Anonymous Loan Applications and Racial Disparities*

Poorya Kabir and Tianyue Ruan†

First Draft: October 15, 2022
This Draft: February 1, 2023

Abstract

Using a unique experiment in the credit market, we find that anonymous loan applications mitigate racial disparities. When names are on applications, ethnic minorities are 10.7% less likely to receive online loan offers than otherwise identical majority applicants; anonymizing applications eliminates such disparities. After receiving online loan offers, applicants need to visit lenders in person for identity verification before loan origination. Despite that race is revealed to lenders, racial disparities in loan origination also decrease. We do not find significant racial gaps in loan performance either before or after anonymization. Further tests show that accurate statistical discrimination is unlikely to explain our results.

*We thank Sumit Agarwal, Will Shuo Liu (discussant), Karsten Müller, Jessica Pan, Ivan Png, and Arkodipta Sarkar for their insightful comments. We are grateful to seminar participants at the National University of Singapore, Fudan University, IIM-Ahmedabad, Naikai University, and the 25th EdukCircle International Convention on Business: Accounting and Finance, as well as conference participants at the Asia Pacific Online Corporate Finance Workshop for their helpful comments and discussions. We are also grateful to Mingyang Sun for his excellent research assistance. This research is supported by the Ministry of Education, Singapore, under its Social Science Research Thematic Grant (MOE2019-SSRTG-024). Kabir acknowledges financial support from the NUS Start-Up Grant No. A-0003875-00-00 and Singapore Ministry of Education AcRF Tier 1 Research Grant No. A-8000758-00-00. Ruan acknowledges financial support from the Singapore Ministry of Education AcRF Tier 1 Research Grant No. A-8000757-00-00.

†National University of Singapore; emails are poorya.kabir@nus.edu.sg and tianyue.ruan@nus.edu.sg.

1 Introduction

Racial disparities are prevalent worldwide, and reducing them is an ongoing concern (Bertrand and Duflo, 2017). Restricting the use of information that predicts race has been implemented to mitigate racial disparities in various settings. The Fair Housing Act in the US prohibits using neighborhood racial composition for lending decisions, and the exclusion of zip codes from variables permissible for insurance pricing by California's Proposition 103 (Pope and Sydnor, 2011*a*) are two notable examples.[1] One such policy that has received considerable attention from policymakers is the removal of applicant names as a source of racially identifying information (Bertrand and Duflo, 2017). However, there is limited evidence on the effectiveness of such policies in the credit market. This is increasingly important since the growing use of technology in lending (Bartlett, Morse, Stanton, and Wallace, 2022; Dobbie, Liberman, Paravisini, and Pathania, 2021; Howell, Kuchler, Snitkof, Stroebel, and Wong, 2021; Fuster, Goldsmith-Pinkham, Ramadorai, and Walther, 2022; D'acunto, Ghosh, Jain, and Rossi, 2022) enables the cost-effective implementation of anonymous applications.

In this paper, we use a unique experiment to study the effect of anonymous applications on racial disparities in the consumer credit market. We analyze online loan offers, origination, and performance using data from a leading online consumer loan comparison platform in Singapore. Consumer loans are short-term unsecured loans to individual borrowers made by licensed money lenders. The online platform sends an individual application to multiple lenders simultaneously; lenders review online applications and make initial loan offers; the individual chooses one offer online and visits the lender in person for identity verification before the final loan origination. Initially, applicant names were shown to lenders on loan applications. To protect customer privacy, the platform removed applicant names from loan applications sent to lenders from September 28, 2021.

[1] Policies such as "ban-the-box" policy that restricts employers from asking about job applicants' criminal histories (Agan and Starr, 2018) and restrictions on pre-employment credit checks (Bartik and Nelson, 2020) are also often presented as tools for reducing racial disparities.

We refer to this change in policy as anonymization.

Whether anonymous loan applications can successfully reduce racial gaps in access to credit is ambiguous. Restricting the use of information predictive of race has unclear distributional consequences: it can help (Pope and Sydnor, 2011*a*) or hurt minorities (Agan and Starr, 2018; Bartik and Nelson, 2020). In our setting, anonymous applications delay (rather than permanently remove) lender access to race; the removed information on race at the initial evaluation stage will be revealed at the subsequent stage, when applicants visit lenders in person for loan origination. On the one hand, by the time lenders originate loans, they have observed applicant race, and hence the information available to lenders at the time of origination has not changed. This could lead to no change in origination. On the other hand, the delay could reduce racial disparities in loan origination by helping minorities if rejecting is harder when a personal relationship, at the in-person stage, is formed (Love, 2011; Agan and Starr, 2018).

We find that before September 28, 2021, when names were on applications, ethnic minority applicants, including Malays, Indians, and other races, are 10% less likely to receive initial loan offers than otherwise identical Chinese applicants (the ethnic majority in Singapore). Because we observe *all* application characteristics available to lenders at the time of initial online screening, the omitted variables bias is unlikely to explain our findings. Furthermore, when the platform changed its policy to anonymize loan applications, racial disparities in offer rates disappear.

Anonymizing applications delays the revelation of information about race to the in-person stage – when applicants visit lenders. Lenders can use the newly available information on race to fully undo the effect of anonymous applications at the in-person stage. We find that the significant racial gap in the origination probability is also reduced by 8% by anonymization, despite that lenders learn about applicant race prior to origination. In other words, the reduction in the racial gap in initial offer rates is attenuated by approximately 20% in the origination stage. We interpret these findings as lenders partially undoing their prior decisions in response to the newly available information on race. Overall, we find that anonymizing applications is an effective way to reduce racial disparities in consumer loans, even if the revelation of applicant race in the process is only delayed rather than removed from the entire process.

We study differences in the default rate of minorities and Chinese before and after anonymization using data from one lender who originates 14% of the loans (due to data limitations). We find that the average default rate for minorities and Chinese is the same both before and after anonymization. This suggests that the reduced racial gap due to anonymization is not driven by an increase in lower-quality loans.

The significant racial disparities in offer rates before anonymization and the elimination of such racial disparities by anonymization imply the existence of discrimination in this market. This claim is further strengthened since we observe and control for all application characteristics available to the lender. To understand whether taste-based or statistical discrimination best describes our results, we perform two additional tests: First, we find similar racial gaps in offers across different levels of income and income-to-debt ratios, implying higher repayment ability does not reduce racial gaps. Second, following (Agan and Starr, 2018), we use a model to assess the accuracy of lenders' beliefs by race in lending decisions and find that lenders' beliefs are inaccurate. These findings suggest that accurate statistical discrimination is an unlikely explanation for our findings. Similar to most prior studies, we cannot further distinguish between inaccurate statistical discrimination and taste-based discrimination due to data limitations (Bohren, Haggag, Imas, and Pope, 2021).

Our study contributes to the large and growing literature on the racial disparities in credit markets (Bartlett, Morse, Stanton, and Wallace, 2022; Butler, Mayer, and Weston, 2022; Bhutta, Hizmo, and Ringo, 2022; Pope and Sydnor, 2011*b*). A distinguishing feature of our study is that we trace out the entire process of obtaining credit, from initial loan offers to loan origination. In consumer credit markets, initial loan offers sometimes take place before formal applications are submitted (also known as pre-approvals). For instance, potential home buyers may seek mortgage pre-approvals to facilitate their property search and only submit formal mortgage applications after they find target properties. Most other studies use formal applications as the starting point and therefore miss initial credit evaluations. However, discrimination can occur at this stage (e.g., Hanson, Hawley, Martin, and Liu (2016)), similar to the lower callbacks faced by minority applicants in the labor and rental markets. By assessing the process from initial loan offers to loan originations, we overcome a crucial data limitation of previous studies and provide a more complete assessment of racial disparities in access to credit.

The prevalence of disparities and concerns for fairness and efficiency call for effective remedies. Existing studies find that anti-discrimination enforcement policies (Butler, Mayer, and Weston, 2022) and algorithmic decision-making (Dobbie, Liberman, Paravisini, and Pathania, 2021; Howell, Kuchler, Snitkof, Stroebel, and Wong, 2021; Fuster, Goldsmith-Pinkham, Ramadorai, and Walther, 2022; D'acunto, Ghosh, Jain, and Rossi, 2022) can reduce racial disparities. We show anonymous loan applications are another effective policy. With the growing use of technology, implementing anonymous applications is increasingly cost-effective. Anonymization in our setting, as is common in labor, rental, and credit markets, delays the revelation of information about race. Our analysis also highlights that despite the eventual revelation of the withheld information on race at the in-person stage, racial disparities in loan origination also decrease.

2 Institutional settings

Using a unique experiment, we study racial disparities in the consumer lending market in Singapore. Consumer loans are loans borrowed for personal use such as medical treatment, credit card debt repayment, education, wedding, etc. These loans are uncollateralized and are short-term, with a median maturity of six months in our data. The loan repayment structure follows an equated monthly installment repayment schedule, similar to a mortgage loan. The loan offer amounts average about S$4,300 (1 S$ = 0.75 USD as of January 2021), and all-in effective annual interest rates (nominal interest rates plus processing fees) average about 99%. These features are broadly consistent with high-cost consumer lending in other economies: For instance, the typical payday loan in the US is below $300 with an effective annual rate of 400 to 1000% and a 7- to 30-day maturity and a typical consumer loan in the UK ranges from £200 to £2,000 with an average effective annual rate of 600% and a maturity from a few weeks to six months (Dobbie, Liberman, Paravisini, and Pathania, 2021). Furthermore, no ubiquitous credit scoring for loans in this market exists.

We study racial disparities in this market using information from a leading online consumer loan comparison platform in Singapore. The process for an applicant to apply for and obtain a loan is as follows: An applicant fills out a loan application on the online platform; the application is sent to multiple lenders partnering with the platform. Next, lenders decide whether to extend an initial offer and the offer terms (the offer stage). The applicant receives the initial offer(s) online, compares offers, and selects one offer online. Afterward, the applicant visits the lender in person, and upon successful further verification of personal documents, a loan agreement is signed (the origination stage). If this process is unsuccessful, the borrower can choose another initial offer from a different lender. The last step, in-person identity verification and contract signing, is required by the main regulator of money lenders in Singapore, the Ministry of Law.

The matching between applicants and lenders is not completely random. Examining how the platform matches applicants and lenders reveal no matching based on applicant race. The platform orders moneylenders and sends each application to a few at the top of the list. Depending on the outcome of the decisions, they may move on to moneylenders down the list. The ordering is manually changed by the platform staff from time to time. Crucially, the same ordering applies to all applications irrespective of their characteristics as long as the ordering is effective. Given the stable distribution of applicant race over time, minority and Chinese applications are unlikely to be systematically matched to different lenders. We conduct a formal test to assess whether the ordering of lenders is different for minority vs. Chinese applicants in Appendix OA.1 and further verify that the ordering is identical for Chinese and minority applicants.

Initially, applicant names were shown to lenders on loan applications. To protect customer privacy, the platform removed applicant names from loan applications sent to lenders from September 28, 2021. This change effectively anonymizes loan applications as applicant names contain racially identifying information (Bertrand and Mullainathan, 2004).[2]

For our analysis, we consider all non-Chinese applicants as minority applicants. The Chinese are the ethnic majority in Singapore. According to the Singapore Census of Population, as of 2020, Singapore's resident population includes 74.3% of Chinese, 13.5% of Malays, and 9.0% of Indians.

3 Data and Summary Statistics

The main data set we use in this paper contains detailed data from the leading online consumer loan comparison platform in Singapore.[3] We observe application characteristics, initial offers, loan originations, and loan performance for the period from October 2020 to January 2022. Below, we briefly describe each part of the data in more detail.

We observe detailed application characteristics, including applicant name, age, nationality, residency status, income, marital status, postal code, occupation, housing status, and existing borrowing from banks and money lenders. This list of variables fully contain the set of application characteristics the lenders observe at the time of application.

[2]Experiments of anonymous hiring procedures in several European countries (Krause, Rinne, and Zimmermann, 2012; Behaghel, Crépon, and Le Barbanchon, 2015) involve removal of address in addition to names to implement anonymization. In our context, applicant location, which remains observable to lenders, is not predictive of race due to the Singapore government's housing policy that prevents granular ethnic segregation (Agarwal, Choi, He, and Sing, 2019; Wong, 2013).

[3]This platform is active in Singapore, Hong Kong, and Australia. As of January 2022, the platform partners with 37 out of 156 licensed moneylenders in Singapore.

