Reducing Racial Disparities in Consumer Credit: Evidence from Anonymous Loan Applications *

Poorya Kabir and Tianyue Ruan†

December 2023

Abstract

Using a unique experiment of anonymizing online loan applications, we test whether race-blind loan screening procedures reduce racial disparities in consumer credit. With names on applications, ethnic minority applicants are 10.6% less likely to receive online loan offers and receive worse offer terms than otherwise identical ethnic majority applicants. Anonymizing applications reduces such disparities substantially. High-income minority applicants benefit more than low-income minorities. We show that the racial disparities are not driven by differences in socioeconomic status or credit demand. Overall, anonymous loan applications reduce racial disparities in access to credit by increasing lender reliance on objective credit risk measures.

JEL Classification: D63, G21, J15, J71, O16.

Keywords: consumer credit, fair lending, racial discrimination, anonymity, privacy

*We thank Sumit Agarwal, Scott Frame (discussant), Pulak Ghosh (discussant), Sebastian Krul (discussant), Will Shuo Liu (discussant), Dan Luo (discussant), Karsten Müller, Jessica Pan, Ivan Png, Vatsala Shreeti (discussant), Arkodipta Sarkar, and Greg Weitzner (discussant) for their insightful comments. We are grateful to seminar participants at the National University of Singapore, Fudan University, IIM-Ahmedabad, Naikai University, Virtual Household Finance Seminar, Peking University, Renmin University of China, and the 25th EdukCircle, as well as conference participants at the Asia Pacific Online Corporate Finance Workshop, Applied Economic Workshop, Research Symposium on Finance and Economics, ABFER "Innovation, Productivity and Challenges in the Digital Era: Asia and Beyond" Webinar, Boulder Summer Conference on Consumer Financial Decision Making, North America Summer Meeting of Econometric Society, European Economic Association, European Financial Management Association, European Finance Association (poster), Asian Meeting of the Econometric Society, Queensland Corporate Finance Conference, CEAR-RSI Household Finance Workshop, and Economics of FinTech 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. 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

Access to credit is unequal across racial groups (Butler et al., 2022; Bhutta and Hizmo, 2021; Bartlett et al., 2022; Pope and Sydnor, 2011b). Restricting the use of information predictive of race is used in various settings to mitigate racial disparities. One particular 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). While anonymizing applications by humans can be time-consuming and error-prone (Krause et al., 2012), the growing use of information technology in lending (Berg et al., 2022; Fuster et al., 2019) can achieve scalable and cost-effective anonymization: FinTech platforms can serve as intermediaries between lenders and applicants, verify applicants, and withhold applicants' racial identities from lenders. Even though FinTech's impact on racial disparities through algorithmic decision-making is studied (D'Acunto et al., 2023; Dobbie et al., 2021; Howell et al., 2021; Fuster et al., 2022), evidence on its influence through anonymous applications remains limited.

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 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 lenders. The online platform sends an individual's application to multiple lenders simultaneously. After the lenders review online applications and make initial loan offers, the individual chooses one offer online and visits the lender in person for identity verification, as required by customer due diligence regulation, before the 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. We refer to this change in policy as anonymization.

Whether anonymous loan applications can successfully reduce racial gaps in access to credit is ambiguous. If race, beyond objective credit risk criteria, factors into lending decisions, removing lenders' access to race can alleviate the racial disparities (Goldin and Rouse, 2000; Pope and Sydnor, 2011a). If lenders rely solely on names as a race proxy, removing names can effectively remove lenders' access to race. Conversely, if race correlates with the remaining application characteristics, lenders may be able to infer race, and racial disparities could persist. Restricting the use of information predictive of race can even hurt racial minorities due to biased beliefs (Agan and Starr, 2018) or reduced signal precision (Bartik and Nelson, 2020).

We find that before implementing anonymization, when names are on applications, ethnic minority applicants, including Malays, Indians, and other races, are 10.6% 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 probability disappear.

In the intensive margins, minority applicants receive smaller, lower maturity, and more expensive offers than otherwise identical Chinese applicants when applicant racial identities are shown to lenders. Such racial disparities across the intensive margins are significantly reduced after anonymization.

Before anonymization, high-income minorities are only slightly more likely to receive offers than low-income minorities. While anonymization improves access to credit for both low and high-income minorities, high-income minorities benefit more, suggesting increased lender reliance on objective credit risk measures, such as income.

We conduct a battery of robustness checks to rule out several alternative explanations. We control for socioeconomic status signals in names (Fryer and Levitt, 2004) and obtain similar findings. We also directly test for race-specific demand for loan characteristics

(Bhutta and Hizmo, 2021; Willen and Zhang, 2022) in our setting and find no discernible differences in demand between minority and Chinese applicants. In our context, it's common for applicants to receive several offers. Therefore, the relevance of differences in offer probability depends on whether the most favorable offers for minorities and Chinese applicants are comparable. We consider three intuitive measures of best offers (defined using offer amount, interest rate, or maturity) and find that before anonymization, the best offers received by minorities are significantly worse relative to Chinese, and such disparities are reduced substantially after anonymization.

Anonymizing applications delays revealing race to the point when applicants visit lenders for the required identity verification. At this in-person stage, lenders can use the newly available information on race to fully undo the effect of anonymous applications. We find that the reduction in the racial gap in initial offer probability is attenuated by approximately 20% in the origination stage, when lenders learn about applicant race. The net reduction in the racial gap in loan origination implies that anonymous loan applications reduce disparities in access to credit.

We study the loan performance of minority and Chinese applicants using data from one lender who originates 14% of the loans (due to data limitations). We find that the average delinquency rate for minority and Chinese applicants 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 before anonymization and the substantial reduction of such 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 observable to lenders at the time of initial screening. Anonymization can be effective in reducing racial disparities regardless of the nature of discrimination present in this market. Nevertheless, to understand whether taste-based or statistical discrimination best describes our results, we perform two additional tests. First, we find similar racial gaps across different levels of income in the pre-period, 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' racial beliefs in lending decisions and find that lenders' beliefs are inaccurate. These findings suggest that statistical discrimination is unlikely to explain our findings.

Our study contributes to the large and growing literature on racial disparities in credit markets (Butler et al., 2022; Bartlett et al., 2022; Bhutta et al., 2022; Pope and Sydnor, 2011b). 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 et al. (2016)), similar to the lower callbacks faced by minority applicants in 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 et al., 2022) and minority loan officers (Frame et al., 2022; Jiang et al., 2022) can reduce racial disparities. Several studies show how two main flavors of technology (Berg et al., 2022) can reduce racial disparities. First, the use of alternative data or machine learning can reduce human bias (Berg et al., 2020; Dobbie et al., 2021; Howell et al., 2021; Fuster et al., 2022). Second, by changing how borrowers and lenders interact, technology can enable cost-effective changes in lenders' information sets. We show that anonymous loan applications, which can be implemented in scale by simple technologies, can effectively reduce racial disparities in credit access. Our analysis also highlights how removing lender access to borrower race leads to an increase in lender reliance on objective credit risk measures and therefore can effectively reduce racial disparities.

D'Acunto et al. (2023) analyze a related technology-based change in lenders' information sets–adding suggestions from algorithmic decisions in P2P credit. Our paper differs in several ways. First, all lenders in our setting are affected by anonymization, alleviating the confounding factor of self-selection into participation (Behaghel et al., 2015). Moreover, lenders in our setting are professional and specialized lenders and hence are skilled in screening borrowers, incentivized to make profitable loans, and able to learn from experience, whereas P2P lenders are likely to be households with arguably little expertise in credit evaluation. Finally, algorithmic suggestions may embed racial biases due to triangulation of otherwise excluded characteristics (Fuster et al., 2022).

2 Institutional settings

Using a unique experiment of anonymizing applications, we study racial disparities in the consumer loan market in Singapore. Consumer loans are uncollateralized short-term installment loans borrowed for personal uses such as medical treatment, credit card debt repayment, education, wedding, etc. These loans are extended by licensed and regulated lenders. The FinTech platform we study allows potential borrowers to apply to multiple lenders simultaneously and compare loan offers. Below, we describe the consumer loan market and the FinTech platform in great detail.

2.A Structure of consumer loans

In our setting, each loan contract is characterized by four standardized dimensions: (1) loan amount in Singapore Dollars (1 SGD = 0.75 USD as of January 2021), (2) loan maturity in months, (3) nominal annual interest rate, and (4) processing fee as a percentage of the loan amount.

The structure of a loan follows an equated monthly installment repayment schedule, similar to a mortgage loan or an auto 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.

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

With processing fee f, the applicant receives $B \times (1-f)$ as opposed to B upon loan origination. The existence of a processing fee implies that the nominal interest rate only incompletely captures the true borrowing cost. To analyze the true borrowing cost, we calculate the effective interest rate. The monthly effective interest rate r is determined by:

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

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 accounts 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.

In our data, the average loan amount is S$4,300, the average maturity is 6 months, and the average effective annual interest rate is 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 et al., 2021).

The loans are extended by licensed lenders who are small financial institutions whose main line of business is extending consumer loans. These lending decisions are done by loan officers and not by sophisticated machine learning algorithms. Singapore's Ministry of Law is the main regulator of the money lenders. It sets price limits for consumer loans: effective from 1 October 2015, the legal upper limit for the interest rate is 4% per month and the legal upper limit for the processing fee is 10% of the loan amount. It also requires that lenders perform in-person identity verification of borrowers before loan origination, even if the loan offers can be extended online (e.g., through the FinTech platform which we describe in the next subsection). Furthermore, no ubiquitous credit scoring for loans in this market exists.

2.B The FinTech platform and the anonymization experiment

The setting for our analysis is a leading FinTech platform that allows potential borrowers to obtain a consumer loan after comparing loan offers. As of January 2022, the platform partners with 37 of 156 licensed lenders of consumer loans in Singapore.[1] The process for an applicant to apply for and obtain a loan through the platform is as follows: an applicant fills out a standardized loan application on the platform; the application is sent to multiple lenders. Lenders review online applications and decide whether to extend initial offers 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 as required by the regulation, 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.

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

[1] This platform is also active in Hong Kong and Australia.

as applicant names contain racially identifying information (Bertrand and Mullainathan, 2004). Applicant location remains observable to lenders. In our context, the location information is not predictive of race and therefore does not undermine anonymization,[2] as locations in a small city-state such as Singapore are relatively homogeneous. The Singapore government's housing policy further prevents granular ethnic segregation (Agarwal et al., 2019; Wong, 2013).

To study racial disparities in consumer loans, we match applicant names to races and validate our hand matching using official data. According to the Singapore Census of Population, as of 2020, Singapore's resident population consists of 74.3% being Chinese, 13.5% being Malays, and 9.0% being Indians. As the Chinese are the ethnic majority in Singapore, we consider all non-Chinese applicants to be minority applicants for our analysis.

3 Data and summary statistics

We obtain detailed data from the FinTech platform on application characteristics, initial offers, loan originations, and loan performance for the period from October 2020 to January 2022.

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

We also observe lenders' decisions such as whether the lender approves an initial offer to the applicant and the offer terms (amount, maturity, interest rate, and processing fee). These lending decisions are done by loan officers and not by sophisticated machine

[2] Experiments of anonymous hiring procedures in several European countries (Krause et al., 2012; Behaghel et al., 2015) involve the removal of addresses in addition to names to implement anonymization.

learning algorithms. Additionally, we observe whether a loan was originated and, if so, the origination terms.

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 a late payment if the lender expects a repayment according to the repayment schedule but receives none.

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 three 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.[3] Third, we can cross-check the accuracy of race classifications which we will elaborate on below. Our final sample includes a total of 322,847 lender-applicant pairs for 16,281 applicants with 2,733 originated loans.

