

D'Acunto, Francesco; Ghosh, Pulak; Jain, Rajiv; Rossi, Alberto G.

Working Paper

How costly are cultural biases?

LawFin Working Paper, No. 34

Provided in Cooperation with:

Center for Advanced Studies on the Foundations of Law and Finance (LawFin), Goethe University

Suggested Citation: D'Acunto, Francesco; Ghosh, Pulak; Jain, Rajiv; Rossi, Alberto G. (2022) : How costly are cultural biases?, LawFin Working Paper, No. 34, Goethe University, Center for Advanced Studies on the Foundations of Law and Finance (LawFin), Frankfurt a. M., <https://nbn-resolving.de/urn:nbn:de:hebis:30:3-651757>

This Version is available at:

<https://hdl.handle.net/10419/261464>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



LawFin Working Paper No. 34

How Costly Are Cultural Biases?

Francesco D'Acunto | Pulak Ghosh | Rajiv Jain |
Alberto G. Rossi

How Costly Are Cultural Biases?*

Francesco D’Acunto[†]
Georgetown University

Pulak Ghosh[§]
IIMB

Rajiv Jain[‡]
Faircent

Alberto G. Rossi[¶]
Georgetown University

This version: June 2022

Abstract

We estimate the cost of cultural biases in high-stake economic decisions by comparing agents’ peer-to-peer lending choices with those the same agents make under the assistance of an automated robo-advisor. We first confirm substantial in-group vs. out-group and stereotypical discrimination, which are stronger for lenders who reside where historical cultural biases are higher. We then exploit our unique setting to document that cultural biases are costly: agents face 8% higher default rates on favored-group borrowers when unassisted. The returns they earn on favored groups increase by 7.3 percentage points when assisted. The high riskiness of the marginal borrowers from favorite groups largely explains the bad performance of culturally-biased choices. Because varying economic incentives do not reduce agents’ biases, inaccurate statistical discrimination—unconscious biased beliefs about borrowers’ quality—can explain our results better than taste-based discrimination.

Keywords: Trust, Social Capital, Discrimination, Cultural Norms, Robo-Advising, Biased Beliefs, Inter-ethnic Conflict, Social Conditioning, Religion, Caste.

*The paper benefited significantly from a fellow visit of Francesco D’Acunto at the Center for Advanced Studies on the Foundations of Law and Finance funded by the German Research Foundation (Deutsche Forschungsgemeinschaft, DFG)—project FOR 2774. For very helpful comments, we thank Tetyana Balyuk, Vicki Bogan, Emily Breza, Paul Goldsmith-Pinkham, Luigi Guiso, Apoorv Gupta, Sasha Indarte, Filippo Mezzanotti, Tim McQuade, Anirban Mitra, Adair Morse, Daniel Paravisini, Chris Parsons, Vesa Pursiainen, Paola Sapienza, Antoinette Schoar, Kelly Shue, Lea Stern, Huan Tang, David Thesmar, Ansgar Walther, Luigi Zingales as well as seminar participants at the 2021 NBER Corporate Finance meeting, the NYU Race, Discrimination, and Inequality workshop, the 2021 Western Finance Association, the 2021 SFS Finance Cavalcade, the 2021 Midwest Finance Association, the CEBI Workshop on Subjective Beliefs and Household Finance, the 2021 Cornell Household Finance workshop, and the University of California at Berkeley, University of Rochester, Boston College (Bartunek Research Forum), EPFL, Université de Lausanne, HEC Paris, Purdue University, Aarhus University, Arizona State University, City University of Hong Kong, University of Alberta, University of Missouri, the WEFIDEV online seminar series, and the JILAEE webinar series. All errors are our own.

[†]McDonough School of Business, Georgetown University. e-Mail: francesco.dacunto@georgetown.edu

[‡]Fairassets Technologies India Private Ltd, Gurugram, Haryana, India.

[§]Indian Institute of Management, Bannerghatta Road- 560076, Bangalore, India. e-Mail: pulak.ghosh@iimb.ac.in

[¶]McDonough School of Business, Georgetown University, Washington, DC, USA. Fellow of the Luohan Academy. e-Mail: agr60@georgetown.edu.

1 Introduction

The cultural norms to which agents are exposed, especially when deep-rooted and based on centuries-long societal customs, can have a long-lived influence on beliefs and economic decision-making (Guiso, Sapienza, and Zingales (2006), Alesina, Giuliano, and Nunn (2013), D’Acunto et al. (2019)). Norms can also produce cultural biases, whereby agents’ choices deviate from those of a neoclassical agent, thus lowering consumption utility (Guiso, Sapienza, and Zingales (2009), Pursiainen (2020), D’Acunto (2019)).¹ Examples of cultural biases include *taste-based discrimination*, whereby discriminators willingly take costly actions to avoid interacting with certain social groups (e.g., see (Becker (1957), Akerlof and Kranton (2000), Parsons et al. (2011), Hjort (2014))) and *inaccurate statistical discrimination*—ex-post systematically incorrect beliefs about the quality of counterparts based on social stereotypes of which discriminators are unaware (Bohren, Haggag, Imas, and Pope (2019)).

Detecting and quantifying the effects of cultural biases on discriminators’ economic choices is challenging because, in the absence of full information about agents’ quality, the mere fact that somebody belongs to a discriminated group might provide a reliable signal of their quality (*statistical discrimination*) (Phelps (1972), Borjas and Goldberg (1978)), thus affecting decision-makers’ performance positively. Moreover, belonging to the same social group might improve principals’ ability to monitor agents, which makes discrimination economically valuable to discriminators (e.g., see (Fisman et al. (2017) and Fisman et al. (2020))). Isolating and quantifying the negative effects of cultural biases on discriminators’ choices, if any, requires a setting in which these channels are muted. The ideal setting includes high-stake economic decisions that have a direct and quantifiable effect on discriminators’ consumption utility to assess if economic incentives (un)successfully reduce cultural biases.

This paper proposes a field setting to estimate the effects of cultural biases on discriminators’ performance—a FinTech peer-to-peer (P2P) lending platform paired with an automated robo-advising tool (D’Acunto et al. (2019), Rossi and Utkus (2020), D’Acunto and Rossi (2020)). We compare lenders’ unassisted choices with those the tool proposes, which lenders can decide to implement or override. We show that the tool’s proposed choices of borrowers are uncorrelated with (unobserved to the tool) borrowers’ social groups. Consistent with earlier research, in our setting lenders discriminate substantially: they are systematically more likely to provide credit to borrowers that belong to their own social group rather than to other available borrowers. This imbalance, though, disappears once lenders observe borrower matches suggested by the robo-advising tool.

We then show that discrimination worsens lenders’ performance. On average, lenders face 8% higher default rates when making unassisted choices and we estimate that they earn up to 7.3 percentage-point higher returns after cultural debiasing. These results refer to actual choices in the field rather than hypothetical algorithmic decisions (Kleinberg et al. (2018), Tantri (2021)) and hence account for the possibility that lenders might not

¹These deviations can be optimal if the agent’s utility decreases when he/she makes choices that conflict with the cultural norms to which they adhere. Deviations are only suboptimal if individuals would have preferred behaving like a neoclassical agent had they been aware of their cultural bias.

adopt and implement advice even when they are provided with it (Bhattacharya et al. (2012)). The fact that the vast majority of lenders do not override the robo-advisor's suggestions reduces the possibility that lenders obtained sizable non-consumption utility from biased choices, for instance due to taste-based discrimination or kin altruism and warm glow towards their favorite group.

Our platform is a leading P2P, *Faircent*, whose lenders and borrowers reside all over India. Contrary to marketplace lending in the US (Paravisini, Rappoport, and Ravina (2017); Vallee and Zeng (2019); Chiu et al. (2018); Dobbie et al. (2020)), *Faircent* admits only individuals who invest their own capital, and hence the incentives of capital owners and lenders coincide in the same agent. The platform screens prospective borrowers based on their risk profile *before* borrowers are visible to lenders. If lenders could screen borrowers more successfully than the platform (Iyer et al. (2016); Tang (2019); Balyuk (2019)), they should perform weakly better when making decisions unassisted than after using the robo-advising tool.² Moreover, contrary to local bank branches, the platform has little scope for relationship lending: 90% of the lenders provide funds to borrowers who reside in at least 5 different Indian states and the platform engages in screening, monitoring, and servicing after loans are issued, without any interactions between lenders and borrowers.

The Indian setting allows us to study two forms of cultural biases. We start with in-group vs. out-group discrimination, whereby the members of two social groups favor members of their own group (in-group) and disfavor members of a conflicting group (out-group). In the Indian context, this form of discrimination has been detected between Hindus and Muslims (Brass, 2011; Tajfel et al., 1979). We then study stereotypical discrimination, whereby everybody discriminates against one social group—so much so that even the members of that group have excessively negative beliefs about the quality of their similar and discriminate among each other. We argue that this form of discrimination arises due to the centuries-old caste-based system in India.

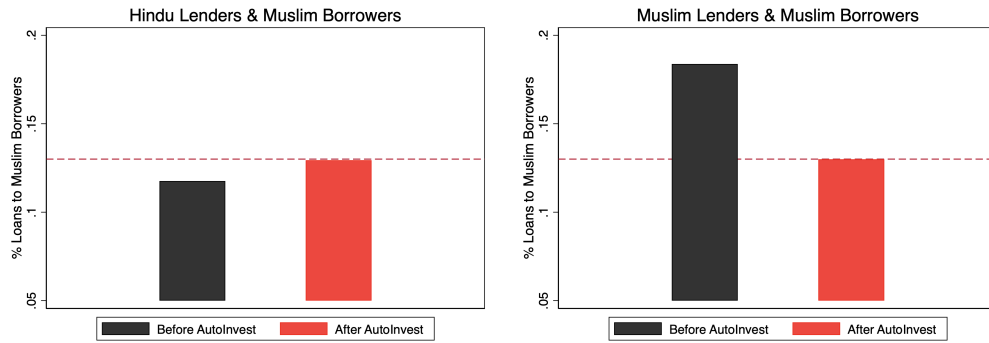
Figure 1 previews our baseline results in the raw data for the case of in-group vs. out-group discrimination. We compare the average share of Hindu and Muslim borrowers who are financed by Hindu and Muslim lenders before (black bars) and after lenders access the robo-advising tool (red bars). The data reveal three facts. First, both Hindu and Muslim lenders tend to lend more to same-religion borrowers, which is consistent with an in-group vs. out-group bias for at least one group. Second, once the robo-advising tool proposes matches, the shares of Muslim and Hindu borrowers change for both groups in opposite directions, which suggests that all lenders discriminated against out-group borrowers when unassisted. Third, the shares of borrowers by religion are virtually identical across lenders after robo-advice adoption and correspond to the borrower population shares on the platform. This result corroborates the notion that the robo-advising tool does not use information about borrowers' religion when allocating them across lenders.³

²If lenders engaged in accurate statistical discrimination, on average the performance of the borrowers they pick on their own should be better than that of the borrowers the tool assigns to them based on the platform's screening procedure.