We also observe lenders' decisions: whether the lender approves an initial offer to the applicant, and if an initial offer is given, the offer terms, including amount, maturity, interest rate, and processing fee. These Lending decisions are done by credit officers and not by sophisticated machine learning algorithms.

Additionally, we observe whether a loan was originated and, if so, the origination terms, including amount, maturity, interest rate, and processing fee.

We observe loan performance for a subset of originated loans by one of the lenders, which accounts for 14% of the loans in the sample. The lender has a profit-sharing arrangement with the platform, providing part of the loan profits to the platform monthly. For each loan, we observe monthly payments to the platform, which allow us to measure actual repayments. We define an applicant to have late payment if the lender expects a repayment but receives zero.

Finally, to control for neighborhood characteristics, we map the location of each individual to a planning area, the main urban planning and census division in Singapore. Our sample covers 29 planning areas in total.

For our main analysis, we focus on the sub-sample of applications whose information is pre-filled directly from the Singapore government database. This filtering offers two advantages. First, the official records have higher data quality and fewer measurement errors than self-reported information. Second, applicant consent is required for this pre-filling service, which helps to screen out spam applications in a similar way that the common "captcha" verification works for many web-based services.[4]

We measure applicants' race by matching their names to races following Wong (2013). This approach is feasible as different race groups in Singapore have distinct names. In our classification, we also require consensus among at least two research assistants, who manually reviewed the names to reduce measurement error. We drop names where there is no consensus. Such a procedure yields racially distinctive names commonly used in

[4]Our main results are robust to including individuals without government-verified information.

correspondence studies.

Some individuals apply multiple times through the platform. All our analysis is done at the application level. For simplicity, however, we use applicant and application interchangeably.

Panel A of Table 1 provides summary statistics of application characteristics. There are a total of 16,281 applications during the entire sample period, 61% of them belong to minority racial groups. The average applicant is 36 years old, and 75% of the applicants are male. There are 11,789 applications submitted before September 28, 2021; lenders see the names on applications at the initial evaluation stage. The remaining 4,492 applications submitted on or after September 28, 2021 are anonymized at the initial evaluation stage. Columns (2) and (3) report the mean differences between Chinese and minority applicants before and after anonymization, respectively. Minority applicants are younger, more likely to be female, more likely to live in public housing, and have lower income than Chinese applicants. These differences remain stable over time.

Panel B of Table 1 provides summary statistics of credit outcomes. These loan offers are not collateralized and are short-term: the average offer has a maturity of 6.43 months. Average annual interest rates are 42%. The average effective interest rate, which takes into account the processing fee, is 99%.[5] Columns (2) and (3) of Panel B document unconditional racial gaps (the mean difference between Chinese and minorities) before and after anonymization, respectively. Two unconditional patterns are worth noting: 1) Minorities receive fewer offers, lower loan amounts, shorter maturity, higher annual effective interest rates, and fewer originations.[6] 2) These differences, however, become less pronounced after anonymization. The unconditional comparisons provide first-pass evidence of disparate treatment by race. From the next section, we analyze racial disparities conditional

[5]The annual effective interest rate is calculated using the monthly installment repayment schedule, taking into account the processing fee. Please, refer to Appendix OA.2 for detailed descriptions.

[6]In Appendix OA.2, we discuss the legal limit and the bunching of processing fees and interest rates as a potential reason for why we find small and insignificant differences between Chinese and minorities across these two dimensions of offer terms.

on application characteristics.

4 Empirical Strategy

To estimate the effect of anonymizing loan applications on racial disparities, we compare the minority-Chinese gap in credit outcomes before and after the inception of removing names on September 28, 2021. In our specifications, we control for all application characteristics observable to the lender when making the initial offer decisions, high-frequency time-fixed effects, and lender-fixed effects. The key identifying assumption for attributing the change in the racial disparities to anonymization is that the racial disparities stay stable absent of the change, analogous to a standard parallel trends assumption in difference-in-differences designs.

We estimate the following OLS regression in the dyadic data on loan applications and lenders:

$$y_{i,j} = \pi_t + \alpha_{j,s(t)} + \gamma_{s(t)} X_i + \beta_{pre} \times Minority_i \times pre_t + \beta_{post} \times Minority_i \times post_t + \varepsilon_{i,j} \quad (1)$$

In this specification, i denotes an application filled out at time t, and $y_{i,j}$ is a measure of credit decision/outcome of lender j for application i. $Minority_i$ is an indicator that takes a value of one for applicants that are minority and zero otherwise. pre_t and $post_t$ are indicators for applications filled out before and after September 28, 2021, respectively. β_{pre} and β_{post} reflect the racial disparities in the outcome variable in the pre and post periods, respectively. Their difference, $\Delta\beta = \beta_{post} - \beta_{pre}$, reflects the change in the racial disparities following the the use of anonymous applications and corresponds to the treatment effect of anonymous applications.

We include a host of control variables and fixed effects. We use all the information available to lenders at the time of application as control variables (X_i). In the baseline specification, we convert all continuous numerical characteristics (e.g., income) to categorical variables using their quintiles to allow for non-linear effects in control variables capturing the potential non-linearity in the lending model. We also allow the retention of missing values this way.[7] We allow the effects of these control variables γ to differ in the pre and post-periods (hence the $s(t)$ subscript). Standard errors are clustered at the lender-month level to allow for correlated decision makings across applications by a lender in a month. Year-month fixed effects π_t are included to absorb time-series fluctuation in aggregate credit conditions and the average impact of all other concurrent aggregate factors. $\alpha_{j,s(t)}$ for $s(t) \in \{pre, post\}$ are lender fixed effects separately for the pre and post-periods. By including this set of fixed effects, we absorb lender-specific practices that can differ in the pre and post-periods.

A common critique of running a regression similar to equation (1) on observational data is the omitted variables bias; namely, relevant covariates for lending outcomes are unobserved by the researchers, and the inability to include these covariates leads to biases in the coefficient estimates of the included covariates. In our setting, the applications and the decisions of initial offer approvals/rejections are completely online. We observe all the information available to lenders at the time of their decisions on initial offers and include them in the regression. Therefore, our analysis is unlikely to suffer from the omitted variables bias.

5 Results

5.A Probability of receiving initial offers

Table 2 shows the estimated effects of anonymous applications on the probability of receiving initial offers. In Column 1, we estimate equation (1) in the dyadic data on loan

[7] The list of control variables includes: the age of the applicant, applied amount, applied loan tenure, length of stay in current residence, loan purpose, marital status, housing type (e.g., public housing, condominium apartment, etc.), housing status (e.g., rented, owned-mortgaged, etc.), job title, job industry, current employment duration, previous employment duration, whether the applicant owns a property, monthly income, remaining bank-loan balance, remaining moneylender-loan balance, monthly bank-loan payment, monthly moneylender loan payment, and the planning area.

applications and lenders where the left-hand side variable is an offer dummy that takes the value of one if a lender extends an offer to the borrower and zero otherwise, multiplied by 100. We include lender-fixed effects separately for the pre- and post-periods to absorb lender-specific practices that are allowed to differ in the pre- and post-periods. We also convert all continuous numerical characteristics (e.g., income) to categorical variables using their sample quintiles to allow for potential non-linearity and allow the impacts of these observable characteristics to differ flexibly in the pre- and post-periods. We find that a coefficient on the interaction term between $Minority_i$ and Pre_t of -3.81, implying that in the pre-period, minority applicants are 3.81 percent points less likely to receive initial offers than the otherwise observably identical Chinese applicants. The racial disparity is highly significant and amounts to 10% of the average offer probability. In the post period, however, the racial gap disappears as seen in a statistically insignificant coefficient on the interaction term between $Minority_i$ and $post_t$ of 0.238. The treatment effect of the anonymization change, reflected by $\Delta\beta = \beta_{post} - \beta_{pre} = 4.048$, is highly statistically significant with a p-value less than 0.0001. It is also economically sizable: this effect amounts to 10.6% of the sample average offer probability.

In Column 2, we use an alternative way to include control variables where we impute zero for missing values, add 1 to zero values, and then log-transform all continuous numerical variables. We maintain the inclusion of Lender×Post fixed effects and the flexibility that the impacts of observable characteristics on the outcome variable can differ in the pre- and post-periods. We find similar estimates as in Column 1.

In Column 3, we aggregate the dyadic sample to the application level and examine how anonymization affects the average offer rate analogously. To match this level of aggregation, we now include the Post indicator, as opposed to Lender×Post fixed effects, and cluster standard errors at the month level. Albeit different aggregations, the estimates remain similar and show economically and statistically significant racial gaps. In other words, the racial disparities are not driven by particular lenders or the matching between applicants and lenders.

One advantage of our setting is that initial offers are extended fully online without any in-person interaction. Hence, there are no application characteristics that lenders can observe but are unobservable to us. In other words, the omitted variable bias that often hampers the usefulness of action-based tests of discrimination in observational data is less of an issue for our setting.

We also study the dynamic patterns of racial disparities in offer rates using the following event study specification:

$$y_{i,j} = \pi_t + \alpha_{j,s(t)} + \gamma_{s(t)} X_i + \textstyle\sum_{s \neq 0} \beta_s \times Minority_i \times \mathbb{1}_s + \varepsilon_{i,j} \tag{2}$$

In this specification, $\mathbb{1}_s$ indicates the timing of application i relative to month 0, the implementation of anonymous applications. We set month 0 as the omitted baseline period, motivated by the zero average racial gap in the post-anonymization months as estimated in Table 2. The coefficient β_s reflects the racial disparity in the initial offer probability in month s. The coefficients for the pre-anonymization months $s < 0$ allow us to test the key identifying assumption of our research design, parallel trends. If our research design is valid, we expect statistically significant and stable racial gaps in pre-anonymization months. Figure 1 plots the entire path of coefficients β_s along with their associated 95% confidence intervals as estimated from equation (2). For each of the four months prior to the anonymization practice, there a statistically significant racial gap; the magnitude of the racial gap stays stable at its average level of 3.81 percent points and is also similar to the level seen in the previous months. This pattern validates the key identifying assumption, i.e., absent anonymization, the racial gap would have remained constant. For the two months following anonymization, we see insignificant coefficients, implying that the racial gap is eliminated.

5.B Heterogeneous racial gaps across income groups

In this subsection, we examine whether the racial gaps in offer rates reflect differences in repayment probability. We consider both the level of income and the income-to-debt ratios as proxies for repayment probability.

We first group applicants into four quartiles of their annual income and separately estimate equation (1) for each of the four quartile groups. We plot the coefficient on the interaction term between $Minority_i$ and Pre_t and their 95% confidence intervals for the four quartile groups in Panel A of Figure 2. We find similar racial disparities across income groups.

Alternatively, we group applicants into four quartiles of their annual income by the applied amount and separately estimate equation (1) for each of the four quartile groups. We plot the coefficient on the interaction term between $Minority_i$ and Pre_t and their 95% confidence intervals for the four quartile groups in Panel B of Figure 2. As in the previous split by income level, we find similar racial disparities across groups with different income-to-debt ratios.

5.C Loan origination and loan performance

Column 1 of Table 3 shows the effect of anonymous applications on disparities in loan origination. The estimates reveal that minority applicants are significantly less likely to receive loan origination than Chinese applicants before anonymization, but such disparity becomes insignificant once applications are anonymized. Comparing the economic magnitude of the treatment effect $\Delta\beta$ for the two outcome variables–initial offer and loan origination–sheds light on whether lenders fully adjust lending in the in-person verification stage. In the post-period, when names are removed from loan applications, lenders do not know the racial identity of the applicants at the time of initial evaluation. The racial identity is revealed to them once applicants visit them to fulfill in-person verification procedures. Lenders can use this information to fully undo their initial decisions without information on race. We find that anonymization is associated with a 10.7% decrease in racial disparities in offer rate (Column 1 of Table 2). The corresponding decrease in racial disparities in origination rate is 8.1% of its sample average according to the estimates from Column 1 of Table 3. In other words, lenders partially adjust their behaviors at in-person visits, and therefore, the reduction of racial disparities in the initial offer stage is attenuated by approximately 20% once the applicants advance to the origination stage. We obtain a similar magnitude in the application-level analysis: comparing the economic magnitude from the estimates in Column 3 of Table 2 and Column 2 of Table 3 reveals an approximately 15% of lenders' partial adjustment in this level of analysis.

