgage lending outcomes. For example, banks that adopt AI may also invest in other technologies that could affect the racial gap in mortgage lending, such as online application platforms or automated underwriting systems. To mitigate this concern, we control for the share of job postings related to other technologies, in particular, the share of job postings related to automation. The results are presented in Table 14. The coefficient estimates on *Bank AI* × *Minority* remain positive and statistically significant, suggesting that the effect of AI adoption on the racial gap in mortgage lending is not driven by other technology investments. Interestingly, the coefficient estimates on the interaction between the minority dummy and the share of automation-related job postings are negative and one of them is statistically significant, suggesting that automation adoption may be associated with a smaller racial gap in mortgage lending. This pattern is consistent with the interpretation that automation adoption could reduce human involvement, which could reduce human bias in mortgage lending.

6.5 Resume-based Bank AI Adoption Measure

We also construct an alternative bank AI adoption measure based on the share of employees with AI-related skills, following the resume-based method in [Babina et al. \(2024\)](#). We then regress mortgage denial and interest rate on the interaction between this alternative measure and the minority dummy, with the same controls and fixed effects as the main specification. Table 15 reports the results. The coefficient estimates on the interaction term are positive and statistically significant for both outcomes, indicating that our findings are robust to an alternative AI measure. The magnitudes are smaller than the job-posting-based estimates, consistent with greater measurement error in resume-classified AI skills than in posted skill requirements.

7 Conclusion

We show that banks’ adoption of artificial intelligence amplifies racial disparities in mortgage lending. Using 2022–2024 HMDA data merged with a bank-year measure of AI adoption derived from job postings, we document that greater AI adoption is associated with significantly worse approval and pricing outcomes for minority borrowers relative to non-minority borrowers. These patterns are robust to rich controls and high-dimensional fixed effects, including bank $\times$ year and census-tract $\times$ year, and are reinforced by instrumental-variable estimates that exploit banks’ historical exposure to the supply of AI graduates.

Our analysis pinpoints where and how these disparities arise. The effects concentrate in underwriting decisions rather than at the pre-application stage, and do not operate through changes in applicant composition or minority loan-officer representation, pointing to AI-

assisted screening and pricing as the central mechanism. The effects are present for purchase loans but absent for refinance loans, consistent with AI having greater scope to amplify disparities when borrower risk is less directly observable. The absence of disparity-widening effects for GSE-backed loans suggests that standardized underwriting guidelines can serve as a check on AI-driven bias. Perhaps most strikingly, AI adoption does not amplify racial disparities among FinTech lenders, where algorithmic underwriting has been embedded from inception. This divergence underscores that the equity implications of AI depend critically on the institutional context (legacy processes, human discretion, and historical data) in which it is deployed.

These findings have implications for financial regulation, bank governance, and the design of AI systems in credit markets. As regulators evaluate the growing role of AI in lending, our results suggest that fair-lending oversight should extend beyond whether protected characteristics are included as model inputs to encompass the distributional consequences of AI-assisted decision-making. For bank managers, the findings highlight the importance of rigorous disparate-impact testing and ongoing monitoring of AI tools, particularly in settings where algorithmic predictions interact with legacy underwriting processes. More broadly, the results illustrate that efficiency-enhancing technologies can have unintended equity consequences when embedded in core financial decision processes, a consideration that applies well beyond mortgage lending.

Several open questions remain for future research. First, confidential HMDA data with borrower credit scores would permit a sharper test of which specific applicant features AI weights differently across racial groups, isolating the flexibility from the triangulation channel. Second, our identification of the bank-AI effect is silent on which AI applications inside

the bank drive the result; firm-level case studies or supervisory data on model documentation could narrow the mechanism further. Third, our findings raise the question of whether disparate-impact audits, model-explainability mandates, or constrained fairness-aware estimators can attenuate the disparity-widening effect we document without sacrificing the efficiency gains of AI-assisted underwriting. Empirical evaluation of such interventions is a natural next step.

References

- Alekseeva, L., J. Azar, M. Giné, S. Samila, and B. Taska. 2021. The demand for AI skills in the labor market. *Labour Economics* 71:102002–.
- Ambrose, B. W., J. N. Conklin, and L. A. Lopez. 2021. Does borrower and broker race affect the cost of mortgage credit? *Review of Financial Studies* 34:790–826.
- Anderson, T. W., and H. Rubin. 1949. Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of Mathematical Statistics* 20:46–63.
- Azar, J. A., I. Marinescu, M. I. Steinbaum, and B. Taska. 2018. Concentration in US labor markets: Evidence from online vacancy data. Working Paper, National Bureau of Economic Research. Working Paper 24395.
- Babina, T., and A. Fedyk. 2025. The effects of AI on firms and workers. Brookings.
- Babina, T., A. Fedyk, A. He, and J. Hodson. 2024. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151:103745–.
- . 2025. Firm investments in artificial intelligence technologies and changes in workforce composition. In S. Basu, L. Eldridge, J. Haltiwanger, and E. Strassner, eds., *Technology, Productivity, and Economic Growth*, Studies in Income and Wealth. University of Chicago Press.
- Babina, T., A. Fedyk, A. X. He, and J. Hodson. 2023. Firm investments in artificial intelligence technologies and changes in workforce composition. Working Paper 31325, National Bureau of Economic Research.