We measure applicants' race by matching their names to races following Wong (2013). In our classification, we 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. This approach is feasible as different race groups in Singapore have distinct names. In Appendix IA.2, we provide two pieces of evidence showing that names tend to uniquely map to race in Singapore using race records obtained by the platform at the time of application. First, we show that race identified by research assistants is the same in 98.4% of observations. Second, we plot the histogram of the percentage of a name used by minorities. Only 5.1% of names are shared between minorities and Chi-

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

nese, 20.3% are only used by Chinese, and 74.6% are only used by minorities. Hence, in our settings, the sample includes racially distinctive names commonly used in correspondence studies. In our final sample, about 61% of the applicants are minorities. For our analysis, we use race classification by research assistants (rather than the one in government records), because it captures whether a name sounds like a minority name more accurately. Nevertheless, the main results are robust to using race classification from government records.

Panel A of Table 1 provides summary statistics of application characteristics. There are a total of 16,281 applications during our sample period from October 2020 to January 2022. The average applicant is 36 years old, and 75% of the applicants are male. There are 11,789 applications submitted before September 28, 2021 when lenders can see applicant names 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 minority and Chinese 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 at the application level. We calculate the average offer probability across all lenders who receive the application. We also count the number of loan offers each application receives. We calculate the average offer amount, maturity, nominal interest rate, processing fee, and effective interest rate across all loan offers the application receives for each application. Lastly, as each applicant can have at most one loan origination, we calculate their origination probability. Column 1 shows that an average applicant has a 44% probability of receiving loan offers across lenders and receives 7.6 loan offers. These loan offers are uncollateralized 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%. Columns (2) and (3) of Panel B report the mean difference between minority and Chinese applicants before and after anonymization, respectively. Column (2) shows stark unconditional racial gaps across the board when lenders know the applicant's race when evaluating loan applications: Minority applicants receive fewer offers, lower loan amounts, shorter maturity, higher annual effective interest rates.[4] They are also less likely to receive eventual loan originations. Column (3) shows that 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 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 anonymization. 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 racial disparities to anonymization is that the racial disparities would stay stable absent of the change, analogous to a standard parallel trends assumption in difference-in-differences designs. In our analyses, we formally validate this assumption and perform a battery of diagnostic tests to rule out several threats to our identification.

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

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

[4]In Appendix IA.3, 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 minority applicants across these two dimensions of offer terms.

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 the 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 anonymization and corresponds to the treatment effect of anonymous applications. If race enters lending decisions, that is $\beta_{pre} < 0$, anonymizing race information could help mitigate disparities (Pope and Sydnor, 2011a). If lenders rely only on names as a race proxy, removing names can effectively remove lenders' access to race ($\beta_{post} = 0$). If race is closely associated with other application attributes, lenders might still infer race, and racial disparities could persist ($\beta_{post} < 0$). Racial disparities could even increase ($\beta_{post} > \beta_{pre}$) if beliefs are biased (Agan and Starr, 2018) or signal precision differs by race (Bartik and Nelson, 2020).

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.[5] As applicant names are visible to lenders in the pre-period but not in the post-period, lenders may use other variables that are available throughout the sample period differently in the two periods when they screen applicants. To capture this potential change in lenders' screening, we allow the effects of these control variables γ to differ in the pre and post-periods (hence the $s(t) \in \{pre, post\}$ subscript). $\alpha_{j,s(t)}$

[5] The list of control variables includes the age of the applicant, applied amount, applied loan maturity, 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, current remaining bank-loan balance and its monthly payment, current remaining consumer loan balance and its monthly payment, and the planning area.

for $s(t)$ 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. 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. Standard errors are clustered at the lender-month level to allow for correlated decision-making across applications by a lender in a month.

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 initial offer decisions and control for them in the regressions. Therefore, the omitted variables bias is unlikely to affect our findings.

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 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 to facilitate a percentage point interpretation. 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 a coefficient on the interaction term between $Minority_i$ and Pre_t of -3.81, implying that when names are on applications, minority applicants are 3.81 percentage points less likely to receive initial offers than 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 one to zero values, and then log-transform all continuous numerical variables. We maintain the inclusion of Lender$\times$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 probability analogously. To match this level of aggregation, we now include the Post indicator, as opposed to Lender$\times$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 prior to anonymization and economically small and statistically insignificant racial gaps after anonymization. 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 unlikely to bias our findings.

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

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

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 parallel trends in our research design. 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 is a statistically significant racial gap; the magnitude of the racial gap stays stable around its average level of 3.81 percentage points and is also similar to the level seen in the previous months. This pattern validates the key identifying assumption where the racial gap would have remained constant absent anonymization. For the two months following anonymization, we see insignificant coefficients, implying that the racial gap in offer probability is eliminated.

5.B Heterogeneity across lenders

The elimination of the *average* racial gaps across lenders masks the potential heterogeneity in lenders. In this subsection, we study whether lenders who are more biased against minorities before anonymization are affected more. Prior research has documented the existence of substantial variation in racial biases across large US firms in the labor market (Kline et al. (2022)). To examine the heterogeneity for 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.

Figure 2 shows the lender-specific β_{pre} (coefficient on the interaction between the minority indicator and the pre indicator) in the horizontal axis against the treatment effect of anonymous applications $\Delta\beta = \beta_{post} - \beta_{pre}$ in the vertical axis. 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 biased lenders, measured as the ones who give fewer offers to minority applicants before anonymization, increase offer probability to minority applicants more relative to other lenders.

An important feature of our setting is that anonymization affected *all* lenders simultaneously. If lenders were given a choice to receive applications with or without names, anonymization would have ambiguous effects on racial disparities. Behaghel et al. (2015) document that self-selection into anonymized job applications by minority-friendly employers increases racial gaps in callback rates. In our setting, if only lenders that favor minorities in pre-period (those with $\beta_{pre} > 0$) would have decided to receive anonymized applications, racial gaps could have widened. In other words, the mandatory participation of lenders in our setting allows us to alleviate the confounding factor of self-selection into participation.

5.C Disparities in the intensive margins

To analyze the effect of anonymization on racial disparities in the intensive margins, we consider three offer terms: (1) offer amount, (2) offer maturity, and (3) effective interest rates. Table 3 documents the results of this analysis. Column 1 focuses on the log of offer amount as the outcome variable. Before anonymization, minority applicants' offer amount is 4.8% lower than Chinese applicants'. As a higher offer amount gives the applicant a higher amount available for personal use, a lower offer amount faced by minority applicants suggests that they receive worse offers in the first offer term we examine. After anonymization, the gap in offer amount between minority and Chinese applicants is reduced to 2.7%. Column 2 reports the estimates for the analysis of the offer maturity. When names are shown on applications before anonymization, the offer maturity is 0.15 months shorter for minority applicants than for Chinese applicants. As a lower maturity corresponds to a lower monthly payment all else equal, a lower offer maturity faced by minority applicants also suggests that they receive worse offers in terms of maturity. The difference reduces to 0.09 months in the post period. Column 3 reports the estimates for the analysis of the effective interest rate. Before anonymization, minority applicants receive more expensive loan offers–the effective interest rates are 0.84 percent points higher than those received by Chinese applicants. After anonymization, the effective interest rate differential is reduced to 0.62 percentage points.

Across the different intensive margins, minority applicants receive worse offers than Chinese applicants when applicant racial identities are shown to lenders.[6] Such racial disparities in the intensive margins are significantly reduced after anonymization.

5.D Who benefits more from anonymization?

Our results so far suggest that minorities benefit from anonymization both in the extensive and intensive margins. These effects, however, could be heterogeneous. For instance, if lenders use more objective variables such as income for their decisions after anonymization, high-income minorities will benefit more. In this subsection, we study the heterogeneity of the effect by income.

To do that, we split the sample of applicants into two groups based on median income

[6] In Section 6.B, we examine applicants' revealed preferences based on what offers they choose and verify that loan offers with a lower offer amount or a shorter offer maturity are indeed worse offers.

and define high and low-income dummy as the applicants with above and below median income. We then interact with these variables $Minority_i \times Pre_t$ and $Minority_i \times Post_t$. We use the following regression specification:

$$y_{i,j} = \pi_t + \alpha_{j,s(t)} + \gamma_{s(t)} X_i + \beta_{pre,H} \times Minority_i \times Pre_t \times HighIncome_i$$
$$+\beta_{pre,L} \times Minority_i \times Pre_t \times LowIncome_i + \beta_{post,H} \times Minority_i \times Post_t \times HighIncome_i$$
$$+\beta_{post,H} \times Minority_i \times Post_t \times LowIncome_i + \varepsilon_{i,j} \tag{3}$$

The controls, fixed effects, and level of clustering are the same as our main specification (equation (1)). Like before, the omitted group is the sample of all Chinese applicants.

Table 4 shows the heterogeneity results by income. In Column (1), we observe that before anonymization, high-income minorities had slightly higher offer rates than low-income minorities. However, after anonymization, the difference between high and low-income minorities becomes substantial. Formally, $(\beta_{post,H} - \beta_{post,L}) - (\beta_{pre,H} - \beta_{pre,L})$, that captures the gains of high relative low-income minorities from anonymization equals 2.56% and is statistically significant at 1% level. Columns (2), (3), and (4) show similar findings for loan amount, tenure, and effective interest rate. In Internet Appendix IA.4, we show that the results in Table 4 is robust to using alternative measurement of objective qualification, different quantiles as cutoffs, or different empirical specifications.

The finding that high-income minorities benefit more from anonymization is consistent with the idea that lenders use income more prominently in their lending decisions. In our setting, without access to names, lenders use an objective measure such as income more prominently in their lending decisions. Consistent with increased reliance on income, we find that the adjusted R^2 of a regression where the left-hand side variable is the offer indicator and the right-hand side is income goes up after anonymization.

6 Robustness checks

6.A Names and socioeconomic status

Names convey information about race but can also signal socioeconomic status (Fryer and Levitt, 2004). Hence, an alternative interpretation of our finding is that lenders use names to proxy for socioeconomic status, not race. This interpretation is unlikely in our setting. Firstly, this alternative explanation is more likely to apply to labor market studies, where socioeconomic status information is unavailable on job applications. In our setting, several applicant characteristics, such as income and residential type – likely to be highly correlated with socioeconomic status – are available to lenders at the time of decision. Secondly, if socioeconomic status can fully explain our results, we expect to find lower racial disparities for applicants with higher social status, i.e., higher income as a result of lenders using names to infer the status. However, in Table 4 (as well as in Table IA.2 and Figure IA.4 in the appendix), we find that racial disparities are similar across different quantiles of income when lenders receive name-bearing applications, suggesting that names are not only used to infer social status.

Lastly, we perform an additional empirical test by controlling for a measure of social status based on applicant names. To do that, we split each full name into separate name parts, and use a leave-one-out approach to calculate the mean income for each name part using the entire sample. For each full name, we then take the average of this variable across different name parts and log transform the outcome to construct a measure of social status based on applicant names.[7] If social status drives our results entirely, we expect racial disparities to disappear after controlling for social status. Table 5 shows that the point estimates for racial disparities barely change after controlling for the measure of social status, suggesting that our results are not driven by names signaling social status.

[7]For unique name parts, the value of mean income will be missing. Hence, we take the average over all name parts with non-missing values.

6.B Testing for differences in demand across racial groups

One alternative explanation for the racial disparities in loan outcomes is differences in demand across racial groups (Bhutta and Hizmo, 2021; Willen and Zhang, 2022). If minorities and Chinese have different preferences, then offers received by minorities and Chinese could differ, not because of lender's biases, but because lenders cater to differences in demand across racial groups. Ideally, for comparing differences in demand, we would like to keep the set of offers identical for minority and Chinese applicants, so that any differences in their choices can be attributed to differences in their demand. Our analyses reveal that with names on applications, minority applicants have lower access to credit in both the extensive and intensive margins than otherwise identical Chinese applicants. As a result, estimating the potential differences in demand in the period prior to anonymization would be affected by problematic selection issues. The disparities are substantially reduced after anonymization, we therefore opt to test for potential differences in demand using the sample in the post-anonymization period to sidestep the selection problems.