Column 3 of Table 3 reports the results for conversion rate, measured as the origination indicator divided by the number of offers. While in the pre-period, minority applicants have a 0.232 percentage points higher conversion rate than the otherwise identical Chinese applicants, they have a 0.076 percentage points lower conversion rate in the post period. Although both estimates are insignificant, the economic magnitude of the treatment effect $\Delta\beta$ corresponds to an 8.33% reduction relative to the sample mean of the conversion rate, reflecting economically meaningful effect size. The reduction in conversion rate provides another piece of evidence for lenders' partial adjustment in the in-person verification stage.

Finally, we examine the relationship between race and loan performance for a sub-sample of originated loans. The data comes from one lender that originates approximately 14% of the loans. Column 4 of Table 3 reports the regression analyses in the loan performance sub-sample. As the sub-sample comes from one lender, the usual Lender$\times$Post fixed effects collapse to the Post indicator. Also, we can only allow the effects of the included control variables to be the same in the pre- and post-periods due to the small sample size. Column 4 shows that the average likelihood of delinquency is lower for minority borrowers in the pre-period, although statistically insignificant; in the post-period, the average difference in loan performance between Chinese and minority applicants remains statistically insignificant from zero.[8] Overall, Chinese and minority borrowers have similar delinquency levels both before and after anonymization.

5.D Additional tests

If all offers given to a borrower are the same, then the offer rate may not be the "right" outcome variable since what matters is whether an individual has an offer or not. Using the offer rate is akin to using call back rate for job applications, which are widely used in correspondence studies. However, in Appendix OA.3 and OA.4, we show two pieces of evidence suggesting the use of the offer rate as an outcome variable is justified. First, in Table OA.1, we document substantial within-application variations in the offer terms. For instance, the average ratio of the maximum offer amount to the minimum offer amount (given to the same application) is 8.71. The large variation in offer terms suggests that not all lenders assess an applicant similarly, and hence more offers are valuable.

Second, even if there is variation in offer terms, it could be that the "best" offer is the same for minorities and Chinese. While it is difficult to know which offer is the best for each applicant, we use three different definitions of the best offer and repeat our main analysis using an applicant's best offer. We use the maximum offer amount and maximum offer maturity[9], and minimum effective interest rate across all lenders. Studying these outcome variables can be justified if applicants have lexicographic preferences over the outcome.

Table OA.2 shows the effect of anonymous applications on disparities in the best initial offer terms; namely, the maximum offer amount, maximum offer maturity, and the minimum annual effective interest rate. For all three aspects, we find that the racial disparities in the pre-period, summarized by the β_{pre} coefficient, are higher in absolute terms than the racial disparities in the post-period, summarized by the β_{post} coefficient. The dif-

[8] In untabulated analyses, we confirm that the coverage in this performance sub-sample is balanced.

[9] In unreported analysis, we find that applicants are more likely to choose offers with longer maturity holding everything else constant. This could be driven by a preference for lower monthly payments. Hence, we interpret longer maturity as a "better" offer.

ference $\Delta\beta$ is statistically significant, as well.

We also assess the robustness of our results to alternative samples and specifications in Appendix OA.6. We obtain similar estimates when we relax the sample filtering and when we estimate a more flexible specification to accommodate for lender-specific changes in lending practices as in our baseline results.

If other application characteristics can not predict race, racial gaps will be eliminated simply because lenders can't adjust their offers to reflect race. In Appendix OA.5, we study the predictability of race and find that race is highly predictable: machine learning models achieve an out-of-sample classification accuracy exceeding 90%. This is similar to findings of Fuster, Goldsmith-Pinkham, Ramadorai, and Walther (2022), who show that race is also highly predictable in the US mortgage market. Hence, eliminating racial disparities after anonymization is not simply driven by race not being predictable.

The analysis so far masks important heterogeneity across lenders. More discriminating lenders (those with higher racial gaps before anonymization) might have more incentives to predict race when extending offers post-anonymization. Hence, it is unclear whether anonymization reduces disparities for the more discriminating lenders. In Appendix OA.7, we study whether such heterogeneity exists and which lenders are more affected by anonymous applications. We estimate lender-specific racial disparities in a specification similar to equation (1). In Panel (A) of Figure OA.4, we plot lender-specific $\beta_{pre,i}$ and $\Delta\beta_i$ for all lenders in the sample. We observe that lenders with higher racial gaps in pre-period (lower β_{pre}) have higher $\Delta\beta$. This finding suggests that anonymization reduces disparities more for lenders with higher pre-anonymization racial gaps. We see a similar pattern for loan origination in Panel (B) of Figure OA.4.

6 Discussion

Our findings strongly suggest the existence of discrimination in the consumer loan market: Before anonymization, minority applicants receive fewer loan offers than otherwise identical Chinese applicants. Controlling for all characteristics observable by lenders mitigates the common concern for the omitted variables bias. In addition, once loan applications are anonymized, the racial disparities in offer rates disappear, further corroborating the existence of discrimination.

Which theory of discrimination best describes our findings? Economists differentiate between taste-based and statistical discrimination. Under taste-based discrimination (Becker, 1957), differential treatment stems from the disutility from providing service to or interacting with members of a particular group. Under statistical discrimination (Phelps, 1972; Arrow, 1973; Aigner and Cain, 1977), differential treatment stems from imperfect information and the use of group membership as a signal of unobserved information. Empirically differentiating the source of discrimination is challenging (Bertrand and Duflo, 2017). The differentiation becomes even more challenging if we allow for the possibility of inaccurate beliefs (Bohren, Haggag, Imas, and Pope, 2021); inaccurate statistical discrimination cannot be differentiated from taste-based discrimination due to data limitations in most settings.

Two pieces of evidence reveal that accurate statistical discrimination cannot explain our results: First, we find similar racial gaps for different quantiles of income and income-to-debt ratios, suggesting that repayment ability does not affect racial disparities. Second, we use a simple model of lender decision-making under accurate statistical discrimination, following Agan and Starr (2018) outlined in Appendix OA.8. Using the model, we can estimate the lender's belief about the probability of someone with a monthly salary in the top quintile belonging to the minority group. We repeat this exercise not only for different levels of income but all application characteristics and compare them with empirical probabilities from the data. We find significant discrepancies between the inferred and empirical probabilities and conclude that lenders' priors are inaccurate. In Appendix OA.8, we test for a specific form of inaccurate beliefs, stereotypes (Bordalo, Coffman, Gennaioli, and Shleifer, 2016), and find that the data does not support such interpretation. We cannot further distinguish between inaccurate statistical discrimination and taste-based discrimination, as doing so requires additional data that is unavailable in this setting (Bohren, Haggag, Imas, and Pope, 2021).

Cornell and Welch (1996), Fisman, Paravisini, and Vig (2017), and D'acunto, Ghosh, Jain, and Rossi (2022) document in-group preferences as an explanation for racial disparities in the lending market. A test of in-group preferences requires variation in lender race. We obtained data on the names of shareholders and authorized officers/representatives of all moneylenders in our sample from Singapore's business registration records. We find that the shareholders and authorized officers/representatives of all but one lender in our sample have Chinese-sounding names. Hence, our data do not have sufficient variation for testing in-group preferences.

7 Conclusion

We study the effect of anonymous loan applications on racial discrimination in the consumer loan market. Initially, with names on loan applications, minority applicants are significantly less likely to receive initial loan offers than otherwise identical Chinese applicants; a system-wide implementation of anonymous loan applications eliminates such disparities. Heterogeneity analyses and analysis of lender beliefs show that our results are inconsistent with accurate statistical discrimination. Thus, inaccurate beliefs or taste-based discrimination are the remaining plausible source of discrimination.

Restricting the use of information correlated with race has been implemented to mitigate racial disparities in various settings. Our study provides evidence of the effectiveness of anonymous evaluation practices for reducing racial disparities. In our setting, as is common in labor, rental, and credit markets, only the first stage of the process is made anonymous. We find that reduced racial disparities in the loan origination stage, as well, although lenders observe applicant race through in-person interaction with applicants. Overall, we find that anonymous applications are effective in reducing racial disparities in access to credit, by delaying the revelation of race.

With the advent and expansion of fintech lenders, the implementation of anonymous applications becomes increasingly cost-effective and feasible. Online credit platforms are prevalent and growing across the world. Serving as an intermediary between borrowers and lenders, these platforms can credibly verify customers and anonymize applications simultaneously. This can potentially increase the allocation of credit to minority applicants. Our quasi-experimental evidence of the benefits of anonymization based in Singapore likely provides a lower bound for other countries as Singapore government has implemented successful policies in promoting racial equity (Agarwal, Choi, He, and Sing, 2019; Wong, 2013). Implementing anonymous loan applications in a country such as the US will likely deliver larger gains to minority borrowers.[10] In addition, with the growth of online-based loan origination (Buchak, Matvos, Piskorski, and Seru, 2018), complete anonymization may become feasible for fintech consumer credit and can help minority borrowers even more. In this context, mandatory information collection of race in regulations designed to detect discriminatory lending practices, such as the Home Mortgage Disclosure Act (HMDA) in the US, may constrain the effective implementation of anonymization.

[10] In a survey by US News, Singapore ranked 13th in the racial equity index in 2022 out of 85 countries and the US ranked 65th. Source: `https://www.usnews.com/news/best-countries/best-countries-for-racial-equality`.

References

Agan, Amanda and Sonja Starr. 2018. "Ban the box, criminal records, and racial discrimination: A field experiment." *Quarterly Journal of Economics*, 133(1): 191–235.

Agarwal, Sumit, Hyun Soo Choi, Jia He, and Tien Foo Sing. 2019. "Matching in housing markets: The role of ethnic social networks." *Review of Financial Studies*, 32(10): 3958–4004.

Aigner, Dennis J. and Glen G. Cain. 1977. "Statistical theories of discrimination in labor markets." *Industrial and Labor Relations Review*, 30(2): 175–187.

Arrow, Kenneth J. 1973. "The theory of discrimination." In *Racial Discrimination in Economic Life*. , ed. Orley Ashenfelter and Albert Rees, 3–33. Princeton University Press.

Bartik, Alexander W and Scott T Nelson. 2020. "Deleting a signal: Evidence from pre-employment credit checks."

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace. 2022. "Consumer-lending discrimination in the FinTech Era." *Journal of Financial Economics*, 143(1): 30–56.

Becker, Gary S. 1957. *The economics of discrimination.* University of Chicago press.

Behaghel, Luc, Bruno Crépon, and Thomas Le Barbanchon. 2015. "Unintended effects of anonymous résumés." *American Economic Journal: Applied Economics*, 7(3): 1–27.

Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri. 2020. "On the rise of FinTechs: Credit scoring using digital footprints." *Review of Financial Studies*, 33(7): 2845–2897.

Bertrand, Marianne and Esther Duflo. 2017. "Field experiments on discrimination." *Handbook of Economic Field Experiments*, 1: 309–393.

Bertrand, Marianne and Sendhil Mullainathan. 2004. "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination." *American Economic Review*, 94(1996): 193–204.

Bhutta, Neil, Aurel Hizmo, and Daniel Ringo. 2022. "How much does racial bias affect mortgage lending? Evidence from human and algorithmic credit decisions."

Bohren, J. Aislinn, Kareem Haggag, Alex Imas, and Devin G. Pope. 2021. "Inaccurate statistical discrimination: An identification problem."

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. 2016. "Stereotypes." *Quarterly Journal of Economics*, 131(4): 1753–1794.

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru. 2018. "Fintech, Regulatory Arbitrage, and the Rise of Shadow Banks." *Journal of Financial Economics*, 130.

Butler, Alexander W., Erik J. Mayer, and James P. Weston. 2022. "Racial disparities in the auto loan market." *Review of Financial Studies*, 36(1): 1–41.

Cornell, Bradford and Ivo Welch. 1996. "Culture, information, and screening discrimination." *Journal of Political Economy*, 104(3): 542–571.

D'acunto, Francesco, Pulak Ghosh, Rajiv Jain, and Alberto G. Rossi. 2022. "How costly are cultural biases?"

Dobbie, Will, Andres Liberman, Daniel Paravisini, and Vikram Pathania. 2021. "Measuring bias in consumer lending." *Review of Economic Studies*, 88(6).

Fisman, Raymond, Daniel Paravisini, and Vikrant Vig. 2017. "Cultural proximity and loan outcomes." *American Economic Review*, 107(2): 457–492.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther. 2022. "Predictably unequal? The effects of machine learning on credit markets." *Journal of Finance*, 77(1): 5 – 47.