³As we discuss below, the robo-advising tool as well as the selection of the borrower pool that accesses the platform might be subject to (accurate) statistical discrimination. We are not arguing that the platform and the tool do not engage in any form of discrimination, but that cultural biases do not appear to be present in the allocation choices of the tool.

Figure 1: Lending to In-Group vs. Out-Group Borrowers: Before and After Debiasing

Panel A. Probability of Choosing Muslim Borrowers



Panel B. Probability of Choosing Hindu Borrowers



The raw-data patterns in Figure 1 are quite robust: they do not change in multivariate analyses that control for the loan-level characteristics we observe, when we restrict the variation within lenders, and hence absorb unobserved time-invariant differences across lenders such as education levels and financial literacy, or when account for different time-varying shocks before and after the robo-advising tool was available to lenders.⁴

To corroborate our cultural-bias interpretation, we perform heterogeneity tests that vary the salience of cultural biases to lenders. The extent of bias and debiasing are higher for lenders in cities with a high occurrence of Hindu-Muslim riots, in states where parties that foment inter-religious conflict obtain higher vote shares, and for lenders exposed to heightened Hindu-Muslim animus during their formative years.

After detecting the presence of in-group vs. out-group discrimination, we move on to study whether discrimination relates to performance and, if so, through which channels. Under cultural biases—taste or inaccurate beliefs—when making unassisted decisions, lenders should dig deeper into the pool of in-group borrowers, who should perform worse than out-group borrowers (Agarwal et al. (2017)). The opposite should be true if lenders engaged in accurate statistical discrimination.

We find that the loans lenders grant to in-group borrowers before adopting the robo-advising tool perform systematically worse than out-group borrowers' loans, in terms of both default rates and returns earned. In-

⁴Note that, because debiasing implies that the shares of Hindus and Muslims move in opposite directions for Hindu and Muslim lenders, systematic cross-sectional or time-series differences between the two groups of borrowers cannot explain our results.

group-borrower loans are about 2.4 percentage points (pp) more likely to default (8% of the average default rate). High-interest-rate loans might provide high returns even if they are not repaid in full. Yet, we find that the returns of out-group loans were higher before lenders adopted the tool, and after adoption the returns of in-group loans increase substantially more than those of out-group loans. In back-of-the-envelope calculations, we estimate that the cost of in-group vs. out-group discrimination amounts to about 6% of the average capital lenders invested in the platform before the robo-advising tool was available.

To dig deeper into the drivers of improved returns, we show that they are mostly driven by a higher performance of the left tail of the distribution of in-group borrowers. Moreover, the lower delinquencies and higher returns lenders enjoy after adopting robo-advising are driven almost exclusively by the different risk profiles of the pools of in-group borrowers before and after using the tool. Ultimately, lenders improve their performance because, when discriminating, they tend to choose high-risk in-group borrowers, which ex-post deliver lower returns on average, and which the robo-advisor does not pick.

Note that in our setting several channels that predict a positive association between discrimination and performance cannot drive the results by construction. For instance, homophily in monitoring borrowers and relationship lending have no scope in our setting given that lenders disburse funds to borrowers all over India and do not interact with them either before or after the loans are issued (Iyer et al. (2016); Schoar (2012); Drexler and Schoar (2014); Fisman et al. (2017); Fisman et al. (2020)). For these reasons, the incentive effects of social collateral (Karlan, Mobius, Rosenblat, and Szeidl (2009); Diep-Nguyen and Dang (2019)), moral incentives and social image (Bursztyn et al. (2018); Bursztyn et al. (2019)), peer effects (Breza (2019)), familiarity through in-person interactions (Rao (2019)), preferences of physical appearance (Duarte, Siegel, and Young (2012); Ravina (2019)), or systematic ethnic differences in housing collateral value (Avenancio-León and Howard (2019); Naaraayanan (2019))⁵ have no scope in our setting either.

The second form of cultural bias we study is stereotypical discrimination, which, in India arises against members of the lower Hindu caste, namely, *Shudra*.⁶ This form of discrimination has received less attention from the recent literature in cultural economics, which typically focuses on in-group vs. out-group discrimination based on inter-ethnic conflict (e.g., see Jha (2014), Hjort (2014), D’Acunto et al. (2019), Hjort et al. (2021)).

We detect substantial stereotypical discrimination: before lenders (including *Shudra* lenders) adopt the tool, *Shudra* borrowers are less likely to appear in their loan portfolios relative to their share in the population of borrowers. Once lenders adopt the robo-advising tool, the share of *Shudra* borrowers they finance increases and the default rates and returns of *Shudra* and other borrowers converge. Moreover, discrimination against *Shudra* borrowers is higher for lenders who reside in states in which the share of crimes against lower castes is higher, which we interpret as settings in which the negative stereotypes against lower castes might be more salient.

Caste-based discrimination provides an additional layer of variation for our tests due to the varying degree

⁵Loans are not backed by collateral on the platform.

⁶Unfortunately, we do not have enough *Dalits* in our sample to test for discrimination against this group.

of recognizability of one's caste when it is not disclosed explicitly as in our setting. On *Faircent*, lenders can only infer borrowers' caste based on a set of individual characteristics—e.g., borrowers' surnames, locations, and occupations—and the extent to which these characteristics are a precise signal of one's caste varies. We thus design an intensive-margin test of discrimination inspired by the experimental literature (Mobius, Rosenblat, and Wang (2016)). We build on Bhagavatula et al. (2017) and Bhagavatula et al. (2018), who design and train an algorithm that replicates an average Indian's inference problem of which caste individuals of known demographic characteristics belong to. In this way, for each borrower we obtain a measure of the likelihood that lenders might recognize whether they are *Shudra*. We find that discrimination against *Shudra* borrowers increases with the likelihood that the borrower is recognizable as a *Shudra*, whereas it virtually disappears for borrowers whose caste is difficult to assess.

The fact that *Shudra* lenders, too, discriminate against *Shudra* borrowers allows us to disentangle our cultural-bias interpretation from kin altruism—the tendency of individuals to take costly actions to favor those who share similar ethnic or other demographic characteristics (e.g., see Simon (1993)).⁷

In terms of performance, we find that lenders face fewer defaults and higher average returns after accessing the tool. Second, the improvement in returns comes largely from the elimination of a left tail of low returns. Third, the performance improvement is driven by a drop in the riskiness of the loans that appear in lenders' portfolios after high-risk non-*Shudra* borrowers are replaced by lower-risk *Shudra* borrowers.

Statistical discrimination, rather than biased discrimination, can arise in settings in which discriminated groups are a minority and decision-makers, from the majority, can assess the characteristics of their similar more easily than of minority groups (Cornell and Welch (1996)). In our setting, this explanation could be consistent with the behavior of Hindu lenders but not of Muslim lenders—who discriminate against the majority, Hindus—or *Shudra* lenders, for whom the favored group (non-*Shudra*) is the majority but does not coincide with the easier-to-assess group (*Shudra*), against whom they discriminate.

Our paper differs from earlier work on the effects of cultural biases on field economic choices. For instance, in Hjort (2014) upstream workers are aware of downstream workers' productivity and hence the setting has no scope for inaccurate statistical discrimination. Consistently, in Hjort (2014) economic incentives reduce biased choices when in-group agents would be affected as negatively as out-group ones. In our setting, instead, inaccurate statistical discrimination is plausible due to the severe asymmetric information about borrowers' quality. Also consistent with inaccurate statistical discrimination and lenders' lack of awareness that they make discriminatory choices, economic incentives do not lead to debiasing in our setting.

In principle, both taste-based and inaccurate statistical discrimination might arise on our platform: lenders might dislike discriminated groups and at the same time form inaccurate beliefs about their performance. Despite the lack of a direct test, inaccurate statistical discrimination is a more coherent explanation. First, high-stake

⁷We thank David Thesmar for raising this point. Kin altruism is also unlikely to explain Hindu-Muslim discrimination: if lenders obtained non-consumption utility from kin altruism they would override the tool's suggestions to pick borrowers from the out-group.

economic incentives alone do not lead to debiasing (Hjort (2014)). Moreover, virtually no lenders override the tool’s proposals,⁸ which they should if they had a conscious distaste for interacting with discriminated groups. Also, we find that the extent of lenders’ in-group bias does not predict adoption of the tool, which is additional anecdotal evidence that our results might be more consistent with an unconscious bias, such as inaccurate statistical discrimination, rather than a conscious one, such as taste-based discrimination.

Our results also emphasize an unintended role of robo-advising tools (D’Acunto and Rossi (2020)), which are diffusing around the world to facilitate consumers’ spending (D’Acunto et al. (2019)), saving (Gargano and Rossi (2020)), borrowing (Agarwal et al. (2019)), and lending decisions. We show that such tools can help discriminating agents avoid the financial losses they face when making culturally biased choices. Robo-advising might be a viable substitute for financial disclosure and financial literacy, because it does not require that agents understand fully the problems they face (Adams et al. (2019); D’Acunto et al. (2019b); D’Acunto et al. (2019a)).

2 Institutional Setting

The setting for our analysis is Faircent, a large FinTech platform that specializes in P2P lending in India.⁹

2.1. Borrowers’ Screening by the Platform

Several features of the platform are crucial to the design and interpretation of our tests. First and foremost, the platform screens applicant borrowers before admitting them to the pool and lenders observe the results of this screening. Even though the platform does not use racial variables, its automated screening procedure might incorporate statistical discrimination (e.g., see Bartlett et al. (2019); Bhutta et al. (2021); Cowgill and Tucker (2020); Fuster et al. (2017); Rambachan et al. (2020)).

Once a prospective borrower signs up, he/she submits a loan application that includes the proposed amount and motivation of the loan, the borrower’s credit score, occupation, geographic location, and whether the borrower has dependents. Borrowers need to provide evidence of a financial account, which excludes unbanked applicants. The first screening step is an automated algorithmic-based assessment of the borrower’s credit viability, which is largely based on the borrower’s credit score, proposed loan amount and maturity, and occupation.¹⁰ Applicants whose credit viability falls below a fixed threshold are dismissed.

The Online Appendix reports raw-data evidence on the outcomes of the automated screening procedure (see Figure A.1). Credit scores are monotonically related to loans’ annual interest rate, maturity, and loan amounts, all of which are assigned by the platform (see Figure A.2).

⁸The fact that agents barely modify the allocations proposed by robo-advisors is common to other forms of robo-advising, such as those for equity portfolio allocations (D’Acunto et al. (2019)) and debt management (Burke et al. (2021)).

⁹This setting is reminiscent of the recent literature on FinTech adoption in developing countries (e.g., see Agarwal et al. (2019); Crouzet, Gupta, and Mezzanotti (2019); Higgins (2019); D’Andrea and Limodio (2019)).