We test potential differences in demand between racial groups in Table 6. We use specifications adapted from equation (1): As we use the sample in the post-anonymization period, the $s(t) \in \{pre, post\}$ subscript in lender fixed effects $\alpha_{j,s(t)}$ and the impacts of application-level controls $\gamma_{s(t)}$ are dropped. We measure demand by an indicator variable that equals 1 if applicant i chooses the offer by lender j, and 0 otherwise. As before, we use the indicator multiplied by 100 as the outcome variable in regressions to facilitate a percentage point interpretation of the estimated coefficients. In Column 1, we report the average relationship between credit demand and minority status. We find that conditional on the same set of application-level controls, minority and Chinese applicants are not statistically different in their probability of choosing an offer. In Column 2, we additionally examine the effects of initial offer characteristics. The estimates suggest that an average applicant is more likely to choose a higher loan amount and a longer maturity (likely due to reduced monthly payments all else equal). We find that effective interest rate plays a muted role in choosing amongst offers, consistent with existing evidence in consumer credit markets (Ponce et al., 2017). As in Column 1, the coefficient on minority status is statistically insignificant and economically small. In Column 3, we estimate the relationship between credit demand and initial offer characteristics using within-application variation by including application fixed effects. In this specification, the minority indicator is fully absorbed by the included application fixed effects. Despite using different levels of variation, we obtain similar estimates as in Column 2: A higher loan amount and a longer maturity are associated with a higher likelihood of choosing an offer, whereas a lower effective interest rate does not make a difference. Lastly, in Column 4, we allow for the relationship between credit demand and loan offer characteristics to differ by applicant race. All interaction terms between loan offer characteristics and applicant race are insignificant, suggesting that there are no discernible differences in demand between applicants of different races.

6.C Relevance of offer probability

The loan offer probability is analogous to the callback rate for job applications used in labor market correspondence studies. We further demonstrate the relevance of offer probability for the credit market setting in this subsection. First, the large within-variation in offer terms (Table IA.1 in the Internet Appendix) suggests that not all lenders assess an applicant similarly, and hence more offers are valuable. Second, to address the concern that the "best" offer might be the same for minority and Chinese applicants, we consider three different definitions of the best offer: (1) maximum offer amount, (2) maximum offer maturity, and (3) minimum effective interest rate across all lenders. These three definitions have intuitive economic meanings. The maximum offer amount gives the applicant the highest amount available for personal use. The maximum offer maturity corresponds to the lowest monthly debt burden, all else equal, as the longer the maturity, the lower the monthly payment. The minimum effective interest rate is the lowest-cost loan offer. We do not ascribe to a particular definition of what constitutes a best offer, instead, we examine all three intuitive definitions of best offer. In the previous subsection, we verify that applicants are less likely to choose offers with a lower offer amount or a shorter offer maturity are inde

Table 7 documents the results of this analysis. Column 1 focuses on the log of maximum offer amount as the outcome variable. Before anonymization, minority applicants' maximum offer amount is 10.5% lower than Chinese applicants'. However, after anonymization, minority applicants' maximum offer amount is 2.1% lower than Chinese applicants'. These two coefficients are both statistically and economically different. Column 2 reports the estimates for the analysis of the maximum offer maturity. While the maximum offer maturity is 0.7 months shorter for minority applicants than for Chinese applicants when names are shown on applications before anonymization, the difference becomes statistically and economically insignificant in the post period. Column 3 reports the estimates for the analysis of the minimum effective interest rate.

Across these definitions of best offers, minority applicants receive worse offers than Chinese applicants when applicant racial identities are shown to lenders. Such racial disparities in the intensive margin are significantly reduced after anonymization.

6.D Other robustness checks

We also assess the robustness of our results to alternative samples, specifications, and measurement in Internet Appendix IA.5. We obtain similar estimates when we relax the sample filtering, when we estimate a more flexible specification to accommodate for lender-specific changes in lending practices, and when we use alternative measures of applicant race as in our baseline results.

One might be concerned that the matching between applicants and lenders is correlated with application characteristics and hence may affect our estimates. In our setting, although the matching between applicants and lenders is not completely random, the workflow of the platform makes endogenous matching unlikely. Specifically, the platform orders lenders 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 lenders 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 versus Chinese applicants in Appendix IA.1 and further verify that the ordering is identical for Chinese and minority applicants.

Anonymization removes information that lenders could potentially use in making their lending decisions, one may be concerned that such a loss of information may reduce overall credit supply and therefore the reduction of racial disparities may not represent an increase of access for minority applicants. To test this possibility, we compare lenders who exhibit different degrees of racial bias (using β_{pre} in Figure 2) in their offer probability in Internet Appendix IA.6. Compared to unbiased lenders, who presumably do not use the race information to make lending decisions prior to anonymization, the biased lenders do not significantly reduce their credit supply in both the extensive and intensive margins. The results suggest that anonymous applications do not affect overall credit supply.

7 Loan origination and loan performance

After anonymization, when names are removed from loan applications, lenders do not know the racial identity of applicants at the time of initial evaluation. The racial identity is revealed to them once applicants visit them in person to fulfill the required verification procedures. Lenders can use this new information to fully undo their initial decisions.

Column 1 of Table 8 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 disparities become 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. We find that anonymization is associated with a 10.6% decrease in racial disparities in offer probability (Column 1 of Table 2). The corresponding decrease in racial disparities in origination rate is 8.0% of its sample average according to the estimates from Column 1 of Table 8. The reduction of racial disparities in the initial offer stage is attenuated by approximately 20% once the applicants advance to the origination stage, suggesting that lenders only partially adjust their behaviors at in-person interactions with applicants. The point estimates in Column 1 of Table 8 are mechanically smaller than point estimates in Column 1 of Table 2 because the number of loan offers for an average applicant is 7.6, whereas there is a maximum of one origination loan per applicant.

We obtain a similar magnitude in the application-level analysis in Column 2. Comparing the economic magnitude from the estimates in Column 3 of Table 2 and Column 2 of Table 8 reveals an approximately 15% of lenders' partial adjustment in this level of analysis.

Finally, we examine the relationship between race and loan performance for a sub-sample of originated loans. The data come from one lender that originates approximately 14% of the loans. Column 3 of Table 8 reports the regression analysis in the loan performance sub-sample. As the sub-sample comes from one lender, the usual Lender×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 3 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.

8 Discussion

8.A Which theory of discrimination best describes our findings?

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 probability are substantially reduced, further corroborating the existence of discrimination.

Anonymization can be effective regardless of which theory of discrimination leads to the existence of discrimination in the first place. Nevertheless, we examine two leading theories of discrimination — taste-based and statistical discrimination — to determine which aligns more closely with our findings. Under taste-based discrimination (Becker, 1957), differential treatment stems from the disutility of providing service to or interacting with members of a particular group. Similar to prejudice, inaccurate beliefs can also give rise to biased behaviors (Bohren et al., 2023). 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.

Two pieces of evidence show that statistical discrimination cannot explain our results. First, in Appendix IA.4 (as well as Tables 4 and IA.2), we find similar racial gaps for different quantiles of income and income-to-debt ratios in the pre-period when names are shown to lenders at initial evaluation, suggesting that repayment ability does not affect racial disparities. Second, we use a simple model of lender decision-making under

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

statistical discrimination following Agan and Starr (2018) and use the observed lender decisions to test whether lenders' beliefs are accurate. We find that lender's beliefs are inaccurate. The next subsections provide the details of these tests.

8.B Accuracy of lender beliefs

One explanation for the existence of racial disparities is inaccurate beliefs held by decision-makers (Bohren et al., 2023). While we do not observe beliefs held by lenders in our setting, we follow (Agan and Starr, 2018) and use a simple model of lender decision-making under accurate statistical discrimination. Using the observed choices by lenders before and after anonymization, the model allows us to test whether lenders' beliefs are accurate.

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 minority applicants and m = 0 for Chinese applicants, x_i is a vector of borrower characteristics, $\varepsilon_{i,j}$ 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

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

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

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

Using Equation (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 (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 3 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.[9] Hence, we conclude that accurate beliefs by race are not supported by the data. In Internet Appendix IA.7, we specifically test for stereotypes (Bordalo et al., 2016) and find that the data do not support such an interpretation.

8.C Predictability of the race information

If race is correlated with the remaining application characteristics (Fuster et al., 2022), lenders may be able to infer race when loan applications are anonymized. Here, we directly assess the predictive power of these application characteristics in identifying race.

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 4 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 equal to one if that characteristic applies to the application 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 non-zero correlation with race.

Panel (B) of Figure 4 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, with (1) one layer of 100 neurons, (2) three layers with 50 neurons each, and (3) one layer of 200 neurons. In addition,

[9] 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.

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 et al., 2016; Berg et al., 2020). Predictability of race information from application characteristics has been documented in other settings, as well (Fuster et al., 2022).

Overall, sophisticated statistical models can predict race using the remaining application characteristics with a high degree of accuracy, consistent with the notion of *triangulation* whereby sophisticated prediction models can effectively "de-anonymize" group identities (Pope and Sydnor, 2011a; Fuster et al., 2022). An implication of the high predictability of race by machine learning is that algorithm-based decisions may generate disparate impacts across racial groups, even if race is explicitly excluded from the training set (Blattner and Nelson, 2021; Gao et al., 2023).

8.D In-group preferences

Cornell and Welch (1996), Fisman et al. (2017), and D'Acunto et al. (2023) 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 lenders 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 Chinesesounding names. Hence, our data do not have sufficient variation for testing in-group preferences.

9 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 substantially reduced such disparities. Furthermore, high-income minorities benefited more from anonymization, suggesting increased lender reliance on objective credit risk measures. Heterogeneity analyses and analysis of lender beliefs show that our results are inconsistent with statistical discrimination.

With the advent and expansion of FinTech credit, the implementation of anonymous applications has become 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 the Singapore government has implemented successful policies in promoting racial equity (Agarwal et al., 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 et al., 2018), complete anonymization may become feasible for FinTech consumer credit and help minority borrowers even more. In this context, mandatory information collection of race in regulations designed to detect

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

discriminatory lending practices, such as the Home Mortgage Disclosure Act (HMDA) in the US, may constrain the effective implementation of anonymization.

In our setting, as is common in labor, rental, and credit markets, only the first stage of the evaluation process is made anonymous. Although lenders observe applicant race through in-person interaction with applicants prior to loan origination, racial disparities in the loan origination stage are also reduced. Overall, we find that anonymous loan applications are effective in reducing racial disparities, even though they merely delay revealing information about race.

More generally, the distributional consequences of policies that restrict the use of race-predictive information are ambiguous. In our setting, anonymous loan applications reduce racial disparities by increasing lender reliance on objective credit risk measures. In other settings, similar policies have been documented to hurt racial minorities due to biased beliefs (Agan and Starr, 2018) or reduced signal precision (Bartik and Nelson, 2020). We believe that much more work needs to be done to understand the interplay of different mechanisms in driving the equilibrium racial disparities.

References

Agan, A., and S. Starr. 2018. Ban the box, criminal records, and racial discrimination: A field experiment. *Quarterly Journal of Economics* 133:191–235.

Agarwal, S., H. S. Choi, J. He, and T. F. Sing. 2019. Matching in housing markets: The role of ethnic social networks. *Review of Financial Studies* 32:3958–4004.

Aigner, D. J., and G. G. Cain. 1977. Statistical theories of discrimination in labor markets. *Industrial and Labor Relations Review* 30:175–187.

Arrow, K. J. 1973. The theory of discrimination. In O. Ashenfelter and A. Rees (eds.), *Racial Discrimination in Economic Life*, pp. 3–33. Princeton University Press.

Bartik, A. W., and S. T. Nelson. 2020. Deleting a signal: Evidence from pre-employment credit checks.

Bartlett, R., A. Morse, R. Stanton, and N. Wallace. 2022. Consumer-lending discrimination in the FinTech Era. *Journal of Financial Economics* 143:30–56.

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

Behaghel, L., B. Crépon, and T. Le Barbanchon. 2015. Unintended effects of anonymous résumés. *American Economic Journal: Applied Economics* 7:1–27.