Hanson, Andrew, Zackary Hawley, Hal Martin, and Bo Liu. 2016. "Discrimination in mortgage lending: Evidence from a correspondence experiment."

Howell, Sabrina T, Theresa Kuchler, David Snitkof, Johannes Stroebel, and Jun Wong. 2021. "Racial disparities in access to small business credit: Evidence from the Paycheck Protection Program." *SSRN Electronic Journal*.

Iyer, Rajkamal, Asim Ijaz Khwaja, Erzo F.P. Luttmer, and Kelly Shue. 2016. "Screening peers softly: Inferring the quality of small borrowers." *Management Science*, 62(6): 1554–1577.

Kline, Patrick, Evan K. Rose, and Christopher R. Walters. 2022. "Systemic discrimination among large U.S. employers." *Quarterly Journal of Economics*, 137(4): 1–74.

Krause, Annabelle, Ulf Rinne, and Klaus F. Zimmermann. 2012. "Anonymous job applications of fresh Ph.D. economists." *Economics Letters*, 117(2): 441–444.

Love, Margaret Colgate. 2011. "Paying their debt to society: Forgiveness, redemption, and the Uniform Collateral Consequences of Conviction Act." *Howard Law Journal*, 54(3): 753–795.

Phelps, Edmund S. 1972. "The Statistical Theory of Racism and Sexism." *American Economic Review*, 62(4): 659–661.

Pope, Devin G. and Justin R. Sydnor. 2011*a*. "Implementing anti-discrimination policies in statistical profiling models." *American Economic Journal: Economic Policy*, 3(3): 206–231.

Pope, Devin G. and Justin R Sydnor. 2011*b*. "What's in a picture? Evidence of discrimination from prosper.com." *Journal of Human Resources*, 46(1): 53–92.

Wolpert, David H. 1992. "Stacked generalization." *Neural Networks*, 5(2): 241–259.

Wong, Maisy. 2013. "Estimating ethnic preferences using ethnic housing quotas in Singapore." *Review of Economic Studies*, 80(3): 1178–1214.

Figure 1: **Estimated dynamic response of racial disparities in offer rate**

This figure plots the entire path of coefficients β_s along with their associated 95% confidence intervals of the racial gap in offer rate as estimated from equation (2). In this specification, we set month 0, the implementation of anonymous applications, as the omitted baseline period, motivated by the zero average racial gap in the post-anonymization months as estimated in Table 2. The *x-axis* denotes the months before and after anonymization; the *y-axis* shows the change in the racial gap in offer rate relative to the omitted baseline period (in percentage points).

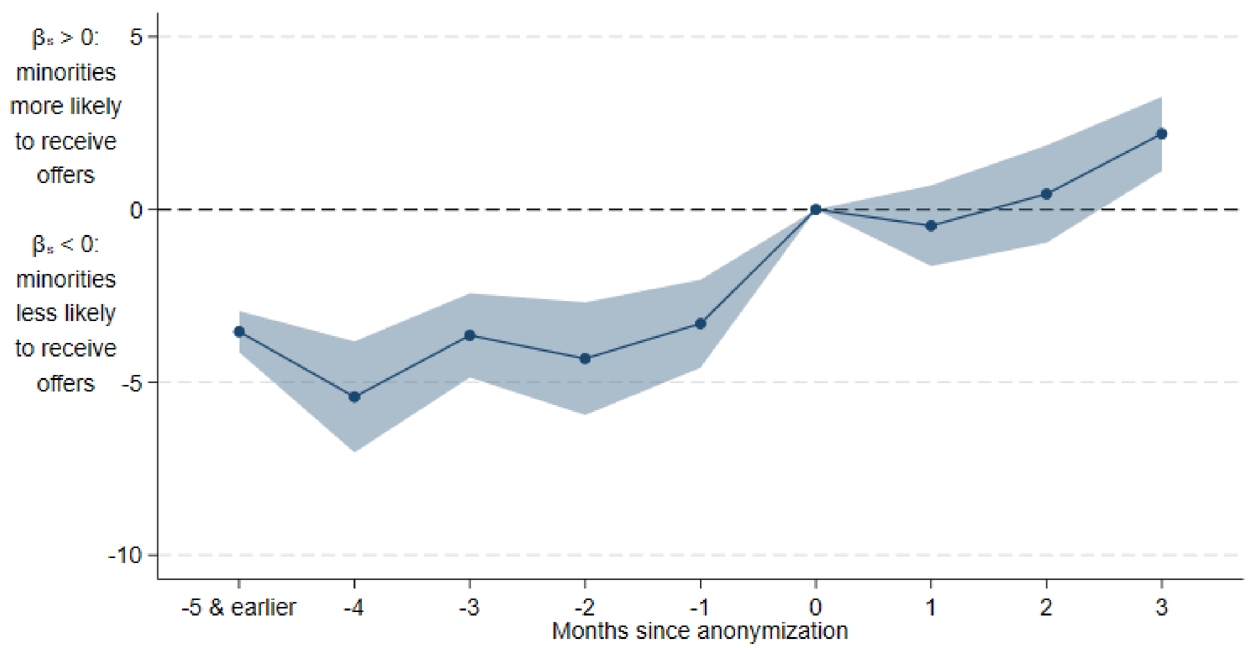

Figure 2: **Heterogeneous racial gaps across income groups**

Panel (A) figure shows the heterogeneous racial disparities by income. We group applicants into four quartiles of their annual income and separately estimate equation (1) for each of the four quartile groups. We plot the coefficient on the interaction term between $Minority_i$ and Pre_t and their 95% confidence intervals for the four quartile groups in this figure. Panel (B) figure shows the heterogeneous racial disparities by income to debt ratio. We group applicants into four quartiles of their annual income by the applied amount and separately estimate equation (1) for each of the four quartile groups. We plot the coefficient on the interaction term between $Minority_i$ and Pre_t and their 95% confidence intervals for the four quartile groups in this figure.

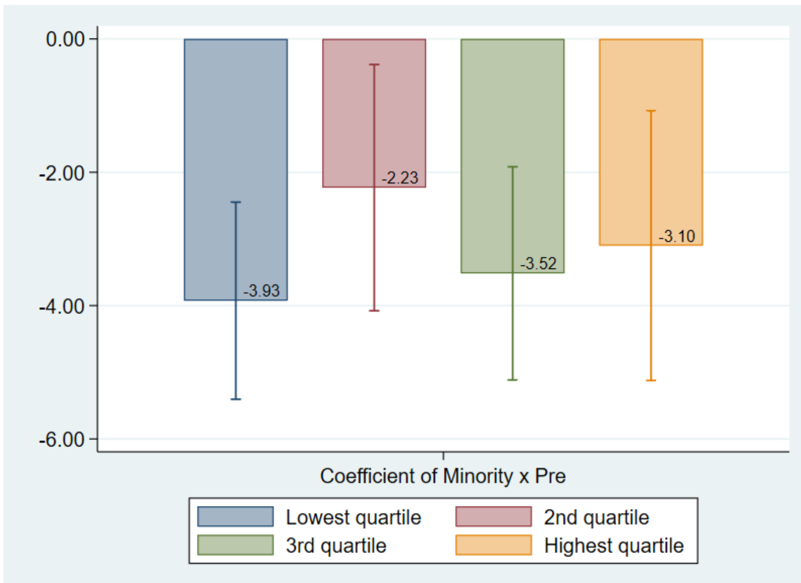

(A) Heterogeneous racial disparities by income

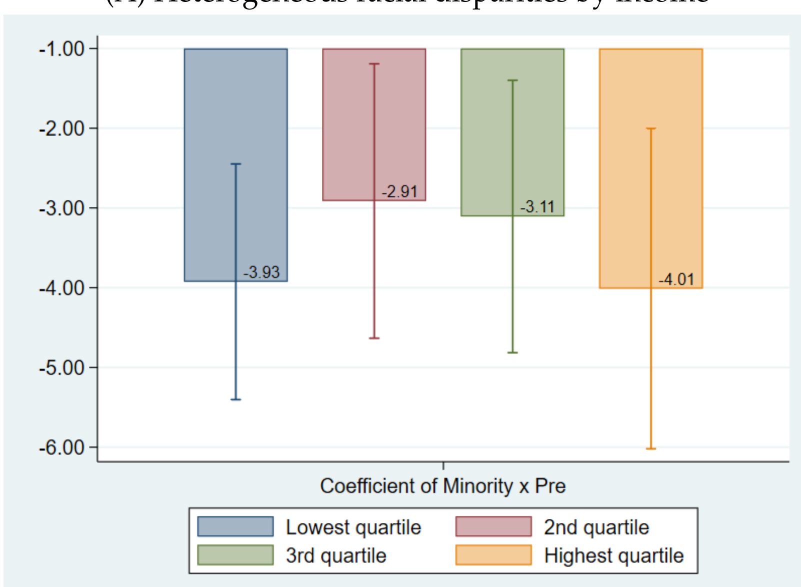

(B) Heterogeneous racial disparities by income-to-debt

Table 1: **Summary statistics of applications and credit outcomes**

This table reports the summary statistics for application characteristics in Panel A and credit outcomes in Panel B. Column 2 of both panels report the mean differences between Chinese and minority applicants in the pre-period, i.e., the period before September 28, 2021, when applicant names were visible to lenders at the initial evaluation stage (non-anonymous applications). Column 2 of both panels report the mean differences in the post-period, i.e., the period after September 28, 2021, when applications are anonymous. The monetary amount is in the local currency Singapore Dollar (SGD), and 1 SGD = 0.75 USD as of January 2021.

Panel A: Application characteristics

	Overall	$\mu_{MIN} - \mu_{CHN}$ (Pre)	$\mu_{MIN} - \mu_{CHN}$ (Post)
Age	35.65	-1.06***	-1.01***
	[9.46]	(0.18)	(0.30)
Female	0.25	0.06***	0.12***
	[0.43]	(0.01)	(0.01)
Living in public housing	0.89	0.04***	0.04***
	[0.31]	(0.01)	(0.01)
Annual income (SGD)	35,974.42	-8,818.68***	-7,185.09***
	[46,533.08]	(895.42)	(1,278.94)
Number of applications	16,281	11,789	4,492

Panel B: Credit outcomes

	Overall	$\mu_{MIN} - \mu_{CHN}$ (Pre)	$\mu_{MIN} - \mu_{CHN}$ (Post)
Average offer rate (%)	43.48	-4.64***	-2.30*
	[30.68]	(0.58)	(0.92)
Number of offers	7.57	-0.72***	-0.18
	[4.69]	(0.09)	(0.12)
Average offer amount (SGD)	4,290.71	-931.15***	-801.60***
	[3,160.12]	(63.01)	(103.88)
Average maturity (months)	6.39	-0.51***	-0.26**
	[2.74]	(0.05)	(0.09)
Average annual nominal interest rate (%)	42.44	-0.00	-0.11
	[4.82]	(0.09)	(0.16)
Average processing fee (%)	9.25	0.02	-0.02
	[0.69]	(0.01)	(0.02)
Average annual effective interest rate (%)	99.02	3.98***	1.42*
	[27.20]	(0.54)	(0.71)
Origination rate (%)	16.79	-1.96**	-1.44
	[37.38]	(0.71)	(1.17)
Number of applications	16,281	11,789	4,492

Table 2: **The effect of anonymous applications on disparities in offer rates**

This table shows the effect of anonymous applications on disparities in offer rates (equation (1)). For each regression, we also report $\Delta\beta = \beta_{post} - \beta_{pre}$, its t-statistic, p-value, and its value divided by the mean of the dependent variable. Fixed effects are included and denoted at the bottom. We include all the information available to lenders at the time of initial screening as control variables and allow the effects of the control variables to differ in the pre- and post-periods. In the baseline specification in Columns 1 & 3, we convert all continuous numerical characteristics (e.g., income) to categorical variables using their quintiles to allow for non-linear effects in control variables and for retention of missing values. In the alternative specification in Column 2, we impute zero for missing values, add one to zero values, and then log-transform all continuous numerical variables. Standard errors are clustered at the lender-month level for the application-lender level analysis and at the month level for the application-level analysis; the corresponding t-statistics are reported in brackets. We use ***, ** and * to denote significance at 1%, 5% and 10% level (two-sided), respectively.