¹⁰Due to confidentiality reasons, we cannot describe how this information is combined in the automated screening.

The prospective borrowers who are approved and accept the parameters proceed to the verification step, whereby Faircent's employees verify borrowers' self-provided information and documents: borrowers' identity (identity cards and a personal picture), two income paystubs or incoming transactions in a bank account under the borrower's name, utility payments, the picture of the borrower's housing and working locations. Borrowers who pass this verification step are admitted to the pool that lenders can browse.

Lenders observe borrowers' demographic characteristics and the qualitative and quantitative risk assessment from the screening process (we discuss the information lenders see in section 2.2. below).

In this setting, lenders make choices *after* a substantive risk assessment of borrowers, whose outcomes lenders observe. Decoupling borrowers' risk screening from lending decisions, which departs from earlier research that studied the choices of loan officers, reduces the scope for accurate statistical discrimination on the part of lenders. Of course, lenders might believe that they can screen borrowers better than the platform (Balyuk and Davydenko (2019); Vallee and Zeng (2019)). If true, lenders' unassisted choices should perform better than robo-advised choices, which is the opposite of what we find and is consistent with biased discrimination.

Execution, servicing, and monitoring after loan approval is also managed by the platform. Once at least 5 lenders choose a borrower,¹¹ the platform approves and executes the loan. The loan agreement is a private contract between the borrower and each lender, but the platform produces the electronic forms that lenders and borrowers have to sign. No lenders enjoy any form of seniority. Upon execution, Faircent tells borrowers their equated monthly installment (EMI)—the monthly payment—and services the loan. Faircent monitors the status of loans each month. Loans' status is declared "closed" after full repayment or after repeated delinquency. Borrowers whose loans are closed while delinquent are dismissed from the platform. Faircent's loans are subject to the Reserve Bank of India (RBI) regulatory policies and oversight like traditional financial institutions.

Figure A.4 of the Online Appendix shows the distribution of the number of states in which the borrowers of each lender reside. The figure reveals that more than 90% of our lenders choose borrowers who live across at least 5 different Indian states. Also, the median lender disburses funds to borrowers who reside in 13 different Indian states. In this respect, our setting is quite different from local-branch relationship lending by professional loan officers, in which soft information about borrowers and social norms/social pressure might influence borrowers' repayment behavior.

2.2. Robo-Advising Tool (Auto Invest)

The second important feature of the platform is that lenders can make their choices unassisted or under the assistance of a robo-advising tool, called Auto Invest.

For unassisted choices, lenders can browse the borrower pool at any point in time. Lenders observe the coarse risk category assigned by the platform (low, medium, or high risk), the detailed risk assessment of the

¹¹The platform imposes that each loan is financed by at least 5 lenders to enhance the diversification of lenders' loan portfolios and to ensure that the platform cannot be used for money laundering purposes.

loan (interest rate, maturity, overall loan amount), and a set of borrowers' demographic characteristics that include names and surnames, location of residence, occupation, education levels, and number of dependents. Upon clicking on the borrowers' profile, lenders can see the verification report about their identity. Lenders can decide who, if anybody, they want to fund and by how much. Because the platform imposes that each loan is financed by at least 5 different lenders to reduce the scope for money laundering and other criminal uses of the platform, lenders can at most finance 20% of the chosen borrower's loan amounts.¹² Lenders need to have the funds they want to commit deposited on the platform before their loan proposals are posted to the borrowers. Loans proceed to the execution phase only if the financing of the full loan amount is committed.

The second mode of investment is with the assistance of an automated robo-advising tool, *Auto Invest*, which was introduced on the platform in the second half of 2018. After its release, lenders can adopt *Auto Invest* at any time. Upon adoption, lenders decide the share of their overall funds deposited on the platform for which they want to obtain suggestions of borrowers by *Auto Invest* and the amount of funds, if any, they want to keep investing unassisted. Lenders are then asked to allocate the funds for which they seek assistance by *Auto Invest* across the same three borrower-risk categories they see when making unassisted decisions—low, medium, and high risk. We report the screenshot of *Auto Invest* in Figure A.3 of the Online Appendix.

For each risk category, the tool matches lenders to the borrowers who are closest to get their loan fully funded. *Auto Invest* does not choose borrowers explicitly based on their interest rate or level of risk within each risk category or based on borrowers' demographic characteristics. Below, we show that *Auto Invest* does not incorporate borrowers' demographics implicitly either. In contrast with the way in which the platform advertises the tool, we find no evidence that *Auto Invest* allocates the best-performing borrowers on the platform faster than would be possible with manual choices. In fact, as we discuss in the last part of the paper, before accessing the tool lenders tended to choose in-group borrowers who performed worse than out-group borrowers available on the platform, whose loans ended up not being funded and hence who would have been available at any time.

Auto Invest might affect lenders' loan portfolios in many ways unrelated to cultural biases, for instance, by increasing their diversification. For this reason, when assessing the cost of cultural biases, we will compare lender-level returns on favored-group borrowers to the same lenders' returns on discriminated borrowers, which capture the effects of *Auto Invest* unrelated to cultural debiasing.

Once *Auto Invest* proposes a set of borrowers, the lender decides if she wants to proceed with the suggested allocation or change it in part or in full. To make changes, she browses the pool of borrowers and finds replacements. This step is crucial for the interpretation of our results, because under conscious forms of discrimination, such as taste-based discrimination, we would expect that lenders override the suggestion of groups they dislike.

In terms of adoption, we do not find that any of the (limited) set of lender-level characteristics predicts the decision to activate *Auto Invest*. Table A.1 in the Online Appendix shows that, in particular, neither the

¹²Within the platform, there is no way to coordinate or communicate among lenders.

extent to which lenders are biased against out-group borrowers or lenders' religion predicts a higher or lower likelihood of adoption.¹³ Moreover, none of the two religious groups is more sensitive to their in-group bias in terms of adopting Auto Invest, where in-group bias is measured as the difference between the share of out-group borrowers in the platform's pool and the share of out-group borrowers in the lender's loan portfolio.

Unfortunately, we do not observe whether each individual loan contribution was made by the lender unassisted or through Auto Invest. For this reason, we cannot compare the choices each individual lender made manually with those based on Auto Invest's suggestions at the same point in time by absorbing lender-by-time variation in our multivariate analyses.

3 Data and Summary Statistics

To perform our analyses, we use seven data sets, each of which covers a different feature of the lending process at Faircent. Our data span the period between January 1, 2018, and March 30, 2020, although, given the large monthly growth of the platform, 60% of the loans in our sample were issued in 2019 and another 19% in the first three months of 2020. Variation in the timing of loan issuance is therefore limited. We limit the sample to the end of March 2020 to avoid covering the COVID-19 pandemic period.

Two data sets—the *Lenders' characteristics data set* and the *Borrowers' characteristics data set*—include cross-sectional data with one observation per individual. Each lender and each borrower is assigned a unique identifier, which allows us to link lenders' and borrowers' characteristics across data sets. For each lender, we observe name and surname, the state of residence,¹⁴ and the date of birth. On top of these characteristics, the borrowers' sample also includes the borrower's residence type (whether owned or rented), number of dependents, employment type (whether self-employed or not), and credit score.

Faircent does not ask any information about religion or caste, and hence lenders have to infer these dimensions based on borrowers' demographic information they observe. We will exploit the extent to which lenders can easily infer castes based on borrowers' characteristics in our analysis. We as econometricians infer these variables using the *Marriage registry data set* (see Bhagavatula et al. (2017)), which includes demographic information about religion and caste elicited at marriage for a random sample of 2,481,158 Indians. It includes names and surnames, date of marriage, state of birth, city of residence, height (in centimeters), religion, and caste.

Assigning religions is quite straightforward, which implies that lenders should also find it easy to infer borrowers' religion, because religion barely varies across individuals who share the same surname. Eighty-nine percent of the unique pairs of surnames and dates of birth in the registry (everybody who was born the same day and shares the surname) belong to same-religion individuals. When we only consider Hindus and Muslims,

¹³Although our full sample includes 2,818 unique lenders, for this analysis we can only include lenders for whom the religion and caste of all borrowers in their portfolios can be retraced, which is 1,567 unique lenders.

¹⁴For a subset of lenders, we also observe the city of residence, but whereas choosing the state from a pre-determined list is compulsory, writing down the city in a blank box is not required.

96% of the surname-date-of-birth pairs are matched uniquely to one of the two religions. For these reasons, we assign religions to lenders and borrowers based on surname and date of birth.

Assignment of castes is less straightforward. The caste information in the registry is dispersed and includes about 540 narrowly-defined partitions, which often merely correspond to the individual's surname. To make our analysis of castes meaningful, we need a reliable way to assess to which of the four main castes borrowers and lenders belong.¹⁵ None of the combinations of characteristics we observe in both the Faircent data and the registry—names, surnames, and dates of birth—restrict the set of castes enough to proceed in the same way as with religions. To infer borrower and lenders' varnas, we thus rely on research in computer science (Bhagavatula et al. (2017) and Bhagavatula et al. (2018)) using the procedure we discuss in Section 5.1..

Fourth, we use is the *Lender-Borrower Mapping*—a cross-sectional data set at the level of lender-borrower-loan triads. This information is critical to merge individual characteristics to borrowers and lenders who match through a loan that is in part funded by the lender. The data are also critical to merge information about loan characteristics and loan performance to each unique lender-borrower-loan triad.

Loan-level information is in cross-sectional and panel formats. The *Loan characteristics data set* is a cross-sectional data set at the loan level that reports the borrower identity, the total amount lent, interest rate, maturity, and proposed monthly payment. We also observe the loan's status as of March 31, 2020 (active or closed) and whether the last payment happened within the 31 days before closure (i.e., whether the loan was in good standing). The *Loan performance data set* is an unbalanced panel at the loan-month level. Unfortunately, we were able to obtain the monthly payments (including zero for missed payments) only for a random subset of the loans in our main sample.

Finally, the *Auto Invest data set* is a cross-sectional data set at the lender level that reports whether lenders have ever activated the robo-advising tool and, if yes, the activation date and the share of the funds for which the lender wanted matching proposals from the tool.

3.1. Sample Selection

To study in-group vs. out-group discrimination, we create a sample at the borrower-lender-loan level that includes all the Hindu and Muslim borrowers and lenders in the data for whom we observe no missing information on loan characteristics, the usage of Auto Invest, or loan performance.

A common concern when studying the adoption of a new technology is that unobserved characteristics and shocks might at the same time cause lenders' adoption as well as a behavioral change after adoption, so that this change cannot be attributed causally to using the new technology. Note that, as we discussed in Section 2.2., we find that none of the observable lenders' characteristics predict adoption in our setting.