Berg, T., V. Burg, A. Gombović, and M. Puri. 2020. On the rise of FinTechs: Credit scoring using digital footprints. *Review of Financial Studies* 33:2845–2897.

Berg, T., A. Fuster, and M. Puri. 2022. FinTech Lending. *Annual Review of Financial Economics* 14:187–207.

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

Bertrand, M., and S. Mullainathan. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review* 94:193–204.

Bhutta, N., and A. Hizmo. 2021. Do minorities pay more for mortgages? *Review of Financial Studies* 34:763–789.

Bhutta, N., A. Hizmo, and D. Ringo. 2022. How much does racial bias affect mortgage lending? Evidence from human and algorithmic credit decisions.

Blattner, L., and S. Nelson. 2021. How costly is noise? Data and disparities in consumer credit.

Bohren, J. A., K. Haggag, A. Imas, and D. G. Pope. 2023. Inaccurate statistical discrimination: An identification problem. *Review of Economics and Statistics* pp. 1–45.

Bordalo, P., K. Coffman, N. Gennaioli, and A. Shleifer. 2016. Stereotypes. *Quarterly Journal of Economics* 131:1753–1794.

Buchak, G., G. Matvos, T. Piskorski, and A. Seru. 2018. Fintech, regulatory arbitrage, and the rise of shadow banks. *Journal of Financial Economics* 130:453–483.

Butler, A. W., E. J. Mayer, and J. P. Weston. 2022. Racial disparities in the auto loan market. *Review of Financial Studies* 36:1–41.

Cornell, B., and I. Welch. 1996. Culture, information, and screening discrimination. *Journal of Political Economy* 104:542–571.

D'Acunto, F., P. Ghosh, and A. G. Rossi. 2023. How costly are cultural biases? Evidence from FinTech.

Dobbie, W., A. Liberman, D. Paravisini, and V. Pathania. 2021. Measuring bias in consumer lending. *Review of Economic Studies* 88:2799–2832.

Fisman, R., D. Paravisini, and V. Vig. 2017. Cultural proximity and loan outcomes. *American Economic Review* 107:457–492.

Frame, W. S., R. Huang, E. J. Mayer, and A. Sunderam. 2022. The impact of minority representation at mortgage lenders.

Fryer, R. G., and S. D. Levitt. 2004. The causes and consequences of distinctively Black names. *Quarterly Journal of Economics* 119:767–805.

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. *The Review of Financial Studies* 32:1854–1899.

Gao, J., H. Yi, and D. H. Zhang. 2023. The role of human underwriting in the Big Data Era.

Goldin, C., and C. Rouse. 2000. Orchestrating Impartiality: The Impact of "Blind" Auditions on Female Musicians. *American Economic Review* 90.

Hanson, A., Z. Hawley, H. Martin, and B. Liu. 2016. Discrimination in mortgage lending: Evidence from a correspondence experiment. *Journal of Urban Economics* 92:48–65.

Howell, S. T., T. Kuchler, D. Snitkof, J. Stroebel, and J. Wong. 2021. Racial disparities in access to small business credit: Evidence from the Paycheck Protection Program.

Iyer, R., A. I. Khwaja, E. F. Luttmer, and K. Shue. 2016. Screening peers softly: Inferring the quality of small borrowers. *Management Science* 62:1554–1577.

Jiang, E. X., Y. Lee, and W. S. Liu. 2022. Reducing racial disparities in consumer credit: The role of loan minority officers in the era of algorithmic underwriting.

Kline, P., E. K. Rose, and C. R. Walters. 2022. Systemic discrimination among large U.S. employers. *Quarterly Journal of Economics* 137:1–74.

Krause, A., U. Rinne, and K. F. Zimmermann. 2012. Anonymous job applications of fresh Ph.D. economists. *Economics Letters* 117:441–444.

Phelps, E. S. 1972. The statistical theory of racism and sexism. *American Economic Review* 62:659–661.

Ponce, A., E. Seira, and G. Zamarripa. 2017. Borrowing on the wrong credit card? Evidence from Mexico. *American Economic Review* 107:1335–1361.

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

Pope, D. G., and J. R. Sydnor. 2011b. What's in a picture? Evidence of discrimination from prosper.com. *Journal of Human Resources* 46:53–92.

Willen, P., and D. Zhang. 2022. Testing for discrimination in menus.

Wolpert, D. H. 1992. Stacked generalization. *Neural Networks* 5:241–259.

Wong, M. 2013. Estimating ethnic preferences using ethnic housing quotas in Singapore. *Review of Economic Studies* 80:1178–1214.

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

This figure plots the entire path of coefficients β_s along with their associated 95% confidence intervals of the racial gap in offer probability 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. A negative $\beta_s < 0$ indicates a minority-Chinese gap (i.e., minority applicants are less likely to receive loan offers than Chinese applicants). The *x-axis* denotes the months before and after anonymization; the *y-axis* shows the change in the racial gap in offer probability relative to the omitted baseline period (in percentage points).

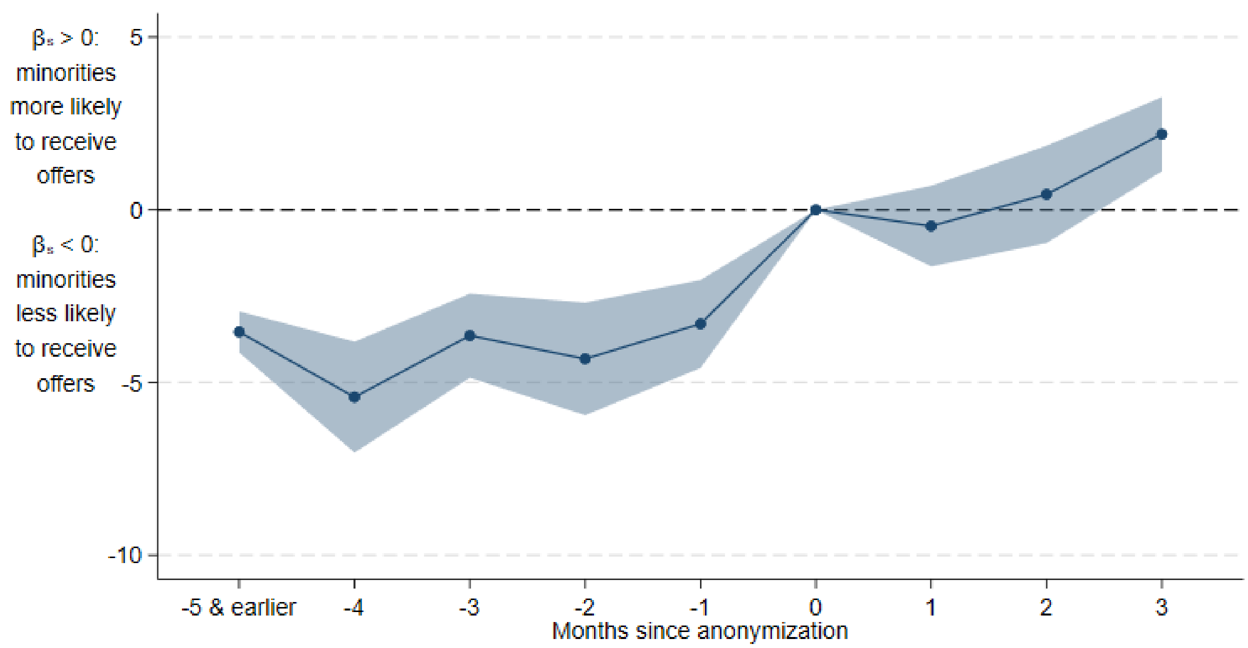

Figure 2: **Lender-specific racial disparities in offer probability**

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 probability. 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.

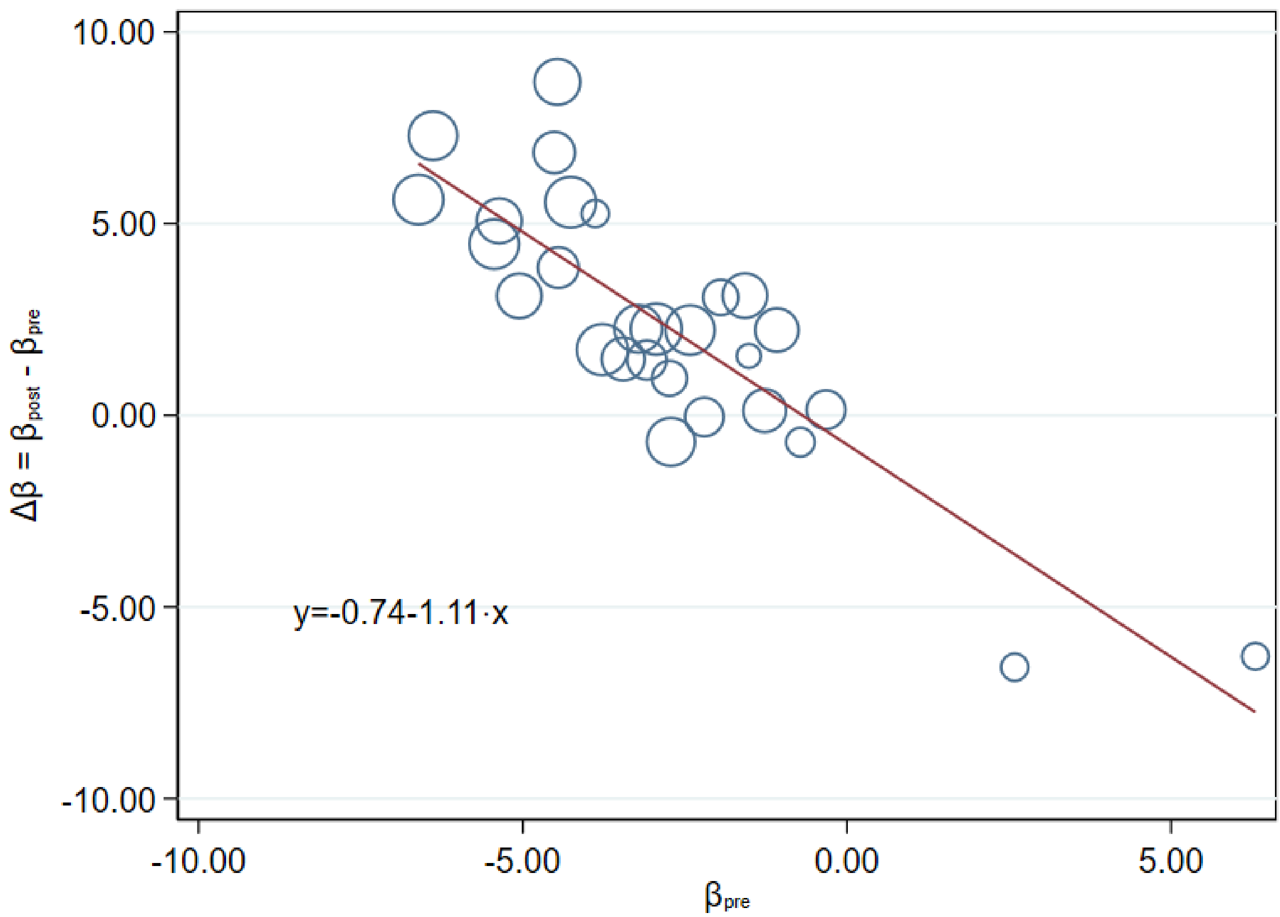

Figure 3: **Histogram of the difference between the inferred and the empirical probabilities**

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' priors are more inaccurate. Each point used in the plot corresponds to one application characteristic.