- Bartlett, R., A. Morse, R. Stanton, and N. Wallace. 2022. Consumer-lending discrimination in the fintech era. *Journal of Financial Economics* 143:30–56.
- Berg, T., V. Burg, A. Gombović, and M. Puri. 2020. On the rise of fintechs—credit scoring using digital footprints. *Review of Financial Studies* 33:2845–97.
- Bhutta, N., and A. Hizmo. 2021. Do minorities pay more for mortgages? *Review of Financial Studies* 34:763–806.
- Blair, P. Q., and D. J. Deming. 2020. Structural increases in skill demand after the great recession. Working Paper 26680, National Bureau of Economic Research.
- Bowen, D. E., S. M. Price, L. C. Stein, and K. Yang. 2025. Measuring and mitigating racial disparities in LLMs: Evidence from a mortgage underwriting experiment. Working paper, Lehigh University.
- Buchak, G., G. Matvos, T. Piskorski, and A. Seru. 2018. Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics* 130:453–83.
- Chu, Y. 2025. Partisan loan officers and the racial gap in mortgage lending. Working paper, University of North Carolina at Charlotte.
- Chu, Y., B. Huang, and C. Zhang. 2025. The color of hedge fund activism. *Management Science* 71:5142–61.
- Deming, D. J., and L. B. Kahn. 2018. Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics* 36:337–69.

- Dwork, C., M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference (ITCS)*, 214–26.
- Federal Reserve Bank of Minneapolis. 2024. Lender-reported reasons for mortgage denials don’t explain racial disparities. Federal Reserve Bank of Minneapolis article.
- Fetter, D. K., G. Howard, L. Kletzer, and B. Taska. 2023. Higher mortgage denials for solo applicants feed racial disparities in lending. Federal Reserve Bank of Minneapolis (Minneapolis Fed) article.
- Frame, W. S., R. Huang, E. X. Jiang, Y. Lee, W. S. Liu, E. J. Mayer, and A. Sunderam. 2025. The impact of minority representation at mortgage lenders. *Journal of Finance* 80:1209–60.
- Fuster, A., P. Goldsmith-Pinkham, T. Ramadorai, and A. Walther. 2022. Predictably unequal? The effects of machine learning on credit markets. *Journal of Finance* 77:5–47.
- Fuster, A., M. Plosser, P. Schnabl, and J. Vickery. 2019. The role of technology in mortgage lending. *Review of Financial Studies* 32:1854–99.
- Goldfarb, A., B. Taska, and F. Teodoridis. 2023. Could machine learning be a general purpose technology? a comparison of emerging technologies using data from online job postings. *Research Policy* 52:104651–.
- Goodman, L., and J. Zhu. 2023. Single borrowers, joint borrowers, and mortgage credit access: Evidence from the U.S. housing market. *Journal of Housing Economics* 62:102002–. doi:10.1016/j.jhe.2023.102002.

- Hardt, M., E. Price, and N. Srebro. 2016. Equality of opportunity in supervised learning. In *Advances in Neural Information Processing Systems*.
- Howell, S. T., T. Kuchler, D. Snitkof, J. Stroebel, and J. Wong. 2024. Lender automation and racial disparities in credit access. *Journal of Finance* 79:1457–512. doi:10.1111/jofi.13303.
- Hunter, W. C. 1996. The cultural maze of mortgage lending. *Economic Commentary* .
- Jagtiani, J., and C. Lemieux. 2018. The roles of alternative data and machine learning in fintech lending: Evidence from the lendingclub consumer platform. Working paper, Federal Reserve Bank of Philadelphia.
- Kleinberg, J., H. Lakkaraju, J. Leskovec, J. Ludwig, and S. Mullainathan. 2018. Human decisions and machine predictions. *Quarterly Journal of Economics* 133:237–93.
- Lewis-Faupel, S., and N. Tenev. 2024. Racial and ethnic disparities in mortgage lending: New evidence from expanded HMDA data. *arXiv preprint arXiv:2405.00895* .
- Munnell, A. H., G. M. B. Tootell, L. E. Browne, and J. McEneaney. 1996. Mortgage lending in boston: Interpreting hmda data. *American Economic Review* 86:25–53.
- Yinger, J. 1995. *Closed doors, opportunities lost: The continuing costs of housing discrimination*. New York: Russell Sage Foundation.

Table 1: Summary Statistics

This table presents the summary statistics of variables. *Denial* is a dummy variable equal to one if a loan application is rejected, and zero otherwise. *Interest Rate* is the contractual interest rate of the loan. *Rate Spread* is the difference between a loan's Annual Percentage Rate (APR) and the Average Prime Offer Rate (APOR) on a comparable mortgage. *Minority* is a dummy variable equal to one if the borrower is African American or Hispanic, and zero otherwise. *Co-applicant* is a dummy variable equal to one if the loan application includes a co-applicant, and zero otherwise. *Loan-to-Income* is defined as the ratio of the loan amount to the applicant's income. *Log Income* is the natural logarithm of the applicant's reported income. *Loan Term* is the loan maturity in months. *CLTV* is the combined loan-to-value ratio. *Female* is a dummy variable equal to one if the applicant is female, and zero otherwise. *Applicant Age* categories are dummy variables for age groups. *Withdrawn or Incomplete* is a dummy variable for applications withdrawn or closed for incomplete information. *AI Job Number* is the number of AI skills-related job postings. *AI Job Share* is the proportion of AI skills-related job postings.