Moreover, in our setting, this endogeneity concern is less compelling than in others, because we test whether

¹⁵We provide a primer on Indian castes and on how they relate to our analysis below.

different lenders change their behavior in opposite ways after adopting the robo-advising tool. For instance, if Hindu and Muslim borrowers faced different economic shocks over time and for this reason they became more or less viable borrowers, all lenders should react in the same direction in terms of choosing more or less of one of the two categories of borrowers, because borrower-level shocks affect the borrowers available in the pool in the same way, irrespective of whether they end up being chosen by a Hindu or a Muslim lender. Similarly, lender-level shocks might affect the extent to which lenders engage with the platform or the amount of money they invest on the platform over time, but once these dimensions are fixed at the lender level they should not impact the allocation of funds across borrowers of different religions. One possibility is that the lenders who think they have been performing worse on the platform are more likely to adopt the Auto Invest earlier, but then at a minimum these lenders were not realizing that cultural biases were the source of their low performance, or otherwise they would have changed their allocation across religions and castes before the introduction of Auto Invest. Even in this case, comparing the allocation of unassisted choices with the allocation of choices implemented through Auto Invest would allow us to measure lender-level cultural biases.

Nonetheless, to reduce concerns about the endogenous adoption of Auto Invest, we follow the literature¹⁶ and only consider lenders who activated Auto Invest at some point between its introduction in 2018 and the end of the sample period (March 31, 2020). In this way, we do not compare the choices of lenders who did not adopt Auto Invest with those of lenders who adopted it at some point in time. At each point in time, the sample includes lenders who have already activated the tool and others who have not yet activated it.

A second way in which we reduce concerns about the endogeneity of adoption is by exploiting the intensive-margin of the use of Auto Invest: we compare lenders who allocate a higher or lower share of their funds to the robo-advising tool, and hence lenders that might have activated the tool at the same point in time but who make a higher or lower fraction of choices unassisted.

The second working sample allows us to study the effects of stereotypical biases. We first select the sample using the same steps discussed above. Then, we further restrict the sample to only include Hindus, because other religious groups do not partake the caste hierarchy. Moreover, we only include borrowers for whom we can retrace caste information in the form of one of the four varnas based on the marriage registry data.

3.2. Summary Statistics

The first working sample includes 113,283 unique lender-borrower-loan triads, which are based on 2,818 unique lenders. Panel A of Table 1 reports summary statistics. Borrowers' religion, consistent with the split between Hindu and Muslim individuals in the general Indian population, is tilted toward Hindus—13% of the borrowers are Muslim. The religious imbalance is starker on the lenders' side—99% of lenders are Hindu. Despite the small share of Muslim lenders (1% of the sample), the number of observations is large enough to allow a statistically

¹⁶For instance, see D'Acunto et al. (2019), D'Acunto et al. (2020), and Gargano and Rossi (2021), among others.

meaningful analysis of Muslim lenders' choices. A potential reason for this imbalance is that the precepts of *Sharia* condemn the earning of financial interest on loans and might discourage Muslims to sign up as lenders.¹⁷

About 45% of the loans were issued when lenders had activated the robo-advising tool. The share of funds allocated to the tool is about 60% on average, but substantial cross-sectional variation exists across lenders. As far as loan characteristics are concerned, the average maturity (tenure) is 22 months and the median maturity is 24 months. The average loan amount is slightly above 130,000 rupees, which corresponds to about \$1,770,¹⁸ with a large standard deviation. The average annual interest rate is 24%—similar to the yearly APRs for credit cards in the US over the same period.¹⁹

The summary statistics so far refer to all unique lender-borrower-loan triads—the level of observation in most of our analysis. Considering the lender level, we find that the share of Muslims on the platform is stable throughout the sample period. We cannot reject the null that the shares are equal before and after Auto Invest is available. The shares are also stable in the borrower population. Lenders' engagement with the platform does not vary, either: the average lender issues 40 loan offers to any borrowers both before and after accessing Auto Invest (p-value of t-test for equality=0.683). In terms of loan characteristics, when computed at the lender level, the average interest rate of loans issued before using Auto invest is higher than that of loans issued after using Auto Invest (25% vs. 23%, p-value of t-test for equality=0.003) and the default rates lenders face in the raw data are substantially higher before they use Auto Invest (33% vs. 16%, p-value of t-test for equality < 0.1%). Also, lenders offer lower average amounts per loan after accessing Auto Invest (₹2,980 vs. ₹2,123, p-value of t-test for equality < 0.1%). As we will see in the last part of the paper, the fact that the loans issued through Auto Invest are less risky, on average, than those issued through unassisted choices is important to explain why culturally-biased choices deliver lower returns to lenders.

Our second working sample, which only includes Hindus whose caste we can infer, has 62,831 unique lender-borrower-loan triads. Panel B of Table 1 reports the summary statistics for this sample, in which 39% of borrowers belong to the discriminated *Shudra* varna. Despite the smaller size, the summary statistics for the main variables of interest in this sample are similar to those described in Panel A. The patterns for the statistics at the lender level, including the lack of changes in the composition of borrower pools before and after Auto Invest is available as well as the lower risk of loan offers lenders make when assisted by Auto Invest, are also the same as for the Hindu-Muslim sample.

¹⁷For this reason, the Muslim lenders who sign up might not be less attached to the religious precepts of Islam. Our tests are based on Muslim ethnic and social identity, to which Muslims should relate whatever their religiosity.

¹⁸This conversion does not adjust for varying the monthly exchange rate between the US dollar and Indian rupee.

¹⁹Over our sample period, the nominal interest rate of reference set by the Reserve Bank of India ranged between 4.5% and 6%.

4 In-group vs. Out-group Discrimination: Hindu vs. Muslims

We first analyse in-group vs. out-group discrimination: agents tend to favor members of their own social group (in-group) over members of conflicting social groups (out-group), where social groups' boundaries are defined based on cultural cleavages (Tajfel et al. (1979); Hewstone et al. (2002); Jenkins (2014)).²⁰

The cultural-bias hypothesis predicts not only that lenders should be more likely to pick in-group borrowers, but also that the in-group borrowers they choose should perform *worse* than the out-group borrowers they choose. Once lending decisions are automated, the likelihood that in-group or out-group members appear in lenders' portfolios as well as the average performance of these two groups of lenders should converge. These predictions are in contrast to what we should find if lenders' favoritism toward in-group borrowers were due to lenders' ability to screen and monitor in-group members better than out-group members.

We consider the religious conflict between Hindus—Indian religious majority—and Muslims, one of the religious minorities in post-independence India. Acts of in-group vs. out-group discrimination between these two religious groups are deeply rooted in history and pre-date the independence of modern India in 1947 as well as the British rule on the Indian subcontinent (Lorenzen (1999)). Not only have Hindu and Muslim identities developed in contrast over time, but the identity clash has also manifested in acts of conflict, including violent conflict and riots, for decades (e.g., see Engineer (1997)).

The Hindu-Muslim conflict has been exacerbated over the last two decades (Graff et al. (2012)) and erupted in violent riots, such as the anti-Muslim pogrom in the state of Gujarat in 2002 (Ghassem-Fachandi (2012)). Several political scientists and sociologists argue this conflict was exacerbated because right-wing political parties, such as the BJP, support Hindu nationalism (see Kaul (2017) among others). For instance, the approval of the *Citizenship Amendment Act* by the BJP in 2019 has produced a widely covered wave of riots and violence between Hindus and Muslims (Bhat (2020); (Mitra and Ray, 2014)).

We first assess the extent to which Hindu lenders might have been more inclined to finance Hindu borrowers, and Muslim lenders to finance Muslim borrowers, when making their decisions autonomously. Then, we compute the change in the propensities to lend to Hindus and Muslims after robo-advising adoption. Third, we assess whether the extent of lenders' in-group vs. out-group bias and debiasing were stronger for lenders who resided in areas with higher Hindu-Muslim animus, for whom the Hindu-Muslim conflict was arguably more salient when making lending decisions.

²⁰Unfortunately, we cannot provide a comprehensive description of all the facets and decades-long academic debate about this family of theories in this paper, but due to space constraints, we need to focus on the most relevant implications in terms of what we can test directly in our setting. For more comprehensive reviews, see, for example, Hewstone et al. (2002) and Jenkins (2014).

4.1. In-group vs. Out-group Lending Before and After Robo-advising

Figure 1 plots the average share of Hindu and Muslim borrowers in Hindu and Muslim lenders' portfolios in the period in which lenders made all loan decisions unassisted and after robo-advising adoption. The top graphs consider the choice to lend to Muslim borrowers before (black bar) and after (red bar) using Auto Invest. The bottom graphs repeat the analysis for Hindu borrowers.

Three broad patterns are worth noticing. First, consistent with the presence of an in-group vs. out-group bias for at least one of the two groups, both groups choose a higher share of borrowers of their own religion when unassisted: 86% of Hindu lenders' borrowers are Hindu, whereas only 80% of Muslim lenders' borrowers are Hindu. Conversely, the share of Muslim borrowers is 12% for Hindu lenders and 18% for Muslim lenders.

Second, after using the robo-advising tool, the shares of Muslim and Hindu borrowers change for all lenders in opposite directions: Hindu borrowers decrease from 86% to 84% of Hindu lenders' choices, whereas the share of Muslim borrowers increases from 12% to 13%. At the same time, Hindu borrowers move from 80% to 86% of Muslim lenders' choices, whereas the share of Muslim borrowers drops from 18% to 13%.

The fact that after using the robo-advising tool, lenders of different religions change their borrowing choices in *opposite* directions reduce the endogeneity concerns inherent to the decision of adopting Auto Invest. This decision might be endogenous to individual- and/or economy-wide time-varying shocks, but because we observe that adopting robo-advising affects behavior in opposite directions for lenders of different religions, the most plausible unobserved time-varying economic shocks about which we would worry, such as shocks that might simultaneously explain the adoption of the tool and a behavioral change in the same direction for all lenders, are by construction unable to explain our results.

Third, Figure 1 shows that the share of borrowers of different religions is equalized for Hindu and Muslim lenders after they use Auto Invest (red dashed horizontal lines). Unobserved channels that might make in-group borrowers more profitable to lenders than out-group borrowers are thus unlikely to exist in our setting. Moreover, this result corroborates the fact that the tool does not use information about religion (which it does not have) to allocate borrowers across Hindu and Muslim lenders. Ultimately, the tool debiases lenders in that it equalizes the share of borrowers of each religion in each lenders' portfolios to the share of borrowers of each religion in the broader population of borrowers.

The averages in Figure 1 compare lending behavior before and after lenders adopt the robo-advising tool, but Auto Invest allows lenders to choose the share of their funds on the platform they want to allocate through Auto Invest's suggestions and the share they want to allocate unassisted. We can thus assess the extent of debiasing when comparing lenders who adopt the tool at the same time, but make a different share of their choices through the tool's suggestions.