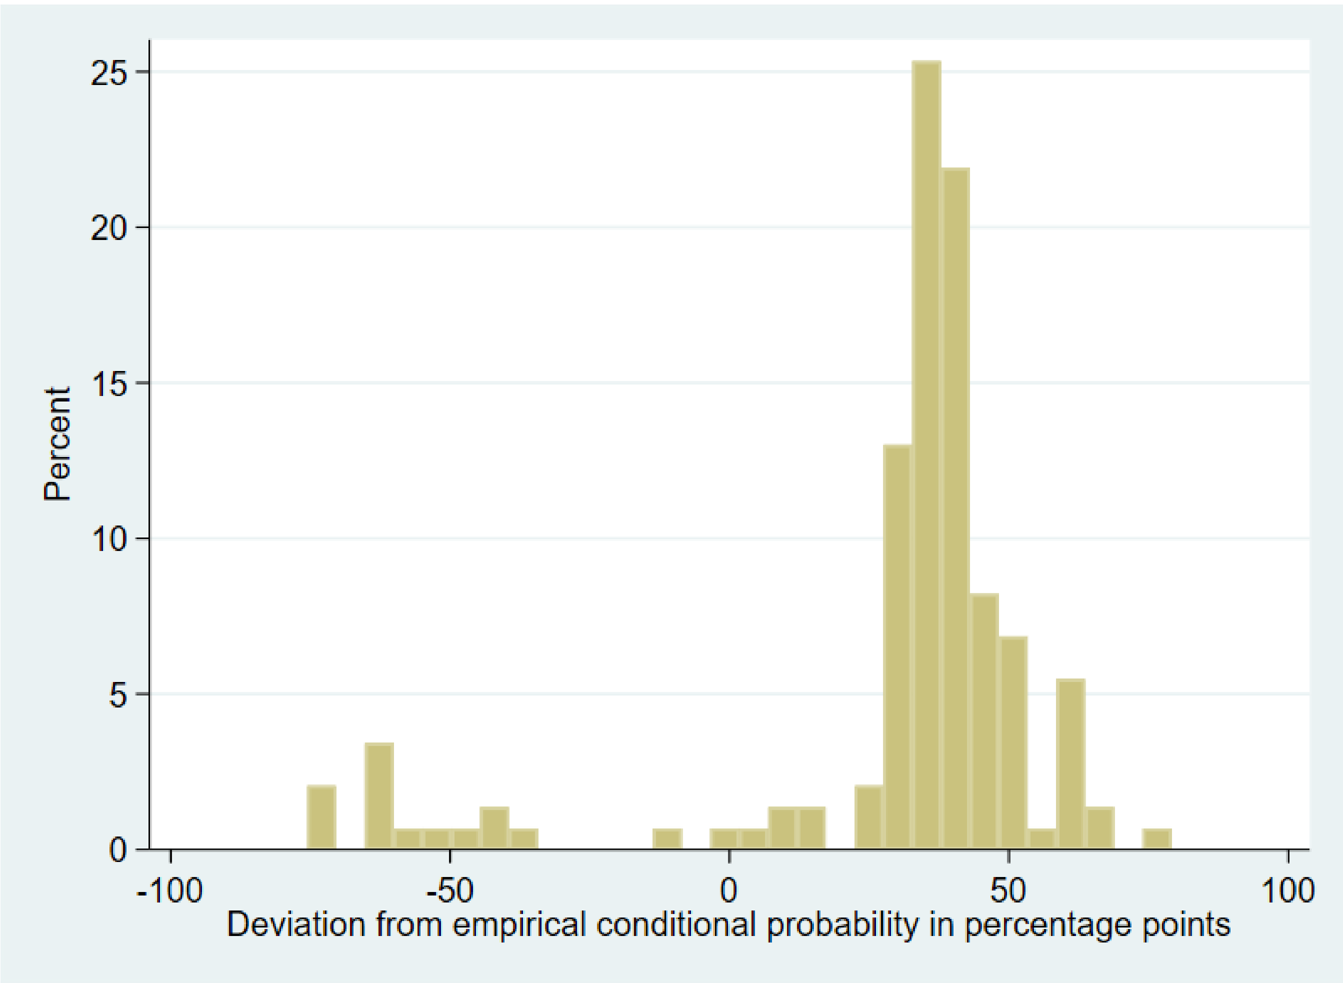

Figure 4: **Predictability of the race information**

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 the curve of various classification analyses for predicting race using other observable application characteristics.

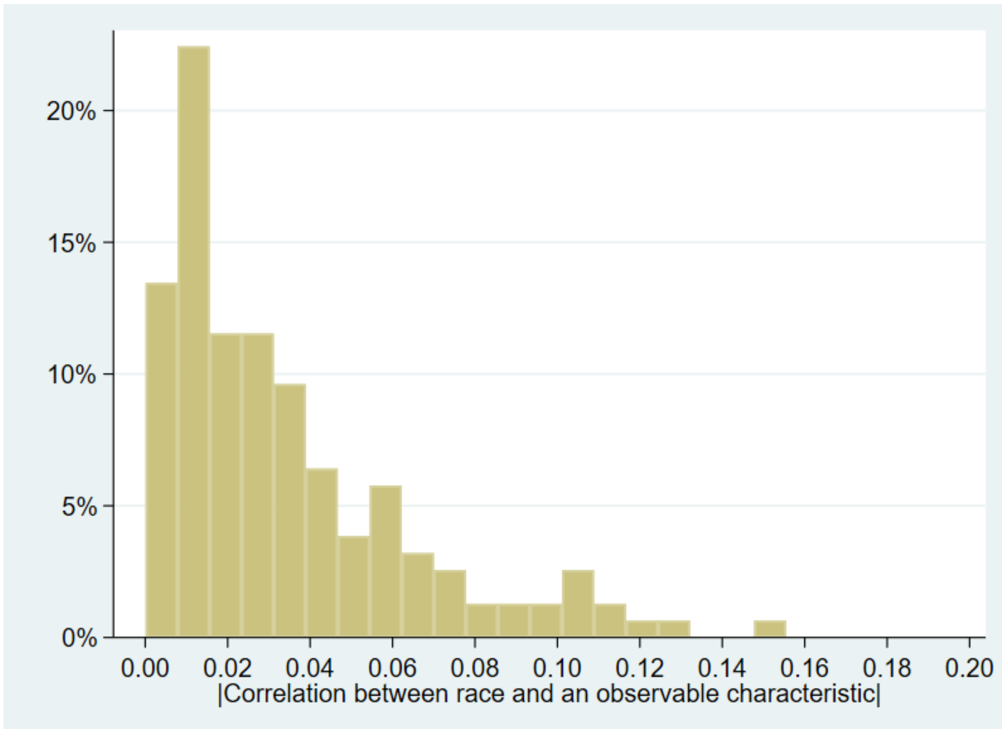

(A) Distribution of bi-variate correlations of race and other application characteristics

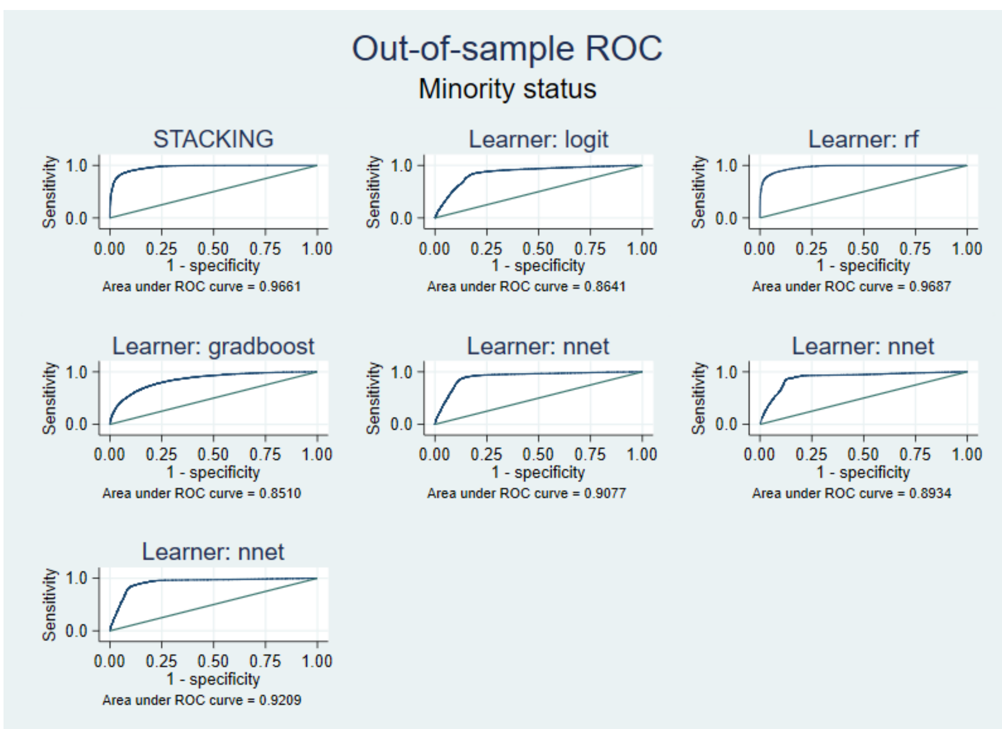

(B) Out-of-sample area under curve of classification analyses

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 reports the mean differences between minority and Chinese applicants in the pre-period (before September 28, 2021), when applicant names were visible to lenders at the initial evaluation stage (non-anonymous applications). Column 3 of both panels reports the mean differences in the post-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 probability (%)	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 probability (%)	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 probability**

This table shows the effect of anonymous applications on disparities in offer probability. We estimate the minority-Chinese gap in offer probability before and after anonymization. In our specifications, we include all application characteristics observable 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. We also include fixed effects suitable for the level of analysis to control for invariant confounding factors; the included fixed effects are denoted at the bottom. For the application-lender level analysis in Columns 1 & 2, we include year-month fixed effects and lender fixed effects separately for the pre- and post-periods. For the application-level analysis in Column 3, we include year-month fixed effects and the Post indicator. 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 probability (%)
	(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]
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 initial offer terms**

This table shows the effect of anonymous applications on disparities in the intensive margins of initial offers. We estimate the minority-Chinese gap in the intensive margins of initial offers before and after anonymization. In our specifications, we include all application characteristics observable 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 functional 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. We also include year-month fixed effects and lender fixed effects separately for the pre- and post-periods to control for invariant confounding factors; the included fixed effects are denoted at the bottom. 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.

	(1)	(2)	(3)
	Log(offer amount)	Offer maturity (months)	Effective interest rate (%)
Minority × Pre	-0.0480***	-0.147***	0.839***
	[-11.31]	[-6.78]	[5.34]
Minority × Post	-0.0267***	-0.0898**	0.619**
	[-4.32]	[-2.53]	[2.50]
Year-Month FEs	Yes	Yes	Yes
Lender × Post FEs	Yes	Yes	Yes
Observable controls	Yes	Yes	Yes
R^2	0.614	0.627	0.723
No. of observations	123,300	123,300	123,300

Table 4: **Heterogeneous racial gaps across income groups**

This table shows the effect of anonymous applications on disparities in credit outcomes for low- and high-income minorities. We estimate the income-group specific minority-Chinese gap in offer probability before and after anonymization. High and low income are dummy variables equal to one for applicants with above and below median income. In our specifications, we include all application characteristics observable 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 functional 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. We also include year-month fixed effects and lender fixed effects separately for the pre- and post-periods to control for invariant confounding factors; the included fixed effects are denoted at the bottom. 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.