	(1) N	(2) mean	(3) std	(4) p5	(5) p10	(6) p25	(7) p50	(8) p75	(9) p90	(10) p95
Loan Characteristics										
Denial	13,915,800	0.201	0.400	0	0	0	0	0	1	1
Interest Rate	11,086,894	6.160	2.019	2.999	3.500	4.990	6.250	7.249	8.500	9.500
Rate Spread	10,463,635	0.465	1.374	-1.325	-0.710	-0.113	0.382	0.989	1.820	2.750
Minority	13,915,800	0.260	0.439	0	0	0	0	1	1	1
Co-applicant	13,915,800	0.413	0.492	0	0	0	0	1	1	1
Loan-to-Income	13,915,800	2.236	1.562	0.280	0.417	0.866	2.050	3.269	4.327	5
Log Income	13,915,800	4.627	0.680	3.584	3.807	4.174	4.595	5.030	5.489	5.820
Loan Term	13,915,800	356.5	106.2	180	240	300	360	360	360	420
CLTV	13,915,800	73.27	22.63	25.94	38.89	60.02	79.67	92.42	97	100
Female	13,915,800	0.370	0.483	0	0	0	0	1	1	1
Applicant Age ≤ 25	13,915,800	0.031	0.173	0	0	0	0	0	0	1
Applicant Age 26–35	13,915,800	0.219	0.414	0	0	0	0	0	1	1
Applicant Age 36–50	13,915,800	0.348	0.476	0	0	0	0	1	1	1
Applicant Age 51–65	13,915,800	0.279	0.448	0	0	0	0	1	1	1
Applicant Age > 65	13,915,800	0.123	0.328	0	0	0	0	0	1	1
Withdrawn or Incomplete	18,975,690	0.191	0.393	0	0	0	0	0	1	1
Bank AI										
AI Job Number(scaled by 1000)	13,915,800	0.0399	0.172	0	0	0	0	0.0100	0.0410	0.190
AI job share	13,915,800	0.0114	0.0202	0	0	0	0	0.0153	0.0358	0.0603

Table 2: Bank AI Adoption and Mortgage Underwriting

This table presents baseline results on the impact of bank AI adoption on mortgage underwriting. The dependent variable is *Denial*, equaling one if the loan is rejected and zero otherwise. The key independent variable is *Bank AI*×*Minority*, which is the interaction term between *Bank AI* and *Minority*. *Minority* equals one if the borrower is African American or Hispanic and zero otherwise. *Bank AI* is the share of the bank's AI skill-related job postings lagged by one year. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)
Bank AI × Minority	0.265*** (2.67)	0.375*** (3.20)	0.375*** (3.24)
Minority	0.045*** (12.89)	0.045*** (12.52)	0.045*** (12.55)
Bank AI	-0.352* (-1.88)		
Co-applicant	-0.018*** (-8.57)	-0.018*** (-8.48)	-0.018*** (-8.51)
Loan-to-Income	-0.007*** (-2.90)	-0.007*** (-2.95)	-0.007*** (-2.94)
Log Income	-0.035*** (-10.48)	-0.035*** (-10.53)	-0.035*** (-10.61)
Loan Term	0.000*** (3.16)	0.000*** (3.08)	0.000*** (3.09)
CLTV	0.001*** (7.76)	0.001*** (7.88)	0.001*** (7.94)
Female	0.001 (1.48)	0.001 (1.46)	0.001 (1.51)
Bank FE	YES		
MSA FE	YES	YES	YES
Bank×Year FE		YES	YES
Census Tract×Year FE			YES
Census Tract FE	YES	YES	
Debt-to-Income FE	YES	YES	YES
Loan Type FE	YES	YES	YES
Loan Purpose FE	YES	YES	YES
Applicant Age Group FE	YES	YES	YES
Adjusted R-squared	0.365	0.369	0.370
Observations	13,918,119	13,918,066	13,916,130