Figure 2 reports the results of this analysis in the raw data. Here, we only consider Hindu lenders who use

the tool.²¹ We sort lenders based on the share of funds with the platform they allocate to Auto Invest, which is strictly larger than zero and lower than or equal to 1. The solid blue line reports smoothed non-parametric estimates of the relationship between the share of Hindu borrowers in lenders' borrower portfolios (measured on the right y-axis) and the percentage of funds that lenders allocate to Auto Invest. Grey bandwidths refer to 95% confidence intervals around the point estimates of the slope of the curve for each percentage of fund allocation.

We find that a greater use of Auto Invest relates to in-group favoritism monotonically. Debiasing is largest for lenders who allocate at least half of their resources through the tool's suggestions.²²

The univariate results discussed thus far suggest automating lending choices reduces lenders' favoritism toward choosing in-group borrowers over out-group borrowers. However, systematically different time-varying trends between the Hindu and Muslim borrower populations might explain the differential shares of lending to borrowers of different religions over time. Note such differential trends should not be deemed relevant to the lending decision by the robo-advising tool to constitute a concern in our setting. Moreover, these trends should change the profitability and risk profiles of borrowers of different religions in the opposite direction for Hindu and Muslim lenders, which is barely plausible and reduces this concern substantially.

To assuage the relevance of this concern, in Table 2 we report the results for estimating the following multivariate specification:

$$\begin{aligned} Muslim\ Borrower_{i,j} = & \alpha + \beta\ Auto\ Invest_j + \gamma\ Hindu\ Lender_j + \\ & \delta\ Hindu\ Lender_j \times Auto\ Invest_j + \zeta\ \mathbf{x}_i + \epsilon_{i,j}, \end{aligned} \quad (1)$$

where $Muslim\ Borrower_{i,j}$ is equal to 1 if borrower i receives funding from lender j is Muslim, and 0 otherwise; $Auto\ Invest_j$ is equal to 1 if the lender made the loans after activating Auto Invest, and 0 otherwise; $Hindu\ Lender_j$ is equal to 1 if lender j is Hindu; and \mathbf{x}_i is a vector of loan-level characteristics that are direct proxies for the risk profiles of the loans that lenders extend to borrowers—loan maturity (measured in months), loan amount, and the interest rate associated with the loan. Importantly, these characteristics are *not* chosen by the lender; rather, the company's algorithm assigns them to borrowers when the loan requests are vetted, borrowers' risk profile is estimated, and requests are approved to be added to the platform.

In terms of statistical inference, in Table 2 we cluster standard errors at the lender level to allow for correlation across the lender-borrower matches that include the same lender. In Table A.2 of the Online Appendix, we show that the results are quite similar if we make different assumptions, such as allowing for double clustering by both lender and borrower, i.e. allowing for correlation also across the residuals of matches within the same loans, allowing for triple clustering at the lender, borrower, and month-of-issuance level, as well as for triple

²¹We assess the extensive margin only for Hindu lenders, because we do not have enough Muslim lenders in the sample to obtain a meaningful mass of them at each value of the percentage of funds allocated to Auto Invest.

²²We detect a mirroring pattern for Muslim borrowers (green dashed line), which is mechanical given that the sample includes only Hindu and Muslim borrowers.

clustering at the level of lender family communities (captured by sharing the same surnames), borrower family communities, and month of issuance.

The coefficients of interest in Table 2 are γ —the likelihood that Hindu lenders had Muslim borrowers in their portfolios relative to Muslim lenders—and δ —the change in the probability of Hindu lenders lending to Muslim borrowers after activating the robo-advising tool.

Column (1) of Table 2 reports the baseline estimates without adding any control variables. This specification assesses the statistical significance of the univariate results. The coefficient estimate $\hat{\gamma}$ is negative and statistically as well as economically significant, indicating Hindu lenders were about 5.8 pp less likely to lend to Muslim borrowers than Muslim lenders before using the robo-advising tool. The constant term—18%—captures the share of Muslim borrowers in Muslim lenders’ portfolios. The insignificant estimate of β shows that the share of Muslim borrowers on the platform is not systematically different before and after lenders access the tool. Moreover, $\hat{\delta}$ is positive and significant, indicating Auto Invest increases the likelihood that Hindu lenders lend to Muslim borrowers by about 4.5 pp. The lack of full debiasing is consistent with the fact that some lenders keep making some loan choices unassisted (see Figure 2).

In the second column of Table 2, we find the coefficient on the interaction between *Auto Invest_j* and *Hindu Lender_j* has the same statistical significance and point estimate as the baseline result, which excludes that the bias before the use of robo-advising (and hence, the extent of debiasing after using the tool) was driven by heterogeneity in objective proxies for the riskiness of borrowers, such as loans’ interest rates and maturity.

We then move on to restricting the variation within lenders by adding a lender fixed effect (column (3)). In this way, we assess the scope for discrimination after absorbing systematic time-invariant differences across lenders, such financial literacy, cognitive skills, and education levels. Our estimates stay almost identical. Absorbing time-varying shocks in column (4) also fails to influence the estimated coefficient $\hat{\delta}$.

The last two columns of Table 2 assess the intensive margin by comparing lenders who allocate less than 40% of their funds to Auto Invest (column (5)) and others (column (6)).²³ $\hat{\delta}$ is small and insignificant for lenders who do not use Auto Invest intensively while positive and significant for others.

4.2. Heterogeneous Hindu-Muslim Animus

To further assess the role of cultural biases in our results, based on earlier research in psychology and sociology, we isolate sources of heterogeneity in the extent to which the Hindu-Muslim conflict should be salient to lenders. Beyond providing an additional test for the effects of cultural biases, these heterogeneity tests further reduce concerns about the endogeneity of adoption: unobserved drivers of both adoption and behavioral changes are unlikely to vary systematically with variation in the salience of Hindu-Muslim conflict if unrelated to it.

²³We chose this threshold based on the shape of the relationship we described in Figure 2, which varies systematically above 40%, but the results are similar irrespective of the choice of threshold.

4.2.1 Hindu-Muslim Riots Across Cities We first consider spatial heterogeneity in Hindu-Muslim riots (Oza (2007); Ticku (2015)). Proximity to the violent riots and local media coverage are likely to make the Hindu-Muslim conflict more salient to exposed lenders (D’Acunto et al. (2019)). We thus test if lenders residing in states that have faced more Hindu-Muslim riots display a stronger in-group bias when making all choices unassisted.²⁴ Panel A of Figure 3 depicts the cross-state variation we employ based on the incidence of riots (Ticku (2015)): dark-green states (Gujarat, Maharashtra, Karnataka, and Uttar Pradesh) are those in which Hindu-Muslim riots have been most prevalent. To make the (many) coefficients of interest across subsamples easier to compare, we report the results of our multivariate heterogeneity tests in graphical form.

Figure 4 plots the relevant coefficients when estimating equation (1) separately for lenders in states with a low incidence of Hindu-Muslim riots during 1980-2000 (Ticku (2015)) and other lenders. The extent of bias was about twice as large for Hindu lenders in states with a higher incidence of Hindu-Muslim riots (6.4 pp), and statistically different from zero, but small and insignificant for other lenders. The left plot of Panel B shows $\hat{\delta}$, which captures the size of de-biasing after adopting Auto Invest for the same sets of lenders. Lenders in states with a high incidence of Hindu-Muslim riots debias by more (4.8 pp, $p < 0.01$). Statistically, though, we cannot reject the null that the estimated coefficients are equal across subsamples: the χ^2 statistic for a Wald-test of equality of the coefficients across the two subsamples is 0.18 in this first split.

4.2.2 Electoral Support for the BJP Across States We move on to consider a second dimension that earlier research has associated with the salience of the Hindu-Muslim conflict: the local vote share for the *Bharatiya Janata Party* (BJP). The BJP’s ideological roots have always been based on the notion of *hindutva*, which implies a coincidence between the spheres of Indian culture and traditional Hindu values (e.g., see Berglund (2004); Chhibber and Verma (2018); Prakash (2007); and Chidambaram et al. (2020)). The BJP is the result of mergers of post-independence parties and has shared a leading role with the *Indian National Congress* since independence (e.g., see Ziegfeld (2020)).²⁵ We exploit state-level variation in BJP vote shares to capture variation in Hindu lenders’ exposure to an ideological bias against Muslims.²⁶ To this aim, we compute the BJP candidates vote share in each election cycle from 1977 to 2015 and state and then the average BJP vote shares within states.²⁷ Panel B of Figure 3 reports the distribution of the average BJP vote share.

In the middle plots of Figure 4, the extent of bias against Muslim borrowers by Hindu lenders—Panel A—was

²⁴As discussed above, all lenders are mapped to an Indian state but not necessarily to a city. Moreover, regulation and policies are often implemented at the state level; for example, see the case of “anti-conversion laws” (Jenkins (2008); Dhattiwala and Biggs (2012)). Also, deep-rooted local cultural norms persistently relate to present-day interreligious violence in India (Jha (2014)).

²⁵In contrast to the BJP, the Indian National Congress has been proposing instances of secularization and has been less supportive of conflict between the Hindu majority and Muslim minority (e.g., see Ganguly (2003) and Verma (2016)).

²⁶We do not argue that the BJP vote share is a precise measure of the extent to which each lender supports the Hindu-Muslim conflict, but that, on average, it captures variation in the extent to which the conflict is salient to lenders.

²⁷We obtain data on the official number of voters, residents, and votes cast for various parties for all elections to the national congress and state-level elections from 1977 to 2015 from Bhavnani (2014), whose data set is based on information from the Indian electoral commission.

9.4 pp before robo-advising in high-BJP-vote share states, but only 3.5 pp for lenders in other states. After robo-advising, lenders in high-BJP-vote-share states debiased fully. A χ^2 test confirms that de-biasing is statistically higher in areas with high BJP support ($\chi^2=10.65$, $p<0.01$).

4.2.3 Cross-Cohort Exposure to Hindu-Muslim Conflict The last dimension we consider uses variation across lender cohorts. We exploit the fact that the electoral support for the BJP has increased substantially since the early 2000s (e.g., see Menon and Nigam (2007)) and reached its peak with the increased visibility and popularity of Narendra Modi after becoming Prime Minister of India in 2014 (Chhibber and Verma (2014)). BJP’s rise to national power has pushed the issue of Hindu-Muslim relations to the top of the agenda in Indian political discourse. We exploit variation across cohorts of lenders who were exposed to the rise of BJP during their formative years or were only exposed to this phenomenon in adulthood, when their political beliefs were likely already cemented (Malmendier and Nagel (2011)). Specifically, we compare lenders born after 1990 (24 or younger when Narendra Modi became Prime Minister) to others.²⁸

The right plots of Figure 4 report the results: pre-robo-advising bias is larger (7.1 pp) and statistically significant for lenders born after 1990, whereas the bias is small and statistically insignificant for others. Panel B shows that debiasing follow the same patterns. These two estimated effects are not only economically but also statistically different ($\chi^2=4.43$, $p<0.01$).