	(1) Offer indicator (× 100)	(2) Log(offer amount)	(3) Offer maturity (months)	(4) Effective interest rate (%)
Minority × Pre × High income	-3.708***	-0.040***	-0.157***	0.739***
	[-12.88]	[-8.19]	[-5.50]	[3.86]
Minority × Pre × Low income	-3.910***	-0.058***	-0.135***	0.969***
	[-14.74]	[-9.25]	[-5.26]	[4.54]
Minority × Post × High income	1.700***	-0.008	-0.048	0.468
	[4.09]	[-1.16]	[-1.10]	[1.52]
Minority × Post × Low income	-1.062***	-0.050***	-0.144***	0.817***
	[-3.03]	[-6.60]	[-3.33]	[2.64]
Year-Month FEs	Yes	Yes	Yes	Yes
Lender × Post FEs	Yes	Yes	Yes	Yes
Observable controls	Yes	Yes	Yes	Yes
R^2	0.306	0.614	0.627	0.723
No. of observations	322,847	123,300	123,300	123,300

Table 5: **Testing for social status signals in names**

This table tests for potential social status signals conveyed in names. We estimate the minority-Chinese gap in offer probability before and after anonymization with and without controlling for social status signals in names. To construct a measure of social status for each applicant, we split each full name into separate name parts, and use a leave-one-out approach to calculate the mean income for each name part using the entire sample. For each full name, we then take the average of this variable across different name parts and log transform the outcome. In our specifications, we include all application characteristics observable 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 functional 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. We also include fixed effects suitable for the level of analysis to control for invariant confounding factors; the included fixed effects are denoted at the bottom. For the application-lender level analysis in Columns 1 & 2, we include year-month fixed effects and lender fixed effects separately for the pre- and post-periods. For the application-level analysis in Columns 3 & 4, we include year-month fixed effects and the Post indicator. 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.

	Application-lender level offer indicator ($\times$ 100)		Application level average offer probability (%)	
	(1)	(2)	(3)	(4)
Minority $\times$ Pre	-3.810***	-3.951***	-3.969***	-4.106***
	[-16.78]	[-16.90]	[-7.93]	[-7.68]
Minority $\times$ Post	0.238	0.147	-0.434	-0.479
	[0.88]	[0.50]	[-0.58]	[-0.52]
Controlling for social status	No	Yes	No	Yes
Year-Month FEs	Yes	Yes	Yes	Yes
Lender $\times$ Post FEs	Yes	Yes	-	-
Post FE			Yes	Yes
Observable controls	Yes	Yes	Yes	Yes
R^2	0.305	0.306	0.569	0.569
No. of observations	322,847	322,847	16,281	16,281

Table 6: **Testing for differences in demand**

This table tests for potential differences in demand. We use the sample in the post-anonymization period to sidestep problematic selection issues stemming from the differences in offers that applicants receive. In Columns 1 & 2, we include all application characteristics observable to lenders at the time of initial screening as control variables. In choosing the functional 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. In Columns 3 & 4, we include application fixed effects to use within-application variation for estimation. In this specification, the minority indicator and other application characteristics are fully absorbed by the included application fixed effects. In all specifications, we also include year-month fixed effects and lender fixed effects to control for invariant confounding factors; the included fixed effects are denoted at the bottom. 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 indicator for choosing loan offer ($\times$ 100)			
	(1)	(2)	(3)	(4)
Minority	0.257	0.376		
	[0.71]	[1.03]		
Log(offer amount)		1.149**	3.059***	3.421***
		[2.25]	[4.92]	[4.49]
Log(offer amount) $\times$ Minority				-0.539
				[-0.68]
Offer maturity (months)		1.104***	1.139***	1.025***
		[10.51]	[9.06]	[8.19]
Offer maturity (months) $\times$ Minority				0.195
				[1.40]
Effective interest rate (%)		0.0167	-0.00725	-0.00330
		[1.33]	[-0.50]	[-0.20]
Effective interest rate (%) $\times$ Minority				-0.00530
				[-0.28]
Application FEs	No	No	Yes	Yes
Year-Month FEs	Yes	Yes	Yes	Yes
Lender FEs	Yes	Yes	Yes	Yes
Observable controls	Yes	Yes	-	-
R^2	0.0622	0.0749	0.231	0.231
No. of observations	31,086	31,086	31,086	31,086

Table 7: **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. We estimate the minority-Chinese gap in best offer terms before and after anonymization. In our specifications, we include all application characteristics observable 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 functional 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. We also include year-month fixed effects and the Post indicator to control for invariant confounding factors; the included fixed effects are denoted at the bottom. 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]
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 8: **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. We estimate the minority-Chinese gap in other credit outcomes before and after anonymization. In our specifications, we include all application characteristics observable 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 3, 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 functional 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. We also include fixed effects suitable for the level of analysis to control for invariant confounding factors; the included fixed effects are denoted at the bottom. For the application-lender level analysis in Column 1, we include year-month fixed effects and lender fixed effects separately for the pre- and post-periods. For the application-level analysis in Columns 2 & 3, we include year-month fixed effects and the Post indicator. 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 ($\times$ 100)	(2) Application level origination indicator ($\times$ 100)	(3) Delinquency indicator ($\times$ 100)
Minority $\times$ Pre	-0.0910*	-1.598*	0.151
	[-1.68]	[-1.89]	[0.03]
Minority $\times$ Post	-0.0236	-0.500	-0.319
	[-0.29]	[-0.33]	[-0.04]
Year-Month FEs	Yes	Yes	Yes
Lender $\times$ Post FEs	Yes	-	-
Post FE		Yes	Yes
Observable controls	Yes	Yes	Yes
R^2	0.00792	0.0677	0.403
No. of observations	322,847	16,281	373

Internet Appendix for " Reducing Racial Disparities in Consumer Credit: Evidence from Anonymous Loan Applications "

This Internet Appendix contains supplementary material, tables, and figures.

IA.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 lenders and send out applications to lenders 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 IA.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 the 5% level, and the other 5 coefficients are statistically significant at the 5% level. Furthermore, the estimated magnitudes are economically small. On average, the absolute value of the estimated coefficient is 0.16. lender rank can be any number between 1 to 36. Hence, even if the coefficients were statistically significant, their corresponding change in lender ordering would be less than 0.16. Hence, the evidence suggests that the platform does not match applications to lenders based on application race.

IA.2 Identifying race from applicant name

In Singapore, names tend to map uniquely to race. That is, it is relatively straightforward to accurately identify applicants' race from names. We support this claim in two ways by using race from Singapore government records obtained by the platform at the time of application. First, we cross-check the race assigned by research assistants with government records and find a close match; only the race of 1.6% of applicants is misclassified.

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

Second, we break each full name into its parts separated by space (for instance, an applicant whose name is "John Doe" will be coded as having two name parts "John" and "Doe".) and then count the percentage of observations for which a name part is used amongst minorities, where minority status is identified using government records. In an extreme case where names are not informative about race, we expect all name parts in the data to have a 50% value. That is, each name part is equally likely to belong to Chinese and minorities. In another extreme case where names fully separate minorities and Chinese, we expect to have a mass of name parts on 0% (these are name parts only used by Chinese) and another mass on 100% (these are name parts only used by minorities), and no mass in between.

Figure IA.2 plots the histogram of the percentage of each name part belonging to minorities. Only 5.1% of names are shared between minorities and Chinese, 20.3% are only used by Chinese, and 74.6% are only used by minorities. Hence, in our settings, the sample includes racially distinctive names commonly used in correspondence studies.

IA.3 Additional results of offer terms

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

Panel (A) of Figure IA.3 plots the empirical cumulative distribution function (CDF) of annual interest rates for all initial offers. The figure shows 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, set by the Singapore Ministry of Law and effective from 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 IA.3 plots the empirical CDF of the processing fee as a percentage of the loan offer amount. Even if lenders cannot 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, set by the Singapore Ministry of Law and effective from 1 October 2015, is 10% of the loan offer amount.

Taken together, the panels of Figures IA.3 suggest that 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 lenders have full discretion.

Table IA.1 documents the dispersion within the same application across three offer terms: offer amount, maturity, and effective interest rate. We use two different measures of dispersion, namely the 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 that 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 probability is indeed a relevant outcome variable for studying racial disparities in lending.

IA.4 Additional results of heterogeneity across applicant characteristics

We present additional results of heterogeneity across applicant characteristics in this appendix subsection.

In Table IA.2, we use an alternative measure of objective qualification, the income-to-debt ratio calculated as the annual income divided by the applied amount. We group applicants into high and low income-to-debt groups using the median as the cutoff and estimate equation (3). As in the previous split by income level in Table 4, we find that high income-to-debt minority applicants benefit more from anonymous loan applications than low income-to-debt minority applicants.

One pattern revealed in the analysis of heterogeneity across applicant characteristics is that minority applicants with above- or below-median income or income-to-debt levels receive similar credit access in the pre-period when names are shown in the loan applications. We further examine the robustness of this pattern using different quantiles to group applicants and separate estimation for different groups.

We first group applicants into four quartiles based on 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 IA.4. We find similar racial disparities across income groups in the pre-period when names are shown in the loan applications.

Alternatively, we group applicants into four quartiles based on their annual income divided 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 IA.4. As in the previous split by income level, we find similar racial disparities across groups with different income-to-debt ratios in the pre-period when names are shown in the loan applications.

IA.5 Alternative samples, specifications, and measurement

We further examine the robustness of our results to alternative samples, specifications, and measurement.

For our main analysis, we focus on the sub-sample of applications whose information is pre-filled directly from the Singapore government database. We do so because official records have higher data quality and less measurement error compared to self-reported information. Nevertheless, we repeat the analysis for the full sample of applications in Column 1 of Table IA.3. We find results similar to the main sample. Hence, our results are not sensitive to the filtering of applications.

In Column 2 of Table IA.3, we fix the coefficients of observable characteristics to be the same in the pre- and post-periods. In other words, the corresponding coefficients γ do not have the $s(t)$ subscript. We obtain similar findings as in our baseline results: The racial gap is substantially reduced by the anonymization.

In Column 3 of Table IA.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.

$$\begin{aligned} y_{i,j} = \pi_t + \alpha_{j,s(t)} + \gamma_{j,s(t)} X_i \\ + \beta_{pre} \times Minority_i \times Pre_t + \beta_{post} \times Minority_i \times Post_t + \varepsilon_{i,j} \end{aligned} \tag{IA.1}$$

We obtain similar estimates as in our baseline results.

We also repeat our analysis using two alternative measures of race. The first measure we use comes from the race information in the government records. Such information is observable to us in the sub-sample of applications whose information is pre-filled directly from the Singapore government database, our main analysis sample. Crucially, even though the platform obtains such information, it does not include this information in the applications it sends to lenders throughout our sample period. Using this measure of applicant race, we repeat our analysis and report the estimates in Column 4 of Table IA.3. We obtain similar estimates as in our baseline results.

The second measure builds upon the official records of race. We analyze the implied race from name parts following a procedure detailed in Internet Appendix IA.2. Using this second alternative measure of applicant race, we repeat our analysis and report the estimates in Column 5 of Table IA.3. We obtain similar estimates as in our baseline results.

IA.6 Testing for the impact of anonymous applications on overall credit supply

Our analysis suggests that racial disparities in offer rates disappear after anonymization. However, anonymizing applications might also affect lender *overall* offers. To estimate the effect of anonymization on overall offer decisions, one would ideally compare lenders who are subject to anonymous applications (treated lenders) with other lenders who continue to receive name-bearing applications (control lenders). In our setting, anonymization affected all lenders simultaneously, making such a direct comparison infeasible. An alternative approach would be to use unbiased lenders as the counterfactual group. Assuming that unbiased lenders do not change their overall lending decisions after anonymization, presumably because they were not using race information to make lending decisions, we can use this group of lenders as the control group. This approach is feasible because, as Figure 2 shows, there is substantial heterogeneity in racial disparities in offer rates across lenders.

To implement this idea, we divide lenders into two groups based on the lender-specific racial gap in offer probability before anonymization as estimated in Section 5.B. We use the median racial gaps before anonymization across lenders as the cutoff to differentiate between biased and unbiased lenders. We then compare lenders in terms of the extensive and intensive margins of credit supply in the dyadic data on loan applications and lenders using the following specification:

$$y_{i,j} = \pi_t + \alpha_j + \gamma_{s(t)} X_i + \beta \times \text{Biased Lender}_j \times Post_t + \varepsilon_{i,j} \tag{IA.2}$$

Here, Biased Lender$_j$ is a dummy variable equal to one if the lender's racial gap in offer probability before anonymization is higher than the median. Controls and fixed effects are the same as equation (1), the only exception being we substitute $\alpha_{j,s(t)}$, pre and post-lender fixed effects, with lender fixed effects α_j. Standard errors are clustered at the lender-month level, the same as equation (1).