Table 3: Bank AI Adoption and Mortgage Interest Rate

This table presents results on the impact of bank AI adoption on mortgage interest rates. The dependent variable is *Interest Rate*, the contractual interest rate of the loan. The key independent variable is *Bank AI*×*Minority*, which is the interaction term between *Bank AI* and *Minority*. *Minority* equals one if the borrower is African American or Hispanic and zero otherwise. *Bank AI* is the share of the bank's AI skill-related job postings lagged by one year. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)
Bank AI × Minority	2.069** (2.33)	1.958** (2.35)	2.019** (2.47)
Minority	-0.004 (-0.32)	-0.001 (-0.11)	0.000 (0.04)
Bank AI	0.558 (0.24)		
Co-applicant	-0.073*** (-9.18)	-0.067*** (-9.04)	-0.066*** (-9.05)
Loan-to-Income	-0.171*** (-6.28)	-0.176*** (-6.44)	-0.173*** (-6.33)
Log Income	0.078** (2.17)	0.067* (1.87)	0.066* (1.84)
Loan Term	0.002*** (6.53)	0.002*** (6.35)	0.002*** (6.35)
CLTV	0.002** (2.44)	0.003*** (3.19)	0.003*** (3.33)
Female	0.005 (1.37)	0.005 (1.56)	0.005 (1.49)
Bank FE	YES		
MSA FE	YES	YES	YES
Bank×Year FE		YES	YES
Census Tract×Year FE			YES
Census Tract FE	YES	YES	
Debt-to-Income FE	YES	YES	YES
Loan Type FE	YES	YES	YES
Loan Purpose FE	YES	YES	YES
Applicant Age Group FE	YES	YES	YES
Adjusted R-squared	0.541	0.578	0.586
Observations	11,088,192	11,088,138	11,085,464

Table 4: The Instrument and the Racial Gap in Mortgage Lending before 2008

This table presents the results of testing the instrument variable, *AI Strong*, on the racial gap in mortgage lending before 2008. The dependent variable is *Denial*, equaling one if the loan is rejected and zero otherwise. The key independent variable is *AI Strong*×*Minority*, which is the interaction term between the instrument *AI Strong* and *Minority*. *Minority* equals one if the borrower is African American or Hispanic and zero otherwise. *AI Strong* equals one if at least one employee in 2008–2010 graduated from an AI strong university, and zero otherwise. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1) 2005	(2) 2006	(3) 2007	(4) 2005–2007
AI Strong × Minority	0.002 (0.31)	0.000 (0.05)	0.012 (1.60)	0.005 (0.77)
Minority	0.046*** (8.53)	0.052*** (9.11)	0.064*** (12.88)	0.054*** (11.14)
Co-applicant	-0.035*** (-12.16)	-0.036*** (-9.27)	-0.042*** (-8.98)	-0.038*** (-10.88)
Loan-to-Income	0.004*** (4.55)	0.004*** (3.41)	0.004*** (3.08)	0.004*** (3.95)
Log Income	-0.045*** (-9.33)	-0.047*** (-8.87)	-0.044*** (-6.99)	-0.045*** (-9.08)
Female	-0.007*** (-4.55)	-0.005*** (-2.89)	-0.008*** (-3.96)	-0.007*** (-4.63)
Bank×Year FE	YES	YES	YES	YES
MSA FE	YES	YES	YES	YES
Census Tract×Year FE	YES	YES	YES	YES
Loan Type FE	YES	YES	YES	YES
Loan Purpose FE	YES	YES	YES	YES
R-squared	0.226	0.187	0.195	0.203
Observations	1,539,583	1,580,696	1,593,514	4,713,793

Table 5: IV 2SLS Results: Bank AI adoption and the Racial Gap in Mortgage Lending

This table presents IV 2SLS results instrumenting Bank AI adoption using banks' *ex ante* exposure to universities' supply of AI graduates following Babina et al. (2024). Column (1) reports the first-stage results. Columns (2) and (3) report the second-stage results for *Denial* and *Interest Rate*, respectively. The instrument *AI Strong* equals one if at least one employee in 2008–2010 graduated from an AI strong university, and zero otherwise. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)
	First stage	Second stage	
	Bank AI $\times$ Minority	Denial	Interest Rate
AI Strong $\times$ Minority	0.027** (2.31)		
Bank AI $\times$ Minority		1.156** (2.23)	3.552** (2.53)
Anderson-Rubin CI		(0.140, 2.171)	(0.798, 6.305)
Minority	0.010*** (6.13)	0.034*** (3.98)	-0.020 (-1.06)
Co-applicant	-0.000 (-1.33)	-0.018*** (-8.08)	-0.065*** (-8.63)
Loan-to-Income	0.000 (0.76)	-0.007*** (-2.81)	-0.188*** (-6.69)
Log Income	0.000** (2.03)	-0.035*** (-9.58)	0.045 (1.22)
Loan Term	0.000* (1.82)	0.000*** (2.93)	0.002*** (6.00)
CLTV	-0.000** (-2.19)	0.001*** (7.68)	0.003*** (3.67)
Female	0.000 (1.55)	0.001 (1.37)	0.004 (1.14)
MSA FE	YES	YES	YES
Bank $\times$ Year FE	YES	YES	YES
Census Tract $\times$ Year FE	YES	YES	YES
Debt-to-Income FE	YES	YES	YES
Loan Type FE	YES	YES	YES
Loan Purpose FE	YES	YES	YES
Applicant Age Group FE	YES	YES	YES
Adjusted R-squared	0.590		
Observations	12,423,664	12,423,664	9,903,992
First-stage F-statistic	10.26		

Table 6: Bank AI Adoption, the Composition of Loan Officers and Loan Applications