5 Stereotypical Discrimination: Shudra (Low-Caste) Borrowers

The second type of cultural bias we study is stereotypical discrimination—the fact that decision-makers systematically discriminate against certain social groups, because society attaches negative stereotypes to them (Becker (1957); Akerlof and Kranton (2000); D’Acunto et al. (2019); and Payne et al. (2019)). Stereotypical discrimination differs from in-group vs. out-group bias because even the members of the discriminated group discriminate against each other (Jost and Banaji (1994); Nosek et al. (2002); Pritlove et al. (2019)).²⁹

The Indian setting is well suited to the study of stereotypical discrimination because of the centuries-long negative stereotypes attached to members of lower castes. Based on a set of traditional and foundational Hindu writings, the Indian society has been divided into five broad social groups for centuries: four *varnas*, or castes, and a fifth group of “outcasts” or untouchables (Fox (1969); Dumont (1980); Sinha and Sinha (1967)).³⁰ In the traditional interpretation, these social groups have a strict hierarchical relation to one other. *Brahmins*, the highest caste, traditionally included Hindu clerics as well as those who dedicated their lives to studying and

²⁸The results are similar if we split the sample based on different years around 1990.

²⁹For example, research finds that not only men but also women tend to rate women’s quality and performance in leadership roles lower than men’s, even when objective measures of performance across genders are similar (Bertrand et al. (2005); Brownstein (2015)).

³⁰Here, we refer to the traditional scriptures-based notion of *varnas*. It does not coincide with the notion of *jati*, which is as richer and more complex sociological system based on which Hindus are further divided into other castes, tribes, and local social groups.

contemplative activities. The second caste (*Kshatriyas*) historically covered governmental and military positions. The third caste, the *Vaishyas*, included farmers, traders, and merchants.³¹ Against the three top varnas stands the (*Shudra*) caste, which has historically included laborers, peasants, and servants (Ambedkar (1947)). This caste was ranked lowest among other varnas and its members were employed in roles that benefited higher castes, which contributed to the diffusion of persistently negative stereotypes attached to (*Shudras*).³²

5.1. Variation in the Recognizability of Borrowers' Caste

Castes are not always easily recognizable based on observational characteristics such as names and surnames, physical appearance, and occupation (Muthukumar (2020)). Variation in caste recognizability provides a natural test for the salience of borrowers' caste in lenders' choices, which would not exist in settings in which social groups are easily recognizable. We exploit predictable variation in caste recognizability to test for the effects of stereotypical discrimination in lending. To capture the extent of recognizability of each borrower's caste, we build on an off-the-shelf algorithm by Bhagavatula et al. (2017) that assigns last names and other characteristics to castes and is designed to mimic the decision that a human would make based on the information at hand.³³ The procedure relies on two features of Indian castes: they are endogamous—marriages occur mainly between individuals that belong to the same caste—and last names are in part indicative of castes.

In the first step, the procedure collects data from 2.5 million individuals registered on online matrimonial agencies. This data contains information on individuals' last names and varna. The possibility of misreporting is virtually non-existent, because prospective spouses search for same-caste matches through matrimonial agencies. In the second step, the procedure assigns one or more castes to each of the last names. We compute the probability of a surname belonging to a given caste as the share of matrimonial website users in each state holding that surname who belong to the caste. In the third step, we assign a caste (and its probability) to each borrower and lender on our platform based on their last name and location. Figure A.6 in the Online Appendix plots the distribution of the probability of being *Shudra* for the borrowers in our sample for whom such a probability is strictly larger than zero—about 80% of all borrowers. Except for a right tail (17%) of borrowers whose probability of being *Shudra* is close to 1, the probability is distributed throughout the support.

5.2. Stereotypical Lending Before and After Robo-advising

The top graphs of Figure 5 plot the percentage of *Shudra* borrowers within all lenders' portfolios. The leftmost graph considers the full set of borrowers including those whose caste is barely recognizable. In this case, we

³¹Historians have emphasized similarities between the notion of *Vaishyas* and the *bourgeoisie* in pre-revolutionary France, for instance.

³²Note that members of the outcast group, the *Dalits*, have faced even stronger discrimination and segregation over the centuries (Maikkēl (1999)). Less than 2% of the borrowers in our platform are *Dalits*, which hinders us from analyzing this group empirically.

³³We thank Manaswini Bhalla for graciously running the algorithm developed in Bhagavatula et al. (2017) and Bhagavatula et al. (2018) on our data. Note that we could not reach out to lenders directly and ask them to assess borrowers' caste, because this priming procedure might have affected their lending behavior on the platform thus invalidating our analysis.

detect no difference in the share of *Shudra* borrowers before and after lenders use the robo-advising tool: both shares equal 31%, which is the share of *Shudras* in the borrower population. This result is consistent with the possibility that, when castes are barely identifiable, no discrimination exists against *Shudra* borrowers.

We obtain different results when we restrict the sample to subgroups in which the borrowers' caste is more easily recognizable. The pre-Auto Invest lending to *Shudra* borrowers is 27% in subfigure (b) and only 17% in subfigure (c), which only includes borrowers whose caste is highly recognizable. On the other hand, lending after robo adoption does not change with caste recognizability.

As discussed above, stereotypical discrimination also arises among members of the same group. And, indeed, in the bottom graphs of Figure 5, in which we restrict the sample to *Shudra* lenders, *Shudra* borrowers are even more discriminated, possibly because *Shudra* lenders can more easily recognize their similar than the average Indian lender. We confirm the robustness of these results with the following multivariate specification:

$$Shudra\ Borrower_{i,j} = \alpha + \beta\ Auto\ Invest_j + \zeta\ \mathbf{x}_i + \epsilon_{i,j}, \quad (2)$$

where the regressors are defined as in equation 1. Because we want to focus on several subsamples of the data based on caste recognizability for both the baseline and heterogeneity results, to facilitate the comparison of many relevant coefficients across specifications, we report these results in graphical form. Panel A of Figure 6 contains four bars representing different estimates of β in equation (2) across different subsamples of the data.

For the first bar to the left, we do not impose any sample restrictions. The β coefficient is positive (0.011) and statistically significant at the 5% level ($t=2.54$), indicating that using Auto Invest increases the probability of lending to Shudra borrowers. On the flip side, the result suggests that, before using Auto Invest, lenders were discriminating against Shudra borrowers. The second bar imposes that the probability of caste recognition is at least 40% and shows an effect similar to the first, both economically and statistically.

The last two columns restrict the estimating sample to lenders whose belonging to the Shudra caste is more easily recognizable—with probability larger than 70% and 80%, respectively. In this case, lending to Shudras after adopting Auto Invest more than doubles. These estimated effects are statistically significant at the 1% level, with t -statistics in excess of 3.5.

5.3. Heterogeneous Salience of Negative Stereotypes Attached to Lower Castes

Recent research documents substantial hatred crimes against members of lower castes, which has been heterogeneous across space (Sharma (2015); Bapuji and Chrispal (2020)). We conjecture that the stereotypical discrimination of *Shudra* borrowers might be higher in areas in which the conflict between higher and lower castes is more salient due to the higher incidence and reporting of acts of violence against lower castes. To operationalize this conjecture, we collect the number of crimes against lower-caste victims per 100,000 inhabitants

of Indian states from the annual report of the *National Crime Records Bureau (NCRB)* (NCRB (2019)). Figure 7 plots the cross-sectional variation in crimes against lower-caste victims.

We estimate equation (2) for lenders in Indian states above and below the median rate of crimes against lower castes per 100,000 inhabitants (18.8). We report the results in Panel B of Figure 6. For low levels of caste recognizability, we find no evidence of discrimination across space: the $\hat{\beta}$ coefficients are not statistically different from zero. As recognizability increases, we start to observe a wedge between the two areas, which becomes more and more marked at higher levels of recognizability. Above 70% recognizability, the share of Shudra borrowers chosen by lenders in areas with high crimes against Shudras increases by 2.7 pp with robo adoption, which is about 12% more of the average share of Shudra borrowers in lenders' portfolios before accessing the tool (unreported in the graph). The effect is less than half for other lenders and statistically insignificant. These patterns become even starker when we condition on higher levels of caste recognizability.

6 Cultural Debiasing and Lenders' Performance

Under cultural biases, lenders should reach deeper into the pool of borrowers of the religion (or caste) they favor when choosing to whom they lend their money. As a result, they should lend to less creditworthy borrowers of the preferred religion (or caste), whereas they should reject more creditworthy borrowers of the religion (or caste) against whom they are discriminating. We therefore conjecture that the performance of favored-group borrowers is worse than the performance of other borrowers before lender access the robo-advising tool, and this difference drop after adoption.

Our setting allows assessing this conjecture while abstracting for economic channels studied in earlier research that predict the opposite effect of cultural biases on performance. For instance, it abstracts from the potential screening and monitoring roles of in-group lending (Fisman, Paravisini, and Vig (2017) and Fisman, Sarkar, Skrastins, and Vig (2020)), because it connects lenders and borrowers all over India who do not live or work in the same small social environment.

6.1 Performance after Cultural Debiasing: Sign, Size, and Channels

We first consider a measure of the extensive margin of lending performance—whether borrowers default on their loans. Loan defaults are a commonly used measures of performance and have also been studied as an outcome for moral financial decision-making (e.g., see Guiso et al. (2013)). In our case, loan defaults allow testing whether in-group borrowers are less likely to default on in-group lenders relative to out-group borrowers due to stigma from the in-group community, which is the opposite of what cultural biases predict.

For high-interest loans, lenders' returns might be high even in case of defaults as long as the borrower pays high interests before defaulting. We will thus also estimate the effects of debiasing on loan returns.

6.1.1 Loan Defaults To assess loans' default, we consider all the loans in our sample that were closed by the last month we have available—March 2020—and categorize as defaulted those loans that had been delinquent for more than 90 days at the time of closure.³⁴

We start by considering in-group vs. out-group discrimination. For Hindu lenders, we estimate variations of the following specification by OLS:

$$\begin{aligned} \text{Delinquent Loan}_{ij} = & \alpha + \gamma \text{Muslim Borrower}_j \\ & + \delta \text{Muslim Borrower}_j \times \text{Auto Invest}_j \\ & + \theta \text{Hindu Borrower}_j \times \text{Auto Invest}_j + \zeta \mathbf{x}_i + \epsilon_{ij}, \end{aligned} \quad (3)$$

where $\text{Delinquent Loan}_{ij}$ is equal to 1 if the loan associated with borrower i and lender j is closed as delinquent, and all other variables are defined as discussed above.

On average, Muslim borrowers are less likely to default relative to Hindu borrowers in Hindu lenders' portfolios *before* Auto Invest is used ($\hat{\gamma} < 0$). We report this result both graphically and in table format.