	Application-lender level offer indicator ($\times$ 100)		Application level average offer rate (%)
	(1)	(2)	(3)
	Baseline controls	Alternative controls	Baseline controls
Minority $\times$ Pre	-3.810***	-3.096***	-3.969***
	[-16.78]	[-14.08]	[-7.93]
Minority $\times$ Post	0.238	0.408	-0.434
	[0.88]	[1.48]	[-0.58]
$\Delta\beta$	4.048	3.504	3.535
t-stat of $\Delta\beta$	11.46	9.970	3.910
p-value of $\Delta\beta$	1.80e-26	4.91e-21	0.00139
$\Delta\beta$ / Mean DV	0.106	0.0918	0.0813
Year-Month FEs	Yes	Yes	Yes
Lender $\times$ Post FEs	Yes	Yes	
Post FE			Yes
Observable controls	Yes	Yes	Yes
R^2	0.305	0.291	0.569
No. of observations	322,847	322,847	16,281

Table 3: **The effect of anonymous applications on disparities in other credit outcomes**

This table shows the effect of anonymous applications on disparities in other credit outcomes (equation (1)). For each regression, we also report $\Delta\beta = \beta_{post} - \beta_{pre}$, its t-statistic, p-value, and its value divided by the mean of the dependent variable. Fixed effects are included and denoted at the bottom. We include all the information available to lenders at the time of initial screening as control variables and allow the effects of the control variables to differ in the pre- and post-periods (except in Column 4, where we can only allow the effects of the included control variables to be the same in the pre- and post-periods due to a smaller sample size). In choosing the function form of the included control variables, we convert all continuous numerical characteristics (e.g., income) to categorical variables using their quintiles to allow for non-linear effects in control variables and for retention of missing values. Standard errors are clustered at the lender-month level for the application-lender level analysis and at the month level for the application-level analysis; the corresponding t-statistics are reported in brackets. We use ***, ** and * to denote significance at 1%, 5% and 10% level (two-sided), respectively.

	(1) Application-lender level origination indicator (× 100)	(2) Application level origination indicator (× 100)	(3) Application level conversion rate (%)	(4) Delinquency indicator (× 100)
Minority × Pre	-0.0910*	-1.598*	0.232	0.151
	[-1.68]	[-1.89]	[0.97]	[0.03]
Minority × Post	-0.0236	-0.500	-0.0764	-0.319
	[-0.29]	[-0.33]	[-0.14]	[-0.04]
$\Delta\beta$	0.0674	1.098	-0.309	-0.471
t-stat of $\Delta\beta$	0.692	0.632	-0.517	-0.0413
p-value of $\Delta\beta$	0.490	0.537	0.613	0.968
$\Delta\beta$ / Mean DV	0.0797	0.0654	-0.0833	-0.0200
Year-Month FEs	Yes	Yes	Yes	Yes
Lender × Post FEs	Yes			
Post FE		Yes	Yes	Yes
Observable controls	Yes	Yes	Yes	Yes
R^2	0.00792	0.0677	0.0806	0.403
No. of observations	322,847	16,281	14,991	373

Online appendix

This appendix contains supplementary material, tables, and figures.

OA.1 Matching between applications and lenders

Does the platform match applications to lenders based on application characteristics and, more specifically, race? Based on our communication with the platform staff, they have a pre-determined ordering of moneylenders and send out applications to moneylenders based on this ordering. The ordering is changed from time to time by the platform. Crucially, the same ordering applies to all applications irrespective of their characteristics as long as a specific ordering is still effective. Nevertheless, we formally test for the possibility of matching between applications and lenders based on applicant race in this Appendix. More specifically, we study whether Chinese and minority applicants are matched to different lenders.

Figure OA.1 shows the estimated coefficients and the 95% confidence intervals of the regressions of lender rank on the minority status across all lenders. For each application, the rank of a lender is an integer (starting from 1) that corresponds to the order in which the application is sent to a given lender. We regress lender rank on minority status for a given lender at a time and repeat this exercise for all lenders.[11] Each colored coefficient and the associated confidence interval correspond to one lender. 31 out of 36 coefficients are statistically insignificant at 5%, and the other 5 coefficients are statistically significant at 5%. Furthermore, the estimated magnitudes are economically small. On average, the absolute value of the estimated coefficient is 0.16. Moneylender rank can be any number between 1 to 36. Hence, even if the coefficients were statistically significant, their corresponding change in money lender order would be less than 0.19 in moneylender ordering. Hence, the evidence suggests that the platform does not match applications to lenders based on application race.

OA.2 Initial offer terms: characteristics and results

Each initial loan offer is characterized by four aspects, offer amount in Singapore Dollars (1 SGD = 0.75 USD as of January 2021), offer maturity in months, annual interest rate, and processing fee as a percentage of the offer amount. While Table 1 shows substantial differences in the average amount and maturity of offers received by Chinese and minority applicants, there seem to be no discernible differences in the average annual interest rate

[11] One lender has a constant lender rank at all times during its partnership with the platform and hence is dropped out from this analysis. Therefore, we have 36 lenders in total.

and processing fees between the two groups. To see why we plot the distribution of these two variables.

Panel (A) of Figure OA.2 plots the empirical cumulative distribution function (CDF) of annual interest rates for all initial offers. The figure shows that there is clear bunching near the legal limit of 48%. 46.42% and 5.38% of initial offers have an interest rate of 47% and 48%, respectively. The legal limit is set by the Ministry of Law in Singapore, and effective 1 October 2015, is 4% per month. The legal limit encompasses all forms of lending, whether collateralized or not, and to all individuals, irrespective of their income.

Panel (B) of Figure OA.2 plots the empirical cumulative distribution function (CDF) of the processing fee as a percentage of the loan offer amount. Even if money lenders can not use interest rates to fully adjust for applicants' "true" underlying risk, they can use processing fees. The figure shows a bunching of observations near the legal limit of 10%. 83.04% of initial offers have a processing fee of exactly 10%. The legal limit is set by the Ministry of Law in Singapore, and effective 1 October 2015, is 10% of the loan offer amount.

Taken together, the panels of Figures OA.2 suggest that money lenders set the interest rate and processing fee equal to the maximum legal limit quite often. Hence, these variables might have been different absent the legal limits. Consequently, comparing these offer terms might not be as informative as other terms where moneylenders have full discretion.

To analyze the borrowing cost, we calculate the effective interest rate. The existence of a processing fee implies that the annual interest rate measures the borrowing cost incompletely. Below, we describe how we calculate effective interest rates in more detail.

The structure of a loan offer follows an equated monthly installment repayment schedule, similar to a mortgage loan. Specifically, if the loan amount is B, the nominal monthly interest rate is i, and the number of months to maturity is N, the monthly payment P is such that the present value of the monthly payments at the monthly interest rate i equals to B.

$$B = \sum_{t=1}^{T} \frac{P}{(1+i)^t} \tag{OA.1}$$

With processing fee f, the applicant receives $B \times (1-f)$ as opposed to B upon loan origination. Therefore, the monthly effective interest rate r is determined by:

$$B \times (1-f) = \sum_{t=1}^{T} \frac{P}{(1+r)^t} \tag{OA.2}$$

We then annualize the monthly effective interest rate by multiplying it by 12 to obtain the annual effective interest rate. To illustrate how the processing fee affects the borrowing cost, consider a typical "zero-interest-rate" loan offer with a maturity of 1 month, a nominal interest rate of 0%, and a processing fee of 10%, which account for approximately 5% of our sample of initial offers. Such an offer has a monthly effective interest rate of 11.11% and an annual effective interest rate of 133.33%, despite having a 0% nominal interest rate.

OA.3 Variation in offer terms for the same applicant

Table OA.1 documents the dispersion within the same application three offer terms: offer amount, maturity, and effective interest rate. We use two different measures of dispersion: coefficient of variation (the standard deviation divided by the mean for all offers of the same applicant) and the ratio of maximum to minimum offer term. Both measures suggest substantial variation in all three offer terms within the same application. For instance, the average coefficient of variation for the offer amount is 0.57. That implies by moving two standard deviations from the mean, the offer amount is 114% higher. A more stark pattern is the ratio of maximum to minimum offer amount. This ratio is on average 8.7, implying substantial differences between the "best" and "worst" offers for the same applicant. Overall, the variation in offer terms suggests that the offer rate is indeed a relevant outcome variable for studying racial disparities in lending.

OA.4 Disparities in "best" offer

Even if there is substantial variation in offer terms (as documented in Appendix OA.3), it could be that the "best" offer is the same for minorities and Chinese. If that is the case, there are no disparities between Chinese and minorities, to begin with. While it is difficult to know which offer is the best offer for each applicant, we use three different definitions to determine an applicant's best offer and repeat our main analysis using the applicant's best offer. For each application, we define the best offer in terms of maximum offer amount, maximum offer maturity, and minimum effective interest rate across all the offers she receives from any lenders. This choice can be justified if applicants have lexicographic preferences in offer amount, maturity, and effective interest rates.

Table OA.2 documents the results of this analysis. Column 1 shows the results for Log(max offer amount). Before anonymization, minorities' maximum offer amount is 10.7% lower than Chinese. However, after anonymization, minorities' maximum offer amount is 3.7% lower than Chinese. These two coefficients are different both statistically and economically. Similar results hold for maximum offer maturity and annual effective interest rates.

OA.5 Predictability of the race information

To analyze whether race is predictable by other observable application characteristics, we first study the bi-variate correlations of race and other application characteristics. Panel (A) of Figure OA.3 shows the histogram of the absolute values of the correlations between race and an observable application characteristic. For an application characteristic such as marital status that takes more than two values (divorced, single, married, etc.), we use N dummy variables (N is the different level that this variable takes) that are equal to one if that characteristic is equal to one of the levels and zero otherwise. Hence, there are N points in the histogram for a variable with N levels. The figure suggests that even at the univariate dimension, application characteristics exhibit a correlation with race.

Panel (B) of Figure OA.3 shows the out-of-sample area under the curve (AUC) for predicting race using various machine learning algorithms with other observable application characteristics serving as the predictors. We start with the logistic regression model and consider several workhorse machine learning approaches: random forest, gradient boosting, and neural nets. In the logistic regression, we include squares and interaction terms of the predictors. For both gradient boosting and random forest, we use 1,000 classification trees. We consider three types of neural nets: (1) with one layer of 100 neurons, (2) three layers with 50 neurons each, (3) one layer of 200 neurons. In addition, we implement the stacked generalization (Wolpert, 1992). Basically, stacking is a way of combining predictions from multiple supervised machine learning models (known as the "base learners") into a final prediction to improve performance. For all machine learning methods considered, we train the model in the randomly drawn training sample of 10,000 applications and assess the classification accuracy in the validation sample of the remaining 6,281 applications using the commonly used AUC metric. Overall, the results suggest that race is highly predictable by other observable characteristics. For instance, the AUC for the stacking method, which is the most powerful prediction model, is 96.6%. For comparison, the rule of thumb cut-off for a "good" AUC in the credit scoring industry is 60% to 70% (Iyer, Khwaja, Luttmer, and Shue, 2016; Berg, Burg, Gombović, and Puri, 2020).

OA.6 Alternative samples and specifications

In this Appendix, we study the robustness of our results to alternative samples and specifications.

In Column 1 of Table OA.3, we include the entire sample of individuals and repeat the analysis using regression specification 1. For our main analysis, we have focused on the sub-sample of applications whose information is pre-filled directly from the Singapore government database. We have done so because the official records have higher data quality and less measurement error compared to self-reported information. Nevertheless, we repeat the analysis for the sample of all individuals. We find results similar to the main sample. Hence, our results are not sensitive to the filtering of applications.

In Column 2 of Table OA.3, we augment our controls for observable characteristics by allowing the effects of control variables X_i to differ *by lender* in the pre and post periods, in other words, the corresponding coefficients γ are now indicated by the $j, s(t)$ subscript.

$$\begin{aligned} y_{i,j} &= \pi_t + \alpha_{j,s(t)} + \gamma_{j,s(t)} X_i \\ &+ \beta_{pre} \times Minority_i \times pre_t + \beta_{post} \times Minority_i \times post_t + \varepsilon_{i,j} \end{aligned} \tag{OA.3}$$

We obtain similar estimates as in our baseline results.