This table presents baseline results examining the impact of bank AI adoption on the composition of loan officers and loan applications. The sample is at the lender-year level. The dependent variables are *% Minority LO*, the percent of minority loan officers of the lender, in column (1), *#Application*, the total number of applications the lender receives, in column (2), and *% Minority Application*, the percent of minority loan applications in column (3). The key independent variable is *Bank AI*, the share of bank's AI skill-related job postings lagged by one year. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)
	%Minority LO	#Application	%Minority Application
Bank AI	0.025 (0.23)	1.853 (1.46)	-0.008 (-0.07)
Bank FE	YES	YES	YES
Year FE	YES	YES	YES
Adjusted R-squared	0.742	0.987	0.93
Observations	7,900	7,900	7,900

Table 7: Bank AI Adoption and Applicant Risk

This table examines whether bank AI adoption is associated with changes in applicant risk. The dependent variables are the combined loan-to-value ratio (*CLTV*) in Columns (1) and (3) and the loan-to-income ratio in Columns (2) and (4). *Bank AI* is the share of the bank's AI skill-related job postings lagged by one year. *Minority* equals one for African American or Hispanic applicants and zero otherwise. Columns (1) and (2) estimate specifications with *Bank AI* only, while Columns (3) and (4) additionally include the interaction term *Bank AI*×*Minority* to test whether AI adoption is associated with differential changes in applicant risk for minority applications. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1) CLTV	(2) Loan-to-Income	(3) CLTV	(4) Loan-to-Income
Bank AI	-17.038 (-1.19)	-0.517 (-0.81)		
Bank AI × Minority			-38.880*** (-5.57)	-0.902** (-2.08)
Minority			5.149*** (30.36)	0.246*** (17.78)
Bank FE	YES	YES		
Year FE	YES	YES		
MSA FE	YES	YES	YES	YES
Census Tract FE	YES	YES		
Bank×Year FE			YES	YES
Census Tract×Year FE			YES	YES
Adjusted R-squared	0.213	0.274	0.226	0.284
Observations	14,337,883	18,213,292	14,335,969	18,211,480

Table 8: Bank AI Adoption and Securitization

This table presents results on the impact of bank AI adoption on mortgage lending outcomes across different loan types. The dependent variable is *Interest Rate*, the mortgage contractual interest rate. The key independent variable is *Bank AI*×*Minority*, which is the interaction term between *Bank AI* and *Minority*. *Minority* equals one if the borrower is African American or Hispanic and zero otherwise. *Bank AI* is the share of the bank's AI skill-related job postings lagged by one year. The sample is split into three groups based on the loan's securitization status: loans that are not sold, loans sold to government-sponsored enterprises (GSEs), and loans sold to others. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1) Not Sold	(2) GSE	(3) Other
Bank AI × Minority	2.446*** (2.83)	0.079 (0.36)	2.342* (1.76)
Minority	0.018 (1.10)	0.037*** (7.97)	-0.047*** (-2.88)
Co-applicant	-0.079*** (-8.03)	-0.023*** (-9.22)	-0.020** (-2.10)
Loan-to-Income	-0.269*** (-7.29)	-0.185*** (-25.12)	0.259*** (3.92)
Log Income	-0.125*** (-3.22)	0.034*** (3.24)	0.721*** (6.81)
Loan Term	0.001** (2.57)	0.005*** (49.38)	0.005*** (8.05)
CLTV	0.005*** (5.87)	0.001*** (2.70)	-0.004** (-2.11)
Female	0.006 (1.28)	0.018*** (8.53)	0.021*** (4.20)
Bank×Year FE	YES	YES	YES
MSA FE	YES	YES	YES
Census Tract×Year FE	YES	YES	YES
Debt-to-Income FE	YES	YES	YES
Loan Type FE	YES	YES	YES
Loan Purpose FE	YES	YES	YES
Applicant Age Group FE	YES	YES	YES
Adjusted R-squared	0.640	0.638	0.573
Observations	5,071,828	3,344,928	1,178,059

Table 9: The Impact of AI on Purchase and Refinance Loans

This table presents results on the impact of bank AI adoption on mortgage lending for purchase and refinance loans separately. The dependent variables are *Denial* in columns (1) and (3) and *Interest Rate* in columns (2) and (4). The key independent variable is *Bank AI*×*Minority*, which is the interaction term between *Bank AI* and *Minority*. *Minority* equals one if the borrower is African American or Hispanic and zero otherwise. *Bank AI* is the share of the bank's AI skill-related job postings lagged by one year. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	Purchase Loans		Refinance Loans	
	(1) Denial	(2) Interest Rate	(3) Denial	(4) Interest Rate
Bank AI × Minority	0.413** (2.51)	1.909*** (3.75)	0.209 (1.45)	0.018 (0.04)
Minority	0.028*** (12.25)	-0.005 (-0.54)	0.054*** (18.78)	0.044*** (5.91)
Co-applicant	-0.005*** (-4.21)	-0.008** (-2.31)	-0.027*** (-16.46)	-0.077*** (-12.12)
Loan-to-Income	-0.007*** (-3.29)	0.064** (2.30)	-0.006*** (-2.64)	-0.436*** (-15.31)
Log Income	-0.033*** (-11.29)	0.389*** (8.85)	-0.043*** (-11.41)	-0.311*** (-10.99)
Loan Term	0.000 (1.03)	0.003*** (6.07)	0.000*** (5.23)	0.003*** (7.58)
CLTV	0.001*** (7.66)	-0.004*** (-4.74)	0.001*** (9.89)	0.010*** (12.14)
Female	-0.001 (-1.19)	0.007*** (3.37)	-0.002* (-1.71)	0.008* (1.95)
MSA FE	YES	YES	YES	YES
Bank×Year FE	YES	YES	YES	YES
Census Tract×Year FE	YES	YES	YES	YES
Debt-to-Income FE	YES	YES	YES	YES
Loan Type FE	YES	YES	YES	YES
Applicant Age Group FE	YES	YES	YES	YES
Adjusted R-squared	0.293	0.432	0.346	0.662
Observations	6,873,021	6,172,708	3,191,518	2,410,863