In panel A of Figure 8, we plot the estimated coefficient $\hat{\gamma}$ for the full sample of Hindu lenders ("All") as well as when estimating Equation 3 separately for the subsamples based on the heterogeneous salience of Hindu-Muslim conflict we discussed in the previous section. Not only are Muslim borrowers on average 4.6 percentage-point less likely to default than Hindu borrowers before lenders access Auto Invest (about 16% of the average rate of default in the sample, i.e., 29%), but the size of this estimate is higher for lenders in states with a higher incidence of Hindu-Muslim riots, those in states with a higher support for the BJP, and younger lenders. Columns (1)-(2) of Table 3 confirm these results and show their statistical significance using two alternative specifications: column (1) does not allow for the interaction between auto-invest and borrower religion while column (2) does.

The second implication of biased discrimination is that, after Hindu lenders access Auto Invest, Hindu borrowers should reduce their likelihood of default by more than Muslim borrowers. That is, we should observe that $\theta < \delta$ when estimating Equation 3. Note that Auto Invest might provide additional benefits in terms of performance than those due to cultural debiasing. For instance, it might provide lenders with a more diversified portfolio of borrowers. For this reason, the likelihood of default by Muslim borrowers might also decline with robo-advising. Discrimination though undoubtedly predicts that the likelihood of default should decline more for Hindu than Muslim borrowers. Results in column (2) of Table 3 are consistent with this implication: Muslim borrowers reduce their likelihood of default by 7.3 percentage points whereas Hindu borrowers by 11.2 percentage

³⁴This definition is close to the regulatory definition by Reserve Bank of India (RBI) as well as to the economic definition of a defaulted loan, whereby the borrower did not pay back in full the loan's future value to the lender at the time the platform stopped monitoring the borrower. Note that we do not observe whether lenders engaged in litigation to collect borrowers' debentures after the loan was closed. If lenders were ultimately able to obtain a higher repayment than what is registered in the company's accounts, unfortunately we cannot know.

points—the effect is 53% larger for Hindu borrowers than for Muslim borrowers.

The third implication we bring to the data relates to the channels that should explain the difference in performance by Hindu and Muslim borrowers before and after robo-advising. If Hindu lenders were willing to dig deeper in the pool of Hindu borrowers and pick riskier borrowers in that group, because on our platform the highest-risk borrowers are disproportionately more likely to default than other borrowers (see Figure A.5 for the relation between ex-post default probabilities and interest rates), once we keep constant proxies for risk—interest rate, maturity, and loan amount—the differences in defaults across Hindu and Muslim borrowers should disappear. Moreover, the estimated improvement after robo-advising should converge for both groups to the improvement we estimated for Muslim borrowers without controls for risk, which under a discrimination story should not be driven by borrower-level risk characteristics.

Column (3) of Table 3 shows results that align with these predictions. Once we control for borrower-level risk, neither Muslim borrowers appear to default less than Hindu borrowers before robo-advising nor do Hindu borrowers appear as improving their performance by more than Muslim lenders after robo-advising. Also, the two point estimates for the decline in default rates—7.2 and 7 percentage points—are indistinguishable from the estimated improvement of Muslim borrowers in column (2), which is 7.3 percentage points.

Moving on to stereotypical discrimination, we assess whether Shudra borrowers were less likely to default than other borrowers before robo-advising, whether the improvement in default after robo-advising was larger for non-Shudra borrowers, and whether controlling for borrower-level risk characteristics reduces the differences in defaults.

Graphically, Panel B of Figure 8 shows that not only Shudra borrowers were less likely to default before lenders accessed Auto Invest, but this difference in default rates increased with borrowers' recognizability as members of the Shudra caste. The left plot shows that Shudra borrowers whose probability of being recognized was high (above 80%) were about 13 percentage points less likely to default than others, whereas when barely recognizable Shudra borrowers enter the sample the lower likelihood of default was only 4.5 percentage points. This result is starker for lenders who reside in states with a higher incidence of crime against Shudras (right plot): among highly recognizable Shudras, those picked by lenders in high-crime states were about twice less likely to default than other Shudra borrowers.

We also find evidence consistent with the second implication in column (5) of Table 3. Non-Shudra borrowers' performance improved by more than Shudra borrowers' performance, although the difference between these two effects (16 and 14.8 percentage points) is not statistically significant. Finally, column (6) shows that, once we control for borrower-level risk, both estimated effects decline and are aligned to about 10 percentage points.³⁵

³⁵Even in this case, before robo-advising, the lower likelihood of default of Shudra borrowers is lower although it does not become insignificant as we found for the case of Muslim borrowers in column (3).

6.1.2 Fraction of Loan Repaid The second dimension of performance we consider is the fraction of the overall amount due (including principal and interest) that the borrower pays back to the lender. This variable aims to capture the extent to which borrowers are willing to default on lenders at the intensive margin.

The share of repaid amount is capped at 1 for borrowers who repay their loan in full. In principle, the share can be as low as 0 if a borrower does not repay anything, but in our data the number of borrowers who repay less than 20% of their amount due is minimal, because the platform expels borrowers who pay less than 20% of the amount on any outstanding loans.

Panel A of Figure 9 focuses on in-group vs. out-group debiasing. The graph to the left plots the cumulative distribution functions (CDFs) for the share of the loans paid by Hindu borrowers (solid green line) and Muslim borrowers (orange dashed line) before Hindu lenders use Auto Invest. The evidence is consistent with the conjecture that Hindu lenders might dig deeper into the pool of Hindu borrowers than into the pool of Muslim borrowers when making choices on their own: Muslim borrowers repay larger shares of their amount due before the loan is closed. In fact, no Muslim borrowers picked by Hindu lenders paid less than 80% of the amount due in our sample. To the contrary, the repayment behavior of Hindu lenders is more volatile: about 20% of them repay less than 40% of the amounts due. Even when considering those who pay at least 80%, the CDF of Hindu borrowers is flatter than that of Muslim borrowers.³⁶

The right graph of Figure 9 plots the CDFs for the shares of amount due repaid by Hindu and Muslim borrowers to Hindu lenders after Hindu lenders start to use Auto Invest. Hindu borrowers improve disproportionately more than Muslim borrowers throughout the distribution. For instance, the share of Hindu borrowers who pay back more than 90% of their loans among those picked by Hindu lenders under Auto Invest increases to about 40% from 30% before Auto Invest.

We consider stereotypical discrimination in Panel B of Figure 9. The left graph splits borrowers whose probability of being Shudra is below the median (solid green line) and above the median (orange dashed line) before lenders access Auto Invest. Throughout the support, and especially in the left part of the distribution, borrowers who are more recognizable as Shudra tend to repay a higher fraction of their loans.

After accessing Auto Invest (right graph of Panel B of Figure 9), the distance between the CDFs decreases and the convergence is largely driven by an improvement of the low-probability Shudra borrowers' repayment behavior—the solid green line shifts to the right. Even for the case of stereotypical discrimination, lenders seem to impose higher standards on highly recognizable Shudra borrowers than on others when making unassisted choices. Once the robo-advising tool debiases lenders' stereotypical discrimination, the standards applied to Shudra and non-Shudra borrowers converge.

³⁶Note that the share of Hindu borrowers who repay intermediate amounts between 50% and 80% is negligible.

6.2. Loan Returns

Loan returns are the most appropriate dimension to assess the costs of cultural biases, because high-interest loans might provide lenders with high cash flows and hence high returns even if borrowers end up defaulting at some point and fail to repay the principal and interest due in full.

We obtained the detailed monthly-servicing information for a random sample of the loans on the platform, which, added to the sample restriction of only using closed loans as discussed in the analysis of defaults, unfortunately reduces the sample size for the analysis in this section substantially relative to the analyses in the first part of the paper. Although the smaller sample size reduces statistical power, we show that we can still reject the null hypotheses we bring to the data both economically and statistically.

6.2.1 Loan Returns Before and After Robo-advising We consider two types of specifications. First, like in the analysis of defaults, we assess the return changes at the loan level before and after lenders access Auto Invest in OLS specifications of the following form:

$$\begin{aligned} \text{Loan Return}_{ij} = & \alpha + \gamma \text{Muslim Borrower}_j \\ & + \delta \text{Muslim Borrower}_j \times \text{Auto Invest}_j \\ & + \theta \text{Hindu Borrower}_j \times \text{Auto Invest}_j + \zeta \mathbf{x}_i + \epsilon_{ij}, \end{aligned} \quad (4)$$

where Loan Return_{ij} is the return associated with the loan borrower i obtains from lender j at the time the loan servicing is closed (whether repaid in full or delinquent). We standardize returns to ease the interpretation of the magnitudes of the results. All other variables are defined as discussed above.

The predictions of cultural debiasing we test are analogous to those discussed for the case of loan defaults except that the predicted signs of coefficients are flipped for the case of loan returns. First, cultural biases would imply that Muslim borrowers provided Hindu lenders with higher returns than Hindu borrowers before robo-advice adoption ($\gamma > 0$), because Hindu lenders were systematically overestimating the expected returns they would obtain from Hindu borrowers and underestimating those from Muslim borrowers.

Second, after robo-advice adoption, Hindu borrowers' returns should increase by more than Muslim borrowers' returns ($\theta > \delta$). Note that, as we discussed in section 2.2., the tool is likely to improve lenders' loan portfolios above and beyond its effect on cultural biases. For instance, the tool might increase the diversification of lenders' loan portfolios. For this reason, Muslim borrowers' returns might also increase after access to robo-advising, but by a lower amount than Hindu borrowers' returns if cultural biases are present. The difference between the change in the average returns of Hindu and Muslim borrowers would bound the return value of cultural biases, because by construction any improvement in Muslim borrowers' returns cannot be explained by cultural biases but might be explained by other effects of robo-advising common to both groups of borrowers.

A third prediction we bring to the data is that, if Hindu borrowers' improvement is driven by a systematic change in the riskiness of Hindu borrowers that are matched with Hindu lenders by the robo-advising tool, once we keep constant proxies for the riskiness of loans, the differences in loan returns across Hindu and Muslim borrowers should be muted. Otherwise, channels other than a systematic difference in the riskiness of the Hindu and Muslim borrowers chosen by Hindu lenders when unassisted might explain the differential performance.

In Table 4, we report the results when bringing these three predictions to the data. Columns (1)-(3) consider in-group vs. out-group discrimination. We find that, on average, Muslim borrowers' loans earned Hindu lenders higher returns before lenders adopted to robo-advising—about half a standard deviation higher returns, or 12.5 percentage-point higher returns. At the same time, Hindu lenders' returns increased on average after adoption by about 20% of a standard deviation, or 5 percentage points.

Column (2) reveals that the improvement in Hindu lenders' returns after robo-advising adoption is almost exclusively driven by an increase in the returns earned on loans issued to Hindu borrowers (22% of a standard deviation higher, or 5.5 percentage points), whereas Muslim borrowers' loan returns, which were already performing better before adoption, did not change in any economically or statistically significant manner after robo-advice adoption. These pattern are consistent with the conjectures about performance under in-group vs. out-group discrimination on the part of Hindu lenders.