OA.7 Heterogeneous racial disparities across lenders

The elimination of *average* racial gaps for all lenders masks the potential heterogeneity in lenders. In this Appendix, we want to study whether lenders who engaged in more discriminatory practices before anonymization are affected more. Prior research has documented the existence of substantial variation in discriminatory practices across large US firms in the labor market (Kline, Rose, and Walters (2022)). To examine the heterogeneity with respect to lenders, we estimate the racial gaps lender-by-lender in a specification analogous to equation (1). To do that, we include the Post indicator instead of Lender×Post fixed effects and cluster standard errors at the month level.

Panel (A) of Figure OA.4 shows the lender-specific β_{pre} (coefficient on the interaction between the minority indicator and the pre indicator) in the horizontal axis against $\Delta\beta = \beta_{post} - \beta_{pre}$ (the difference in the coefficient on the interaction between the minority indicator and the post indicator from the coefficient on the interaction between the minority indicator and the pre indicator, i.e., the treatment effect of anonymous applications) in the vertical axis for offer rates. Each circle in this scatterplot represents a lender in our sample and the size of the circle corresponds to the volume of applications the lender receives. The red line gives the best linear fit. We find a strong negative association between lender-specific racial gaps and the treatment effect of anonymous applications that is approximately one-for-one. In other words, the more discriminatory lenders, measured as the ones who give fewer offers to minorities before anonymization, increase offer rates more to minorities relative to other lenders.

Panel (B) of Figure OA.4 presents the same analysis for origination rates. We find a similar negative association between lender-specific discrimination and the treatment effect of anonymous applications. The more discriminatory lenders, measured as the ones with fewer loan originations to minorities before anonymization, increase loan originations more to minorities.

OA.8 Accuracy of lender beliefs

Our model closely follows the simple model outlined in Agan and Starr (2018). We briefly explain the model and its assumptions here. Assume that lender j offers a loan to applicant i if $u_j(x_i, m) + \varepsilon_{i,j} > u_j^*$ where m = 1 for minorities and m = 0 is for Chinese, x_i is a vector of borrower characteristics, and $\varepsilon_{i,j}$ is the preference parameters of lender j over applicant i; $u_j(x_i, m)$ is the utility of lender j from lending to an applicant with characteristics x_i and m. u_j^* is a fixed threshold above which the lender offers a loan. Then, the expected utility from a loan offer to applicant i by lender j, when not observing race is equal to $u_j(x_i, m = missing) = p(m = 1|x_i) * u_j(x_i, m = 1) + p(m = 0|x_i) * u_j(x_i, m = 0)$. If we make an additional simplifying assumption that $\varepsilon_{i,j}$ is uniformly distributed, that is $Pr(\varepsilon_{i,j} > \varepsilon) = A_j + B_j\varepsilon$, then

$$
\begin{aligned}
&Pr(\text{offer}|x_i, m = missing) \\
&= Pr\Big(\varepsilon_{i,j} > u_j^* - p(m = 1|x_i) * u_j(x_i, m = 1) - p(m = 0|x_i) * u_j(x_i, m = 1)\Big) \\
&= Pr(m = 1|x_i) \times Pr(\varepsilon_{i,j} > u_j^* - u_j(x_i, m = 1)) \\
&+ Pr(m = 0|x_i) \times Pr(\varepsilon_{i,j} > u_j^* - u_j(x_i, m = 0)) \\
&= Pr(m = 1|x_i) \times Pr(\text{offer}|x_i, m = 1) + Pr(m = 0|x_i) \times Pr(\text{offer}|x_i, m = 0)
\end{aligned}
$$

If we assume $x_{i,k}$ is characteristic k for individual i, and H is a level this variable takes, we have:

$$
\begin{aligned}
&Pr(\text{offer}|x_{i,k} = H, m = missing) \\
&= Pr(m = 1|x_{i,k} = H) \times Pr(\text{offer}|x_{i,k} = H, m = 1) \\
&+ (1 - Pr(m = 1|x_{i,k} = H)) \times Pr(\text{offer}|x_{i,k} = H, m = 0)
\end{aligned}
\tag{OA.4}
$$

Using Equation (OA.4), we can infer the subjective probability $Pr(m = 1|x_{i,k} = H)$ that a borrower is a minority for different levels of all control variables. For instance, if we focus on living in a private apartment (condo) as the characteristic, we observe these empirical probabilities in the data:

$$Pr(\text{offer}|x_{i,k} = \text{living in a condo}, m = missing) = 0.32$$
$$Pr(\text{offer}|x_{i,k} = \text{living in a condo}, m = 1) = 0.32$$
$$Pr(\text{offer}|x_{i,k} = \text{living in a condo}, m = 0) = 0.46$$

Using Equation (OA.4), we obtain that lenders infer $Pr(m = 1|x_{i,k} = \text{living in a condo}) = 0.86$. In the data, the empirical probability $Pr(m = 1|x_{i,k} = \text{living in a condo})$ is approximately 50% and is stable in the pre- and post-periods. For this case, the deviation is $0.86 - 0.5 = 0.36 = 36\%$.

We compare lender decisions before and after anonymization to their perceived probability of characteristic signaling minority status for all observable characteristics. Two patterns emerge. First, only 12 out of 146 inferred probabilities are between 0% and 100%. Second, the deviations of the inferred probabilities from their empirical counterparts are also sizable. Figure OA.5 shows the histogram of the difference between the inferred and empirical probabilities in our data. For ease of interpretation, we truncate the inferred probabilities at 0% and 100% before calculating the deviation from empirical probabilities. The absolute deviation, bounded above at 100% due to the truncation, exceeds 20% for 95% of application characteristics.[12] Hence, we conclude that accurate beliefs by race are not supported by the data.

We also test whether the inaccurate beliefs reflect stereotypes (Bordalo, Coffman, Gennaioli, and Shleifer, 2016). In their model, the decision process based on Kahneman and Tversky's representativeness heuristic produce stereotypes. An empirical prediction is that beliefs about a group are biased towards representative types, defined as the types that occur more frequently in that group than in a baseline reference group. We test this prediction in the data. For any application characteristic k and its possible values, we calculate the likelihood ratio $\frac{Pr(x_{i,k}=H|m=1)}{Pr(x_{i,k}=H|m=0)}$. A higher likelihood ratio means that type H for characteristic k occurs with higher relative frequency for minority applicants, hence a more representative type. Figure OA.6 plots the relative frequencies in the vertical axis against the deviation of inferred probabilities from empirical probabilities in the horizontal axis. Contrary to a positive relationship between these two predicted by stereotypes, the relationship is slightly negative. In other words, the representative types of minority applicants are not overweighted in lenders' beliefs.

[12] If we do not truncate the inferred probabilities, the absolute deviation of the inferred probability from the empirical probability exceeds 300% for close to 30% of the characteristics we consider.

Figure OA.1: **Testing for matching between applications and lenders**

This figure shows the coefficients and the associated 95% confidence intervals of the regressions of lender rank on the minority status across all lenders. For each application, lender rank is an integer (starting from 1) corresponding to the order in which the application is sent to the given lender. We regress lender rank on minority status for a given lender at a time and iterate through all lenders. Each colored coefficient and the associated confidence interval correspond to one lender.

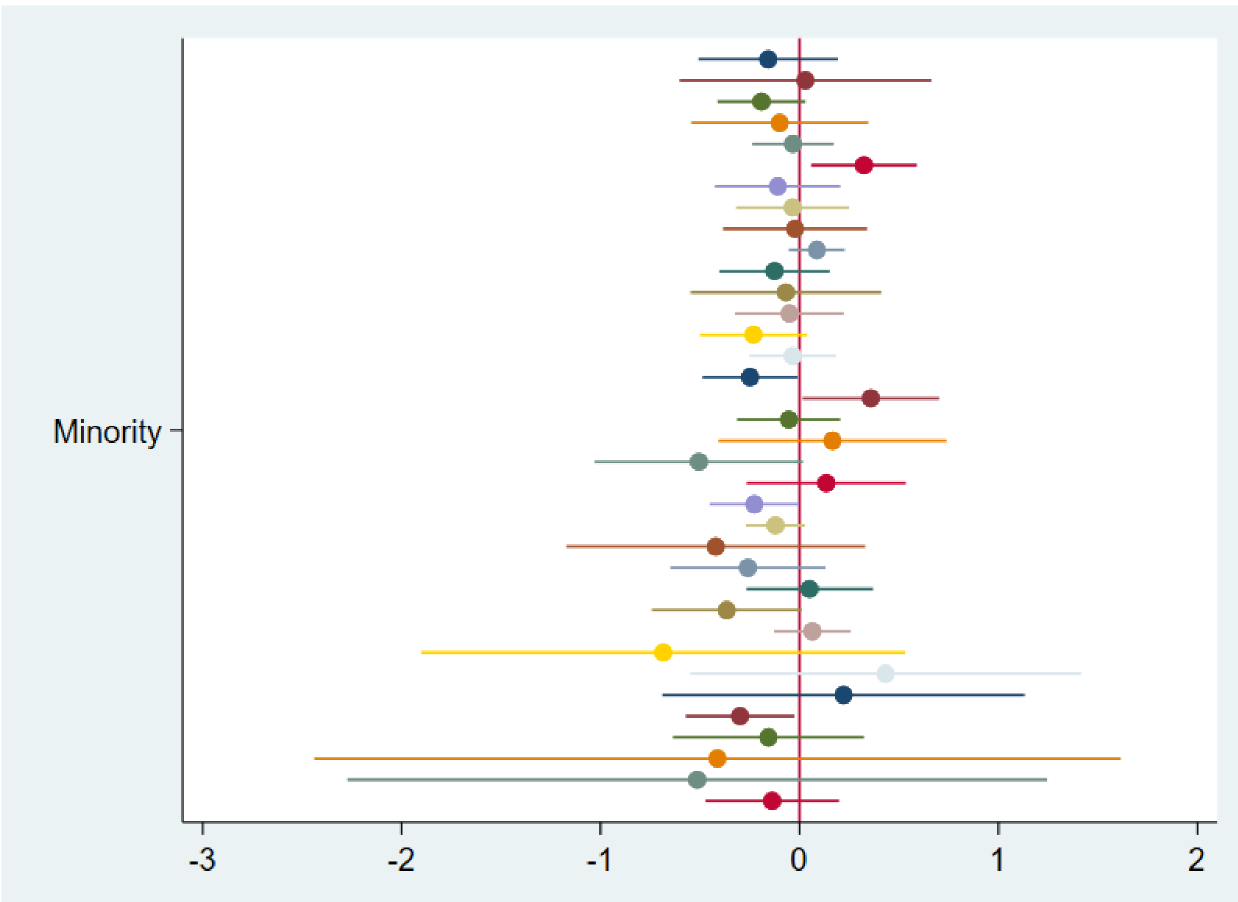

Figure OA.2: **Empirical distribution of offer interest rate and processing fee**

Panel (A) shows the empirical cumulative distribution function of initial offers' interest rates. Panel (B) shows the empirical cumulative distribution function of initial offers' processing fees.