Table 10: Bank AI Adoption and Mortgage Lending: FinTech Lenders

This table presents results on the impact of bank AI adoption on mortgage lending for FinTech lenders. The dependent variable is *Denial* in columns (1) through (3) and *Interest Rate* in columns (4) through (6). The key independent variable is *Bank AI*×*Minority*, which is the interaction term between *Bank AI* and *Minority*. *Minority* equals one if the borrower is African American or Hispanic and zero otherwise. *Bank AI* is the share of the bank's AI skill-related job postings lagged by one year. Columns (1) and (4) include bank and census tract fixed effects. Columns (2) and (5) replace bank fixed effects with bank×year fixed effects. Columns (3) and (6) further replace census tract fixed effects with census tract×year fixed effects. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	Denial			Interest Rate		
	(1)	(2)	(3)	(4)	(5)	(6)
Bank AI × Minority	-0.351 (-1.30)	-0.293 (-1.00)	-0.274 (-0.94)	-2.806 (-1.53)	0.252 (0.21)	0.293 (0.33)
Minority	0.051*** (4.77)	0.050*** (4.30)	0.049*** (4.33)	0.002 (0.05)	-0.064** (-2.12)	-0.061** (-2.36)
Co-applicant	-0.029*** (-6.96)	-0.028*** (-6.92)	-0.028*** (-7.05)	-0.082*** (-4.95)	-0.078*** (-4.56)	-0.076*** (-4.46)
Loan-to-Income	-0.005* (-1.78)	-0.005* (-1.75)	-0.005* (-1.71)	-0.293* (-1.79)	-0.294* (-1.78)	-0.281* (-1.71)
Log Income	-0.039*** (-7.15)	-0.039*** (-7.20)	-0.039*** (-7.85)	-0.016 (-0.07)	-0.019 (-0.09)	-0.014 (-0.07)
Loan Term	0.000 (0.69)	0.000 (0.67)	0.000 (0.69)	0.003** (2.59)	0.003** (2.55)	0.003** (2.68)
CLTV	0.002*** (5.61)	0.002*** (5.68)	0.002*** (5.94)	0.006 (0.92)	0.006 (0.99)	0.006 (1.03)
Female	-0.005* (-1.83)	-0.005* (-1.85)	-0.005* (-1.93)	-0.007 (-0.64)	-0.007 (-0.62)	-0.006 (-0.57)
Bank FE	YES			YES		
MSA FE	YES	YES	YES	YES	YES	YES
Bank×Year FE		YES	YES		YES	YES
Census Tract×Year FE			YES			YES
Census Tract FE	YES	YES		YES	YES	
Debt-to-Income FE	YES	YES	YES	YES	YES	YES
Loan Type FE	YES	YES	YES	YES	YES	YES
Loan Purpose FE	YES	YES	YES	YES	YES	YES
Applicant Age Group FE	YES	YES	YES	YES	YES	YES
Adjusted R-squared	0.307	0.309	0.312	0.494	0.517	0.519
Observations	3,053,454	3,053,454	3,043,112	2,401,901	2,401,901	2,386,945

Table 11: Bank AI adoption, Minority Lending, and the Stages of Mortgage Lending

This table examines the impact of bank AI adoption on mortgage lending at different stages, i.e., applicant withdrawal or incompleteness and denial by underwriting. Column (1) reports results where the dependent variable is *Withdrawn or Incomplete*, an indicator equal to one if a loan application is withdrawn by the applicant or remains incomplete, and zero otherwise. Column (2) reports results where the dependent variable is *Denial Conditional on Underwriting*, defined as an indicator equal to one if a loan application is denied after entering the underwriting process. The key independent variable is *Bank AI $\times$ Minority*, which captures whether bank AI adoption differentially affects minority applicants across decision stages. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Dependent variable	(1) Withdrawn or Incomplete	(2) Denial Conditional on Underwriting
Bank AI $\times$ Minority	0.100 (1.38)	0.417*** (3.49)
Minority	-0.001 (-0.93)	0.044*** (12.31)
Co-applicant	-0.001* (-1.85)	-0.022*** (-10.04)
Loan-to-Income	-0.002 (-1.52)	-0.019*** (-5.74)
Log Income	0.002 (0.74)	-0.049*** (-13.03)
Loan Term	0.000*** (3.06)	0.000 (0.99)
CLTV		0.000*** (3.13)
Female	0.001 (1.50)	-0.002*** (-2.73)
Bank $\times$ Year FE	YES	YES
MSA FE	YES	YES
Census Tract $\times$ Year FE	YES	YES
Debt-to-Income FE	YES	YES
Loan Type FE	YES	YES
Applicant Age Group FE	YES	YES
Adj. R^2	0.912	0.369
Observations	18,211,027	13,418,494