Table IA.4 shows the differences in offers between biased lenders with the rest after anonymization. The number of observations is slightly lower than our main regression tables as only lenders who have partnered with the platform before anonymization have well-defined bias measure and included in the analysis. Column (1) shows that the overall offer probability is not significantly affected by anonymization. Columns 2-4 show similar patterns on the intensive margins. Overall, these tests based on comparing biased and unbiased lenders suggest that anonymous applications do not affect overall credit supply.

IA.7 Stereotypes

We also test whether the inaccurate beliefs reflect stereotypes (Bordalo et al., 2016). In their model, the decision process based on Kahneman and Tversky's representativeness heuristic produces 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 IA.5 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.

Figure IA.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.

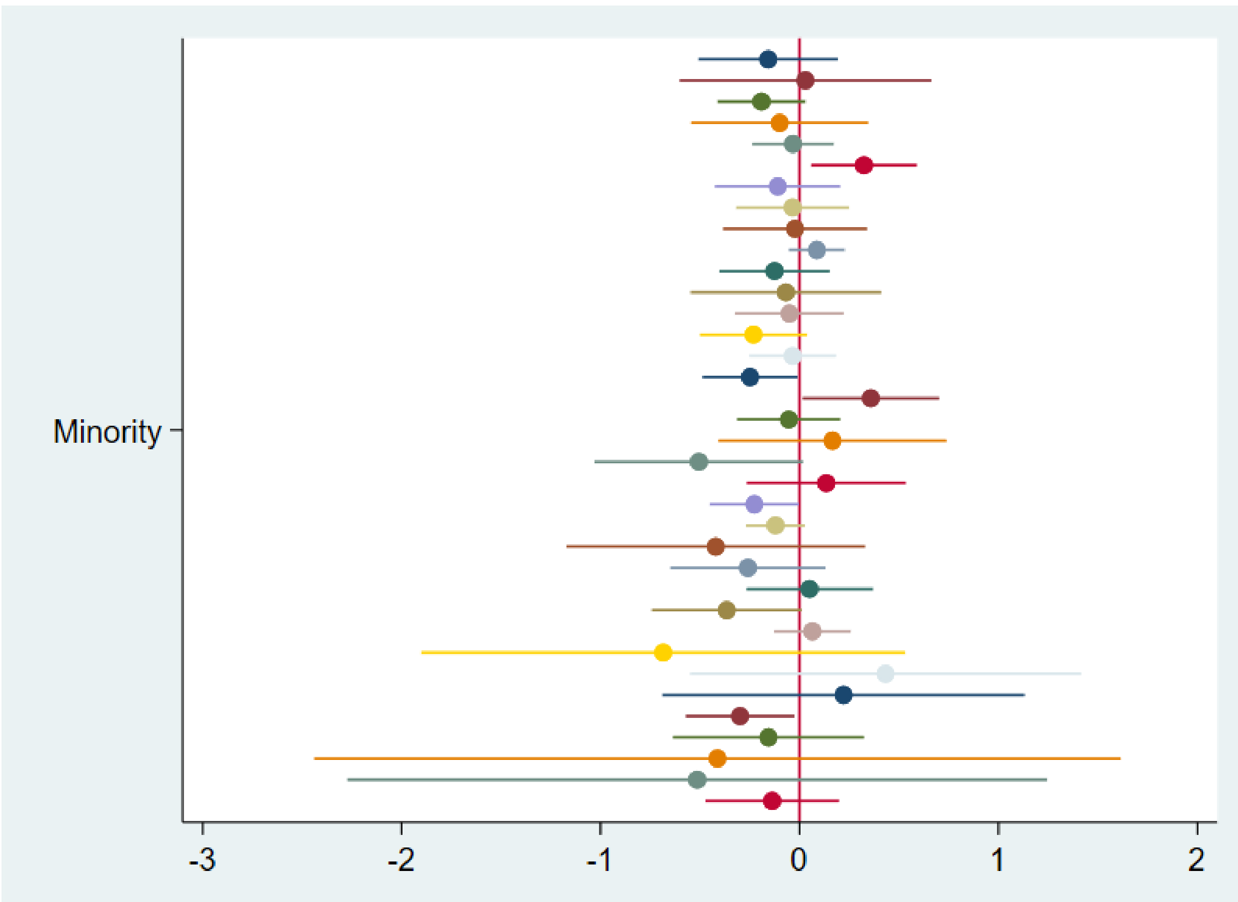

Figure IA.2: **Histogram of minority percentage by unique name part**

This figure is the histogram of minority percentage across unique name parts. For each name part, e.g., "John", we calculate the percentage of John in our dataset that are minority. 0% are name parts only used by Chinese, 100% are name parts only used by minorities, and anything in between is shared between minorities and Chinese. To identify minorities, we use race from government records.

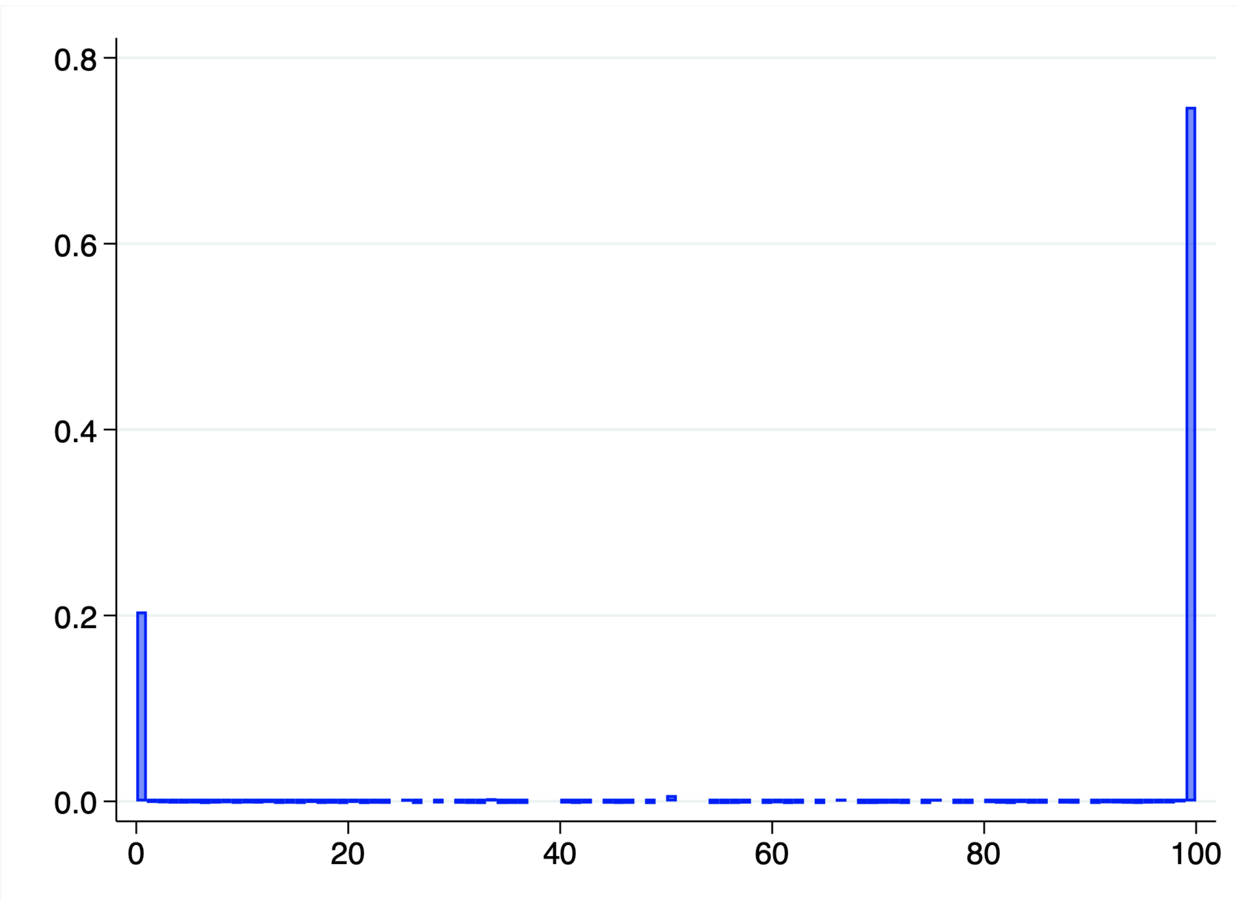

Figure IA.3: **Empirical distribution of offer interest rate and processing fee**

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.