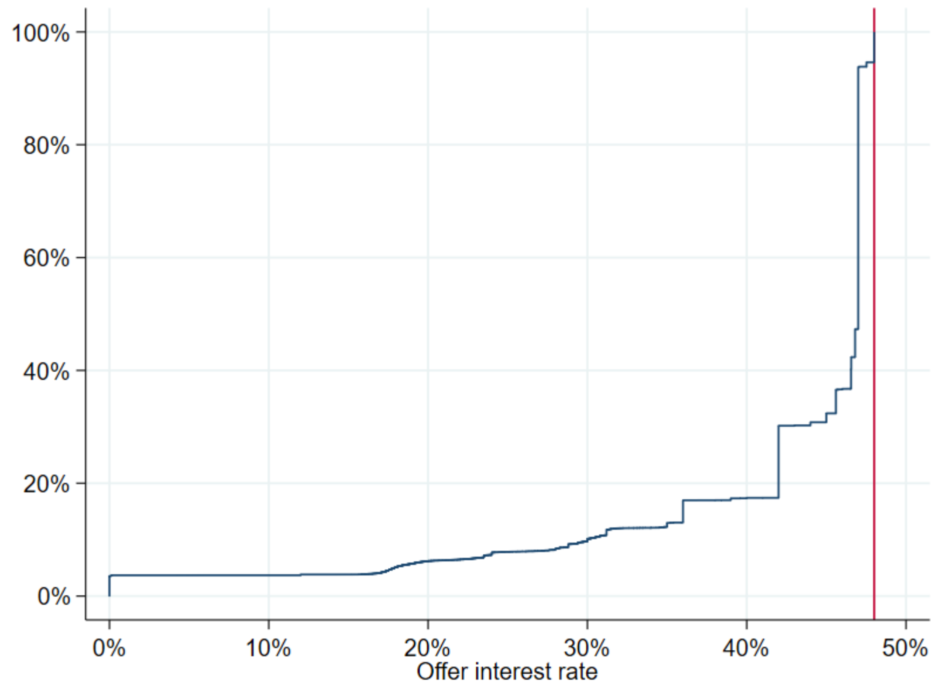

(A) Empirical distribution of offer interest rate

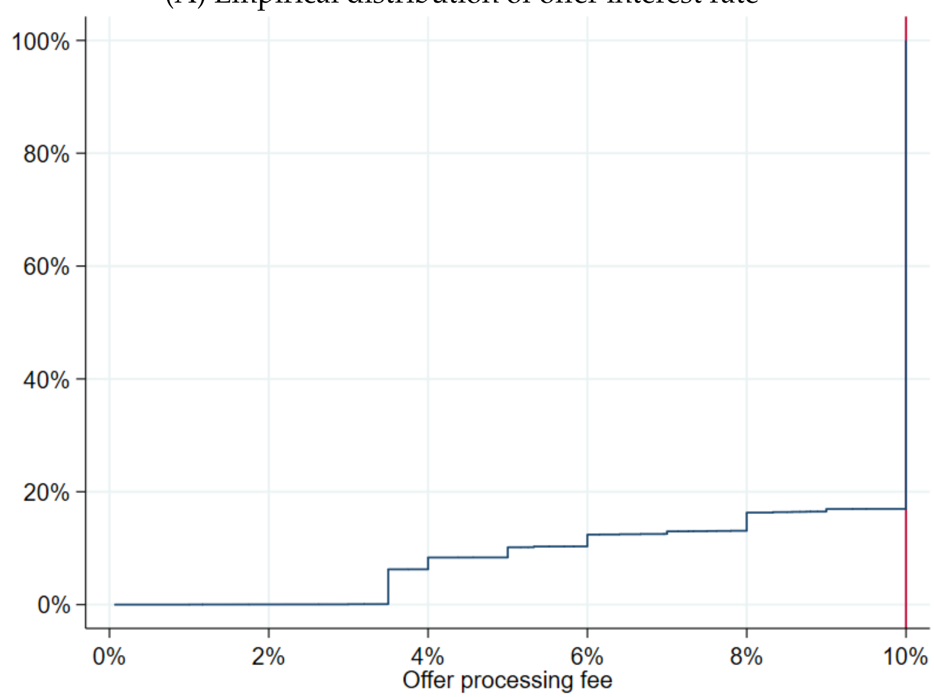

(B) Empirical distribution of offer processing fee

Figure OA.3: **Predictability of the race information**

Panel (A) plots the histogram of the absolute values of the correlations between race and an observable application characteristic. Panel (B) plots the out-of-sample area under curve of various classification analyses for predicting race using other observable application charectristics.

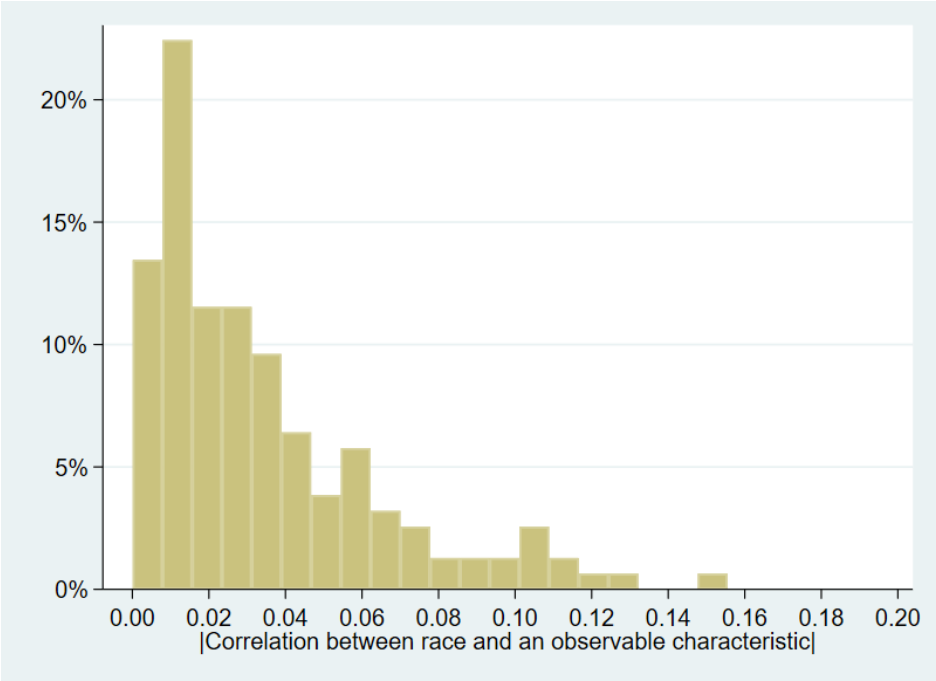

(A) Distribution of bi-variate correlations of race and other application characteristics

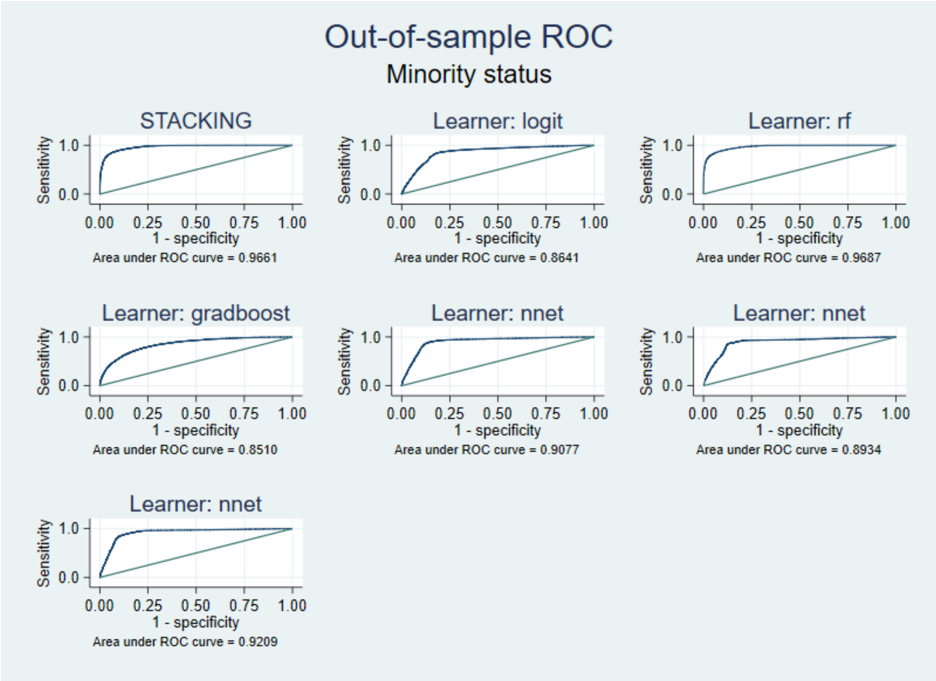

(B) Out-of-sample area under curve of classification analyses

Figure OA.4: **Lender-specific racial disparities**

This figure shows the lender-specific β_{pre} (coefficient on the interaction between the minority indicator and the pre indicator) in the horizontal axis against $\Delta\beta = \beta_{post} - \beta_{pre}$ (the difference in the coefficient on the interaction between the minority indicator and the post indicator from the coefficient on the interaction between the minority indicator and the pre indicator, i.e., the treatment effect of anonymous applications) in the vertical axis for offer rates in Panel (A) and origination rates in Panel (B). In both panels, each circle represents a lender in our sample, and the size of the circle corresponds to the volume of applications the lender receives. The red line gives the best linear fit.

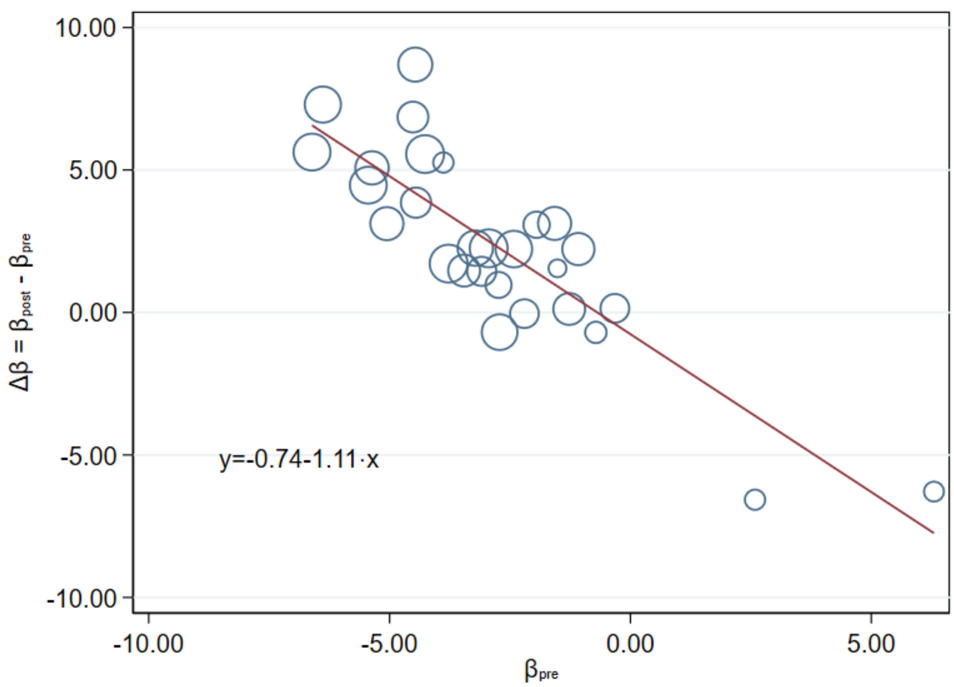

(A) Dependent variable: Application-lender level offer indicator ($\times$ 100)

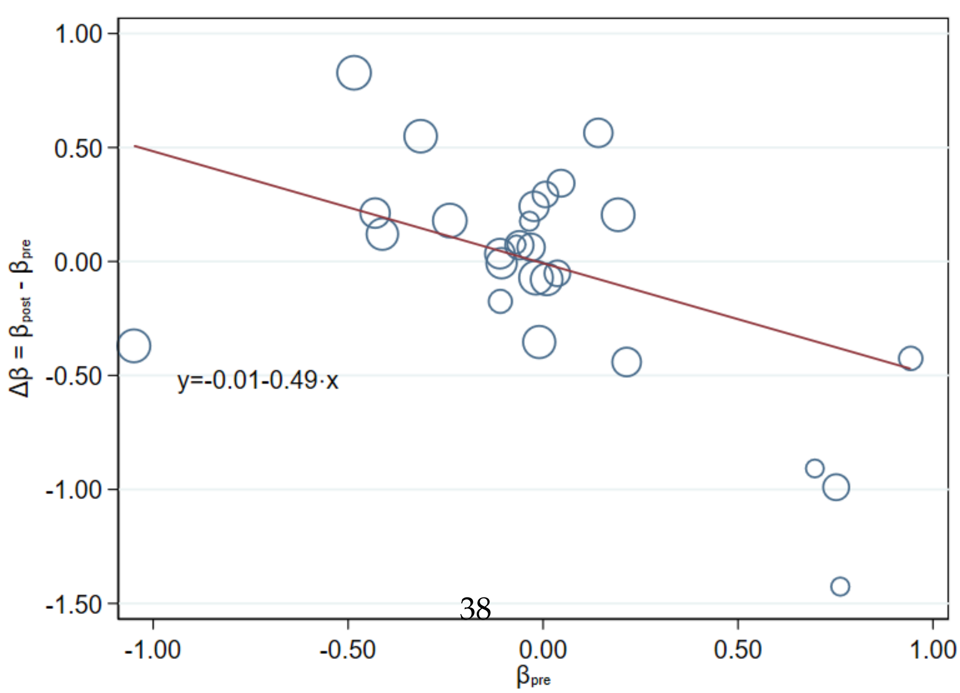

Figure OA.5: **Histogram of the difference between the inferred and the empirical probabilities**

This figure plots the histogram of the difference between the inferred probabilities and the empirical probabilities in our data. The inferred probability is the subjective probability (held by lenders) that the applicant belongs to the minority group after observing a certain characteristic (e.g., the probability that the application belongs to the minority group conditional on observing an applicant owns a property). For ease of interpretation, we truncate the inferred probabilities at 0% and 100% before calculating the deviation from empirical probabilities. A larger deviation on either side implies that lenders' prior are more inaccurate. Each point used in the plot corresponds to one application characteristic.