Table 12: Bank AI Adoption and Mortgage Origination Fees

This table presents results on the impact of bank AI adoption on mortgage origination fees. The dependent variables are *Total Loan Cost*, the sum of borrower paid origination charges and borrower paid third party charges, such as appraisal fees and title insurance fees scaled by loan amount in percentage points, in column (1), *Origination Charges*, the total fees charged by the lender for originating the loan, including discount points designated as borrower paid scaled by loan amount in percentage points, in column (2), *Discount Points*, the total amount spent on discount points scaled by loan amount in percentage points, in column (3), and *Rate Spread*, the difference between the loan's annual percentage rate (APR) and the average prime offering rate (APOR) on comparable loans, in column (4). The key independent variable is *Bank AI*×*Minority*, which is the interaction term between *Bank AI* and *Minority*. *Minority* equals one if the borrower is African American or Hispanic and zero otherwise. *Bank AI* is the share of the bank's AI skill-related job postings lagged by one year. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1) Total Loan Cost	(2) Origination Charges	(3) Discount Points	(4) Rate Spread
Bank AI × Minority	0.312 (0.65)	0.202 (0.86)	-0.033 (-0.10)	1.782** (2.25)
Minority	0.008 (0.85)	0.039*** (5.23)	0.039*** (5.58)	0.006 (0.51)
Co-applicant	0.025*** (4.04)	0.011*** (3.35)	0.011*** (5.25)	-0.036*** (-5.11)
Loan-to-Income	0.007 (0.54)	0.037*** (4.58)	0.055*** (13.16)	-0.091*** (-3.35)
Log Income	-0.395*** (-16.48)	-0.128*** (-11.37)	0.018** (2.44)	-0.041 (-1.11)
Loan Term	0.001*** (2.80)	0.001*** (3.66)	0.000*** (2.74)	-0.001*** (-2.86)
CLTV	-0.005*** (-9.21)	-0.003*** (-8.59)	-0.001*** (-5.66)	0.006*** (7.64)
Female	0.029*** (8.36)	0.020*** (8.82)	0.014*** (12.37)	0.003 (0.87)
Bank×Year FE	YES	YES	YES	YES
MSA FE	YES	YES	YES	YES
Census Tract×Year FE	YES	YES	YES	YES
Debt-to-Income FE	YES	YES	YES	YES
Loan Type FE	YES	YES	YES	YES
Loan Purpose FE	YES	YES	YES	YES
Applicant Age Group FE	YES	YES	YES	YES
Adjusted R-squared	0.571	0.431	0.365	0.416
Observations	11,123,364	11,123,364	11,123,364	10,461,449

Table 13: Bank AI Adoption and Mortgage Lending: McDash Data

This table presents results on the impact of bank AI adoption on mortgage lending using data from McDash Analytics. The dependent variable is *Interest Rate* in all columns. The key independent variable is *Bank AI*×*Minority*, which is the interaction term between *Bank AI* and *Minority*. *Minority* equals one if the borrower is African American or Hispanic and zero otherwise. *Bank AI* is the share of the bank's AI skill-related job postings lagged by one year. Column (1) presents the baseline specification. Column (2) additionally controls for loan-to-value ratio (*LTV*) and borrower credit score. Column (3) includes the loan-to-value ratio (*LTV*) and borrower credit score bucket×year fixed effects to control for time-varying differences across loan segments. Column (4) includes bucket×year-month fixed effects to control for more granular time-varying differences across loan segments. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)
Bank AI × Minority	0.845*** (3.01)	0.885*** (2.99)	0.848*** (2.89)	0.632** (2.27)
Minority	0.041*** (5.37)	0.029*** (3.23)	0.033*** (4.12)	0.005 (0.72)
Co-applicant	-0.022*** (-6.00)	-0.022*** (-5.46)	-0.025*** (-6.47)	-0.007** (-2.19)
Loan-to-Income	-0.187*** (-21.32)	-0.182*** (-22.57)	-0.182*** (-21.09)	-0.130*** (-15.73)
Log Income	0.011 (0.60)	0.003 (0.16)	0.021 (1.14)	-0.034* (-1.80)
Loan Term	0.005*** (23.95)	0.005*** (23.47)	0.005*** (23.68)	0.004*** (18.34)
CLTV	0.001*** (2.96)	-0.002*** (-2.81)	0.001** (2.52)	0.004*** (11.12)
Female	0.018*** (9.80)	0.018*** (9.01)	0.016*** (8.73)	-0.006*** (-4.77)
LTV		0.328*** (4.30)		
Credit Score		-0.001*** (-7.14)		
Bank×Year FE	YES	YES	YES	YES
MSA FE	YES	YES	YES	YES
Census Tract×Year FE	YES	YES	YES	YES
Debt-to-Income FE	YES	YES	YES	YES
Loan Type FE	YES	YES	YES	YES
Loan Purpose FE	YES	YES	YES	YES
Applicant Age Group FE	YES	YES	YES	YES
Bucket×Year FE			YES	
Bucket×Year-Month FE				YES
Observations	2,026,830	1,695,175	2,025,115	2,025,104
Adj R-squared	0.641	0.652	0.643	0.849