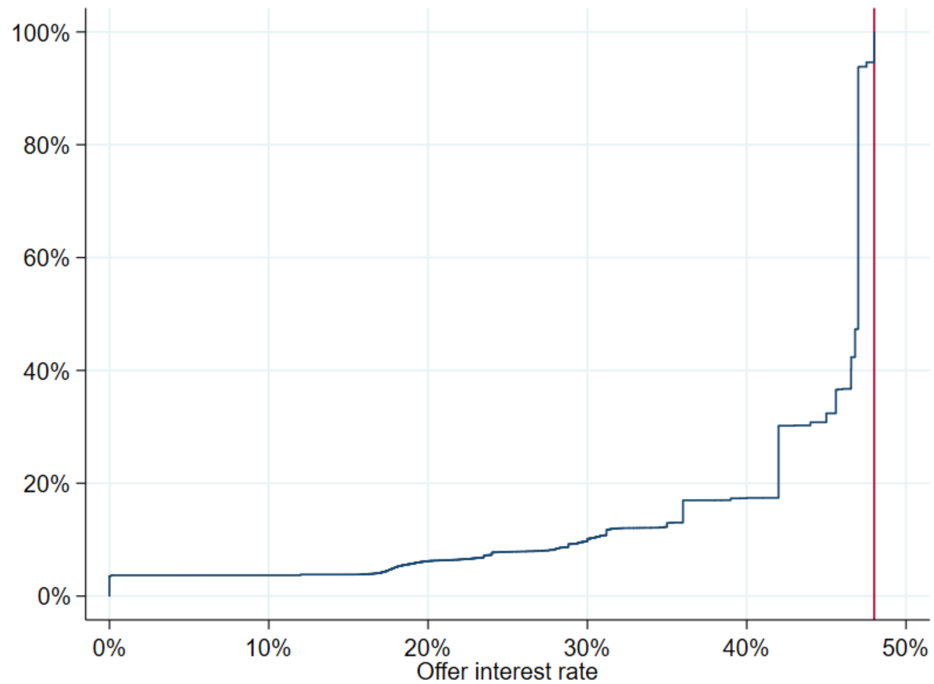

(A) Empirical distribution of offer interest rate

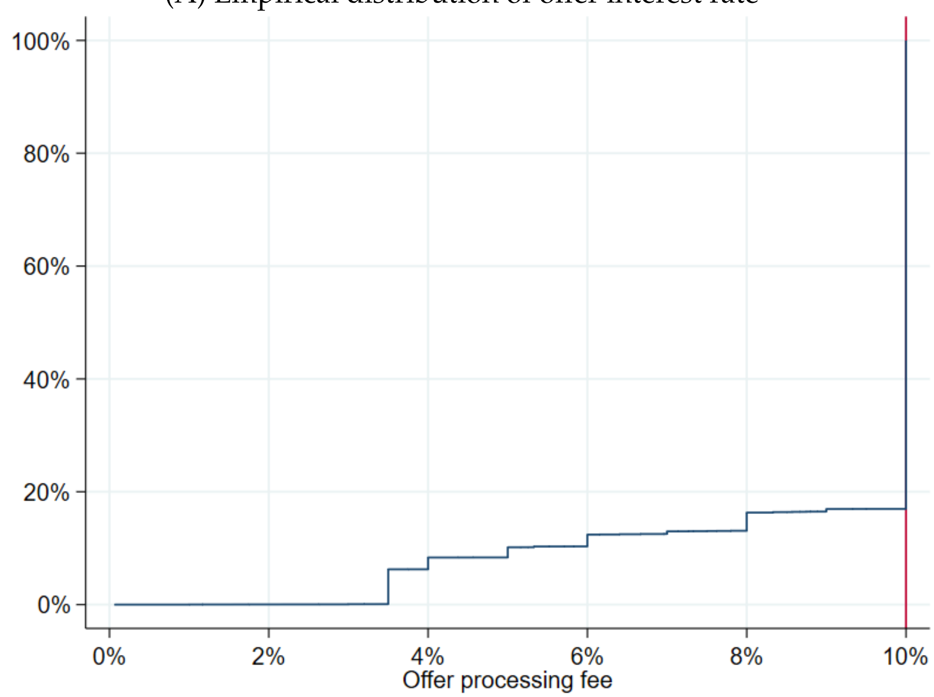

(B) Empirical distribution of offer processing fee

Figure IA.4: **Heterogeneous racial gaps across income groups**

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

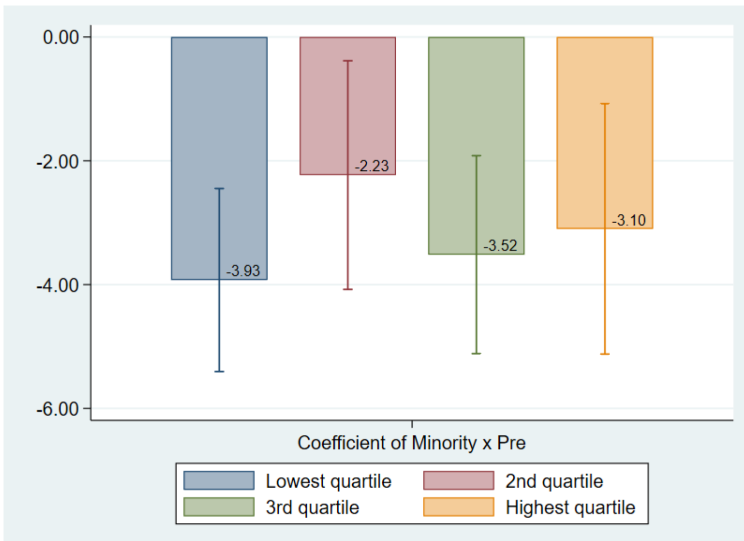

(A) Heterogeneous racial disparities by income

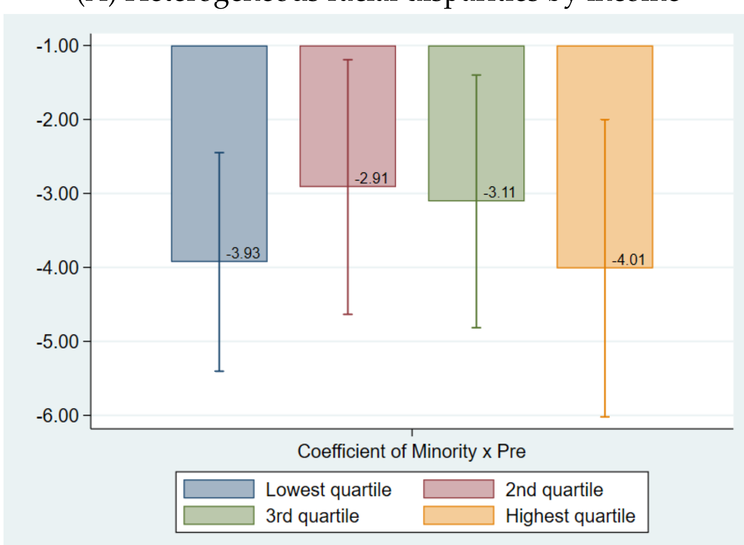

(B) Heterogeneous racial disparities by income-to-debt

Figure IA.5: **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 et al. (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.

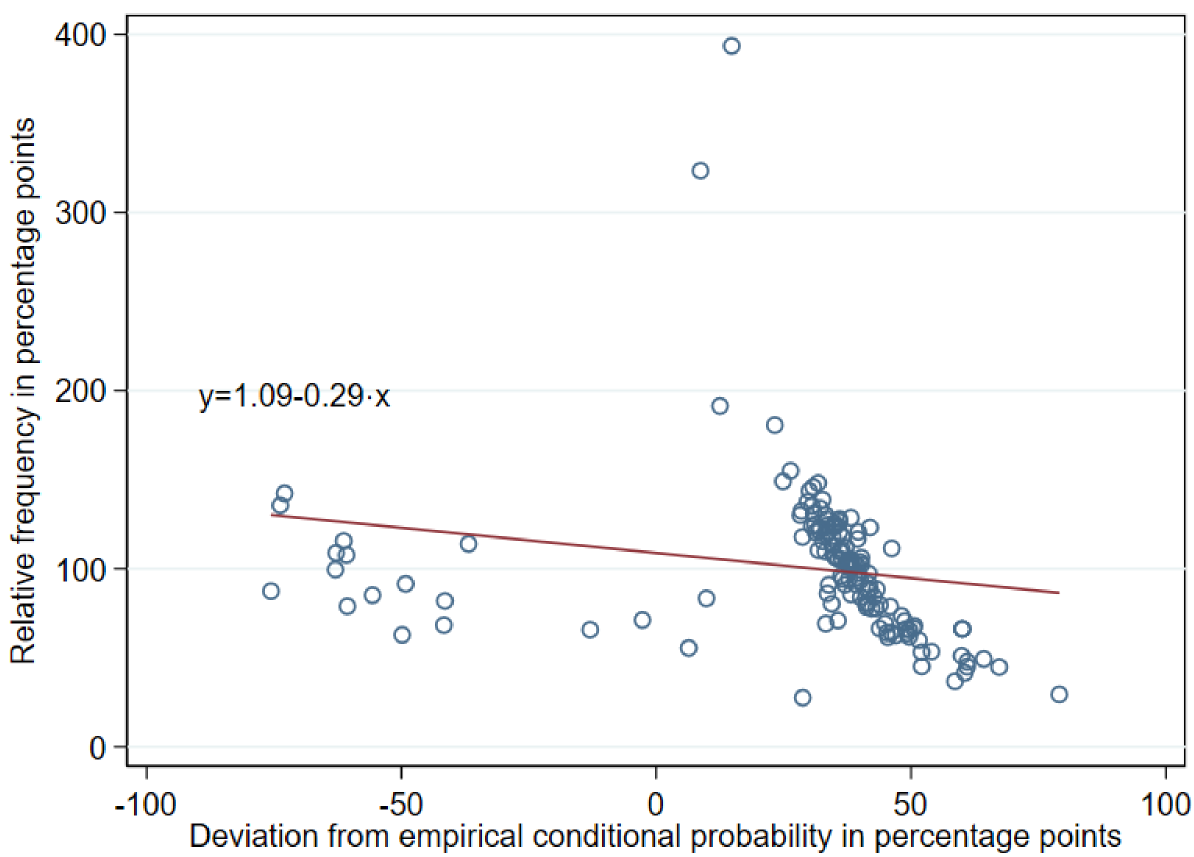

Table IA.1: **Within-application dispersion of initial offer terms**

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 IA.2: **Heterogeneous racial gaps across income-to-debt groups**

This table shows the effect of anonymous applications on disparities in credit outcomes for low- and high-income minorities. We estimate the income-to-debt-group specific minority-Chinese gap in offer probability before and after anonymization. High and low income/debt are dummy variables equal to one for applicants with above and below median levels of income-to-debt. In our specifications, we include all application characteristics observable 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 functional 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. We also include year-month fixed effects and lender fixed effects separately for the pre- and post-periods to control for invariant confounding factors; the included fixed effects are denoted at the bottom. 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.

	(1) Offer indicator (× 100)	(2) Log(offer amount)	(3) Offer maturity (months)	(4) Effective interest rate (%)
Minority × Pre × High income/debt	-3.475***	-0.019***	-0.105***	0.381**
	[-13.35]	[-3.98]	[-3.93]	[2.09]
Minority × Pre × Low income/debt	-4.138***	-0.082***	-0.196***	1.372***
	[-13.96]	[-13.58]	[-6.89]	[6.55]
Minority × Post × High income/debt	1.286***	0.003	-0.005	0.290
	[3.61]	[0.41]	[-0.12]	[0.96]
Minority × Post × Low income/debt	-0.829**	-0.064***	-0.195***	1.027***
	[-2.14]	[-8.03]	[-4.49]	[3.77]
Year-Month FEs	Yes	Yes	Yes	Yes
Lender × Post FEs	Yes	Yes	Yes	Yes
Observable controls	Yes	Yes	Yes	Yes
R^2	0.306	0.614	0.627	0.723
No. of observations	322,847	123,300	123,300	123,300

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

This table shows the robustness checks on the effect of anonymous applications on disparities in offer probability. We estimate the minority-Chinese gap in offer probability before and after anonymization. Column 1 estimates the baseline specification (equation (1)) in the full sample. Columns 2 and 3 report estimates obtained under different parametrization of control variables in the main sample. Column 4 and 5 report the estimates obtained with alternative measures of applicant race. In our specifications, we include all application characteristics observable to lenders at the time of initial screening as control variables. In choosing the functional 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), used in Columns 1, 4, and 5) allows the effects of the control variables to differ in the pre- and post-periods. Column 2 reports the estimates if we fix the coefficients of observable characteristics to be the same in the pre and post periods. The augmented specification (equation (IA.1), used in Column 3) allows the effects of the control variables to differ by lender and by whether the application is in the pre- vs post-period. We also include year-month fixed effects and lender fixed effects separately for the pre- and post-periods to control for invariant confounding factors; the included fixed effects are denoted at the bottom. 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.

	Outcome variable: Application-lender level offer indicator ($\times$ 100)				
	Alternative sample/specifications			Alternative measures of race	
	(1)	(2)	(3)	(4)	(5)
	Full sample	Looser controls	Tighter controls	Using government records	Using empirical frequency from names
Minority $\times$ Pre	-4.006***	-3.340***	-3.779***	-3.963***	-4.083***
	[-19.02]	[-14.13]	[-17.48]	[-17.41]	[-18.00]
Minority $\times$ Post	-0.411*	-0.808**	0.335	-0.289	-0.252
	[-1.73]	[-2.50]	[1.26]	[-1.00]	[-0.90]
Year-Month FEs	Yes	Yes	Yes	Yes	Yes
Lender $\times$ Post FEs	Yes	Yes	Yes	Yes	Yes
Observable controls	Yes	Yes	Yes	Yes	Yes
R^2	0.293	0.302	0.400	0.305	0.305
No. of observations	468,663	322,847	322,847	331,469	331,469

Table IA.4: **Testing for the impact of anonymous applications on overall credit supply**

This table presents tests for the impact of anonymous applications on overall credit supply. We compare how anonymous applications affect biased relative to unbiased lenders (equation (IA.2)). Biased lender is a dummy variable equal to one if the lender's racial gap in offer probability before anonymization is higher than the median. Lender-specific racial gaps are estimated with a specification analogous to equation 1 lender-by-lender, where we include the Post indicator instead of Lender×Post fixed effects and cluster standard errors at the month level for the lender-specific samples. We include all application characteristics observable 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 functional 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. Fixed effects are included and denoted at the bottom. 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.

	(1) Offer indicator (× 100)	(2) Log (offer amount)	(3) Offer maturity (months)	(4) Effective interest rate (%)
Biased lender × Post	-1.597	-0.0373	0.336	-0.154
	[-1.02]	[-0.69]	[1.21]	[-0.09]
Year-Month FEs	Yes	Yes	Yes	Yes
Lender FEs	Yes	Yes	Yes	Yes
Post FE	Yes	Yes	Yes	Yes
Observable controls	Yes	Yes	Yes	Yes
R^2	0.299	0.606	0.621	0.724
No. of observations	306,989	118,876	118,876	118,876