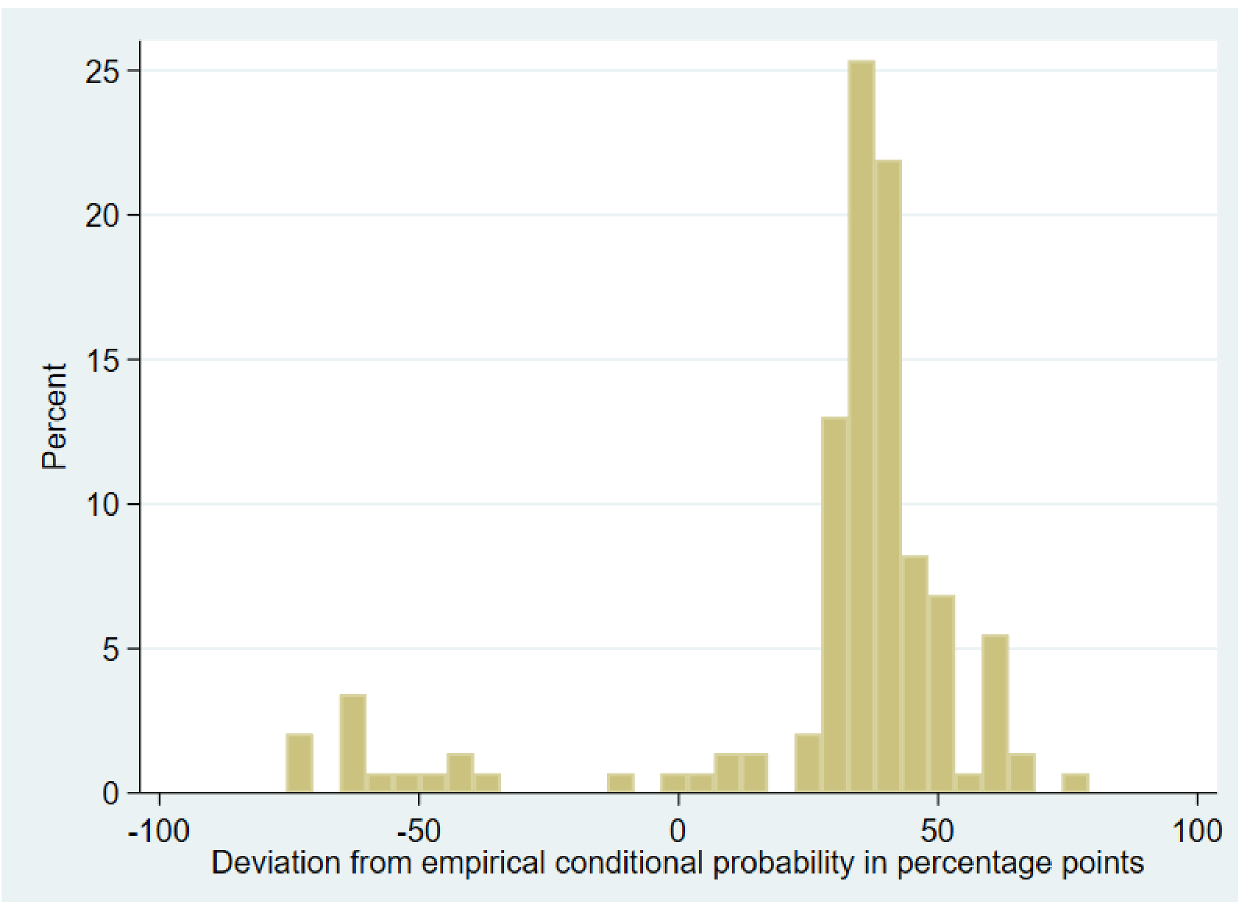

Figure OA.6: **Representative types and lender beliefs**

This figure shows the difference between the inferred probabilities and the empirical probabilities in the horizontal axis against relative frequency in the vertical axis. The inferred probability is the subjective probability (held by lenders) that the applicant belongs to the minority group after observing a certain characteristic (e.g., the probability that the application belongs to the minority group conditional on observing an applicant owns a property). For ease of interpretation, we truncate the inferred probabilities at 0% and 100% before calculating the deviation from empirical probabilities. Relative frequency is calculated following (Bordalo, Coffman, Gennaioli, and Shleifer, 2016) as the ratio of the likelihood of belonging to a type among minorities to the likelihood of belonging to a type among Chinese. A high relative frequency corresponds to a more representative type for minority applicants. Each point used in the plot corresponds to one application characteristic.

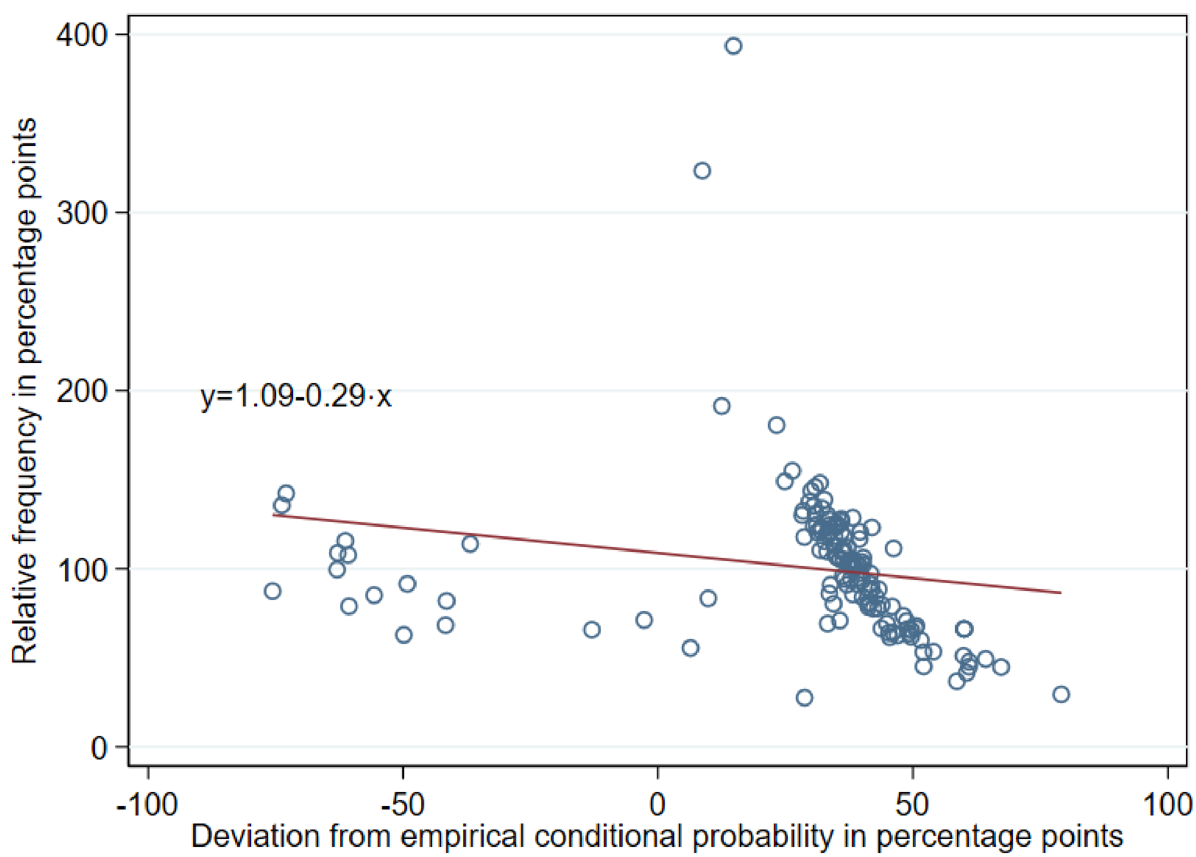

Table OA.1: **Within-application dispersion of initial offer terms**

This table reports the summary statistics for the within-application dispersion of initial offer terms. The coefficient of variation is the standard deviation divided by the mean of all offers given to an application. Offer maturity is measured in months.

	Mean	Std. Dev.	Median
Within-application coefficient of variation:			
offer amount	0.57	0.26	0.54
offer maturity	0.60	0.22	0.59
effective interest rate	0.36	0.14	0.37
Ratio of within-application maximum to minimum:			
offer amount	8.71	11.18	5.33
offer maturity	8.93	5.76	12.00
effective interest rate	3.26	1.80	2.90
Number of applications	14,991		

Table OA.2: **The effect of anonymous applications on disparities in best initial offers**

This table shows the effect of anonymous applications on disparities in best initial offers (equation (1)). For each regression, we also report $\Delta\beta = \beta_{post} - \beta_{pre}$, its t-statistic, p-value, and its value divided by the mean of the dependent variable. Fixed effects are included and denoted at the bottom. We include all the information available to lenders at the time of initial screening as control variables and allow the effects of the control variables to differ in the pre- and post-periods. In choosing the function form of the included control variables, we convert all continuous numerical characteristics (e.g., income) to categorical variables using their quintiles to allow for non-linear effects in control variables and for the retention of missing values. Standard errors are clustered at the month level; the corresponding t-statistics are reported in brackets. We use ***, ** and * to denote significance at 1%, 5% and 10% level (two-sided), respectively.

	(1) Log(max offer amount)	(2) Max offer maturity (months)	(3) Min effective interest rate (%)
Minority × Pre	-0.105***	-0.658***	3.993***
	[-7.27]	[-6.02]	[5.28]
Minority × Post	-0.0211**	0.00990	-0.487
	[-2.76]	[0.13]	[-0.77]
$\Delta\beta$	0.0837	0.667	-4.480
t-stat of $\Delta\beta$	5.288	4.997	-4.641
p-value of $\Delta\beta$	0.0000910	0.000159	0.000320
$\Delta\beta$ / Mean DV	0.00967	0.0567	-0.0755
Year-Month FEs	Yes	Yes	Yes
Post FE	Yes	Yes	Yes
Observable controls	Yes	Yes	Yes
R^2	0.503	0.410	0.281
No. of observations	14,991	14,991	14,991

Table OA.3: **The effect of anonymous applications on disparities in offer rates (robustness)**

This table shows the robustness checks on the effect of anonymous applications on disparities in offer rates. Column 1 estimates the baseline specification (equation (1)) in the full sample. Column 2 estimates the augmented specification (equation (OA.3)). For each regression, we also report $\Delta\beta = \beta_{post} - \beta_{pre}$, its t-statistic, p-value, and its value divided by the mean of the dependent variable. Fixed effects are included and denoted at the bottom. Control variables include all the information available to lenders at the time of initial screening. In choosing the function form of the included control variables, we convert all continuous numerical characteristics (e.g., income) to categorical variables using their quintiles to allow for non-linear effects in control variables and for retention of missing values. The baseline specification (equation (1)) allows the effects of the control variables to differ in the pre- and post-periods. The augmented specification (equation (OA.3)) allows the effects of the control variables to differ by lender and by whether the application is in the pre- vs post-period. Standard errors are clustered at the lender-month level; the corresponding t-statistics are reported in brackets. We use ***, ** and * to denote significance at 1%, 5% and 10% level (two-sided), respectively.

	Application-lender level offer indicator ($\times$ 100)	
	(1)	(2)
	Full sample	Tighter controls
Minority $\times$ Pre	-4.006***	-3.779***
	[-19.02]	[-17.48]
Minority $\times$ Post	-0.411*	0.335
	[-1.73]	[1.26]
$\Delta\beta$	3.595	4.114
t-stat of $\Delta\beta$	11.33	12.08
p-value of $\Delta\beta$	5.73e-26	8.34e-29
$\Delta\beta$ / Mean DV	0.0986	0.108
Year-Month FEs	Yes	Yes
Lender $\times$ Post FEs	Yes	Yes
Observable controls	Yes	Yes
R^2	0.293	0.400
No. of observations	468,663	322,847