Table 14: Bank AI Adoption, Automation, and Minority Lending

This table examines the relationship between bank AI adoption, automation, and minority lending. The dependent variables are *Denial* in column (1) and *Interest Rate* in column (2). The key independent variables are *Bank AI*×*Minority* and *Bank Automation*×*Minority*, which capture the differential impact of bank AI adoption and automation on minority applicants. *Minority* equals one if the borrower is African American or Hispanic and zero otherwise. *Bank AI* is the share of the bank's AI skill-related job postings lagged by one year. *Bank Automation* is the share of the bank's automation-related job postings lagged by one year. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1) Denial	(2) Interest Rate
Bank AI × Minority	0.722*** (6.88)	2.507*** (2.82)
Bank Automation × Minority	-0.060*** (-3.19)	-0.083 (-1.51)
Minority	0.055*** (10.54)	0.016 (1.12)
Co-applicant	-0.018*** (-8.52)	-0.066*** (-9.07)
Loan-to-Income	-0.007*** (-2.94)	-0.173*** (-6.34)
Log Income	-0.035*** (-10.58)	0.066* (1.84)
Loan Term	0.000*** (3.09)	0.002*** (6.35)
CLTV	0.001*** (7.92)	0.003*** (3.33)
Female	0.001 (1.54)	0.005 (1.49)
MSA FE	YES	YES
Bank×Year FE	YES	YES
Census Tract×Year FE	YES	YES
Debt-to-Income FE	YES	YES
Loan Type FE	YES	YES
Loan Purpose FE	YES	YES
Applicant Age Group FE	YES	YES
Adjusted R-squared	0.370	0.576
Observations	13,918,119	11,085,289

Table 15: Resume-Based Bank AI Adoption and Minority Lending

This table examines the relationship between resume-based bank AI adoption and minority lending. The dependent variables are *Denial* in column (1) and *Interest Rate* in column (2). The key independent variable is *Bank AI (Resume) × Minority*, which captures the differential impact of resume-based bank AI adoption on minority applicants. *Minority* equals one if the borrower is African American or Hispanic and zero otherwise. *Bank AI (Resume)* is the share of banks' employees with AI skills based on resume data lagged by one year. The standard errors are clustered at the bank level, and t-statistics are reported in parentheses. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

	(1) Denial	(2) Interest Rate
Bank AI (Resume) × Minority	0.089*** (2.80)	0.250* (1.68)
Minority	0.032*** (5.94)	0.022 (0.98)
Co-applicant	0.007*** (2.61)	-0.048*** (-7.68)
Loan-to-Income	-0.000*** (-2.68)	-0.000*** (-2.80)
Log Income	-0.033*** (-12.57)	0.154*** (7.15)
Loan Term	0.000*** (3.64)	0.002*** (4.06)
CLTV	0.000*** (9.91)	-0.000*** (-9.01)
Female	-0.003** (-2.08)	0.018*** (5.10)
MSA FE	YES	YES
Bank×Year FE	YES	YES
Census Tract×Year FE	YES	YES
Debt-to-Income FE	YES	YES
Loan Type FE	YES	YES
Loan Purpose FE	YES	YES
Applicant Age Group FE	YES	YES
Adjusted R-squared	0.361	0.548
Observations	21,274,275	16,592,933

Table A1: Definitions of Variables

Variable	Definition
Bank AI	The level of bank investment in AI, measured by the share of the bank's AI skills-related job postings in a year.
Denial	An indicator equaling one if the loan is rejected and zero otherwise.
Interest Rate	The contractual interest rate of the loan.
Rate Spread	The difference between a loan's Annual Percentage Rate (APR) and the Average Prime Offer Rate (APOR) on a comparable mortgage.
Minority	An indicator equaling one if the borrower is African American or Hispanic and zero otherwise.
Co-applicant	An indicator equaling one if the loan has a co-applicant and zero otherwise.
Loan-to-Income	The ratio of the loan amount to the applicant's income.
Log Income	The natural logarithm of the applicant's income.
Loan Term	Loan maturity in months.
Loan Purpose	An indicator equaling one if the loan purpose is a home purchase and zero if it is a refinance.
Debt-to-Income	The ratio of the applicant's monthly debt obligations to monthly income.
Applicant Age	Categorical variables representing the applicant's age group: Age ≤ 25 , Age 26–35, Age 36–50, Age 51–65, and Age > 65 .
CLTV	Combined loan-to-value ratio.
Female	An indicator equaling one if the borrower is female and zero otherwise.
Withdrawn or Incomplete	An indicator equaling one if the application was withdrawn or closed for incomplete information and zero otherwise.
AI Job Number	The number of AI skills-related job postings at the bank in a given year.
AI Job Share	The proportion of the bank's job postings that require AI-related skills, calculated as AI job postings divided by total job postings in a year.