Moreover, consistent with the third conjecture, once we absorb dimensions that capture the riskiness of loans in column (3), we fail to detect systematic differences in the change of loan returns between Hindu and Muslim borrowers after lenders adopt the robo-advising tool. This result suggests that the asymmetric improvement in performance across groups of borrowers we detected in column (2) is explained by a systematic change in the riskiness of the two pools of borrowers matched to lenders after robo-advice adoption.

We then move on to assess the effects of debiasing stereotypical discrimination on loan returns. Column (4) of Table 4 reports the results for the baseline specification. Even in this case, the results align with the predictions of cultural debiasing. On the one hand, Shudra borrowers, who were discriminated against before robo-advice adoption, provided higher loan returns to lenders when lenders made their choices unassisted. Shudra borrower returns were almost one quarter of a standard deviation higher than non-Shudra borrower's returns, which in the sample of Hindu borrower we use in this analysis corresponds to about 6.5 percentage-point higher returns.

Moreover, lenders increased their average returns after accessing robo-advice, relative to before. Column (5) compares the contribution of Shudra and non-Shudra borrowers in the increase in average returns after robo-advice. In fact, in this case we do not detect as clear a pattern as we did with in-group bias—we do not find that in column (5) the increase in the returns of Shudra borrowers' loan is substantially smaller than the increase for non-Shudra borrowers. If anything, the point estimate is larger for Shudra borrowers.

When we absorb the proxies for the riskiness of loans in column (6), consistent with our third prediction, we find that the estimated improvements in returns after robo-advising are similar in terms of magnitude and

statistical significance across the two groups of borrowers. This result is consistent with the possibility that a systematically different risk profile of different groups of borrowers after robo-advising drives the improvement in returns.

6.2.2 Loan Return Distribution Before and After Robo-advising The results on the average returns of loans issued to different groups of borrowers before and after robo-advice adoption are broadly in line with the predictions of cultural debiasing.

Yet, the results based on sample averages do not allow us to pin down which parts of the distribution of loan returns are responsible for the improvement. If lenders tended to pick very risky borrowers from their favorite group before accessing robo-advising, and the tool does not pick such risky borrowers, we should for instance observe that the tool improves performance by eliminating the left tails of negative performance by favorite groups' borrower loans. If instead lenders tended to pick riskier in-group borrowers throughout the distribution, we should not detect that any tails improved systematically more than the other.

To assess which parts of the loan return distribution drive our results, we follow D'Acunto and Rossi (2019) and estimate a set of quantile regressions of the following form:

$$Q_{\tau}(Returns_{i,j}) = \alpha(\tau) + \beta(\tau) Auto Invest_j + X'_{i,j} \zeta(\tau) + \epsilon_{i,j}, \quad (5)$$

whose outcome variable is quantile Q_{τ} of the distribution of the return associated with borrower i and j throughout the sample period. All other variables are defined as in equation (3).

To interpret the estimates of equation (5), consider the special case of the median, which is the 50th percentile of the distribution. The coefficient $\hat{\beta}(50)$ estimates that the median return was $\hat{\beta}(50)$ higher after lender j moved to Auto Invest relative to before. A positive $\hat{\beta}(50)$ would suggest the median of the distribution has shifted to the right. The advantage of estimating quantile regressions is that we can assess how the whole intensive margin (distribution) has changed rather than focusing on specific moments, such as the conditional mean.

We report the results for estimating the baseline OLS specification in columns (1)-(2) of Panel A of Table 5 and the quantile regression estimates in columns (3)-(8).

In the first line of each column, we focus on specifications that *do not* control for the risk characteristics of loans, such as interest rates and amount, so that we allow for the possibility that difference between the returns before and after access to the robo-advising tool are driven by a different distribution of the risk of loans in each lender's portfolio.

Based on the hypothesis that Hindu lenders picked in-group borrowers of worse quality before moving to Auto Invest, we should find that they improve their performance after accessing Auto Invest. The same should be true for Muslim lenders, who, under cultural biases, should have picked in-group borrowers of worse quality than otherwise available out-group borrowers. And, indeed, the first line of columns (1)-(2) of Panel A of Table

5 reveals that both Hindu and Muslim lenders improve their performance in terms of average loan returns.

Moving on to the quantile regression results (columns (3)-(8)), they reveal that most of the returns' improvement is driven by substantially higher returns in the left tail of the distribution, as can be seen by the fact that the size of the estimated coefficients is larger for the 25th and 50th percentiles of the return distribution and declines as we move towards the right (the 75th percentile).

In the second line of each column, we add borrower-level loan characteristics as controls to further assess whether any changes before and after Auto Invest might be driven by systematically different choices in terms of borrower risk—for instance, whether lenders tended to choose systematically riskier in-group borrowers relative to out-group borrowers when making unassisted choices.

The evidence is consistent with a risk channel that explains the performance improvement lenders enjoy when moving to Auto Invest. Once we absorb differences in the riskiness of loans, conditional returns do not differ when lenders make choices on their own or when the robo-advising tool makes choices on their behalf, either on average or in terms of the different parts of the distribution of loan returns, including the left tail of the distribution.

We perform the same analysis for the case of stereotypical discrimination. The first line of Panel B of Table 5 estimates the relationship for Shudra borrowers without controlling for the loans' risk characteristics, whereas the second line estimates the relationship conditional on borrower-level risk proxies.

The patterns we uncover for stereotypical discrimination are the same as for the case of in-group vs. out-group discrimination. The bulk of the improvement in returns is driven by the elimination of a left tail of low-return borrowers. Moreover, once we estimate the differential returns conditional on loans' risk measures, we detect no systematic differences before and after Auto Invest, which is consistent with the conjecture that the return improvement under Auto Invest is due to the selection of a less risky set of borrowers and especially to avoiding a left tail of high-risk low-return borrowers.

Ultimately, our analysis of performance suggests that eliminating cultural biases improves lenders' performance and this improvement is driven by a change in the composition of the borrower pool that reduces lenders' risk exposure. This reduction is largely driven by eliminating a left tail of low-return loans. It is consistent with the conjecture that culturally-biased lenders dig deeper into the pool of in-group borrowers and hence select riskier (and lower return) borrowers when making unassisted choices.

7 Quantifying the Cost of Cultural Biases: Lender-Level Returns

In this section, we propose a quantification of the aggregate effects of cultural biases at the aggregate lender level rather than a reduced form analysis at the lender-loan level, on which our analysis has focused thus far.

We start by computing the change in the returns each lender made on their overall invested amounts before

and after accessing the automated robo-advising tool, both in general as well as separately for the amounts lenders disbursed to in-group vs. out group religions as well as to Shudra vs. non-Shudra borrowers. For each lender, we define the total return on the investment before and after Auto Invest as follows:

$$Lender\ Total\ Return_{i,t} = 100 \times \frac{\sum_j Amount\ Disbursed_{i,j,t} \times Loan\ Return_{j,t}}{\sum_j Amount\ Disbursed_{i,j,t}}, \quad (6)$$

where $Lender\ Total\ Return_{i,t}$ is the overall return on the aggregate investment made by lender i earned either before access to Auto Invest ($t = PRE$) or after access to Auto Invest ($t = POST$); $Amount\ Disbursed_{i,j,t}$ is the amount (in rupees) lender i disbursed to loan j , which was issued either before or after access to Auto Invest (t); $Loan\ Return_{j,t}$ is the return of loan j to which lender i contributed.

The quantities defined by equation (6) thus capture the total returns to invested amounts lenders obtained before and after using Auto Invest. We then compute the lender-level change in total return across the two conditions:

$$Change\ Lender\ Return_i = Lender\ Total\ Return_{i,POST} - Lender\ Total\ Return_{i,PRE}, \quad (7)$$

where a positive value indicates that lender i earned a higher total return on their investment after accessing Auto Invest relative to before, and a negative value the opposite.

In Figure 10, we plot the density of the distributions of $Change\ Lender\ Return_i$ for lenders in the in-group vs. out-group discrimination sample (panel A) and for those in the stereotypical discrimination sample (panel B). For each distribution, we indicate the mean of the distribution with a solid vertical line and compare it to a dashed vertical line that indicates no change in returns. We find that the average lender in the in-group vs. out-group discrimination sample earned a 4.5-percentage-point higher total return after accessing Auto Invest relative to before, whereas the average lender in the stereotypical discrimination sample earn a 7.3-percentage-point higher return.

Our regression results at the lender-borrower-loan level suggested that most of the improvement in terms of loan defaults and repayment behavior derived from borrowers belonging to demographics that lenders tended to favor before moving to Auto Invest. As expected, this pattern holds in terms of lender-level total returns. For instance, if we compute the change in lender returns defined in equation 6 separately for Hindu lenders based on the amounts they lent to Hindu borrowers or to Muslim borrowers before and after Auto Invest, we find Hindu lenders on average gained a 4.3-percentage-point higher return on the amounts disbursed to Hindu borrowers, whereas they actually on average “lost” (an insignificant) 0.5 percentage points in returns to the amounts disbursed to Muslim borrowers. Virtually the whole improvement in the average lender-level returns derive from higher returns earned on the loans disbursed to favored demographic groups.

To capture the rupee-level change in performance and hence the lender-level and aggregate value of cultural

biases, we need a measure of lender-level performance in which returns are value-weighted—they are weighted by the rupee amounts each lender disburses to borrowers on the platform. A challenge to define such measure is that the amounts lenders disbursed before and after accessing Auto Invest might differ for many reasons, which are potentially unrelated to cultural biases, lender characteristics, or platform characteristics. We therefore compute the following:

$$\begin{aligned}
 & \text{Change Lender Value}_i \\
 &= [Amount\ Disbursed_{i,POST} \times Return_{i,POST} - Amount\ Disbursed_{i,PRE} \times Return_{i,PRE}] \\
 &\quad - [Amount\ Disbursed_{i,POST} \times Return_{i,PRE} - Amount\ Disbursed_{i,PRE} \times Return_{i,POST}]. \quad (8)
 \end{aligned}$$

The expression defined in equation (8) allows us to purge the difference in total rupee-value earnings at the lender level merely due to the fact that lenders might disburse different amounts before and after accessing Auto Invest, irrespective of the returns they earn in the two periods.

Note that equation (8) is equivalent to $Amount\ Disbursed_{i,PRE} \times Change\ Lender\ Return_i$, and hence this measure can be interpreted as the change in lender-level return after accessing Auto Invest relative to before weighted by the rupee amount the lender disbursed on the platform before accessing Auto Invest, which cannot have been determined by the returns the lender earned after starting to use the tool. Ultimately, this value captures the incremental rupee amount each lender would have earned in the period before accessing Auto Invest had they enjoyed the same return they did enjoy after moving to Auto Invest rather than the actual return they earned over that period.

We find that the average change in lenders' rupee value for lenders in the in-group vs. out-group discrimination sample is ₹457, which is about 6% of the average amount of resources disbursed by each lender in the period before accessing Auto Invest (₹7,543).³⁷ When we consider the stereotypical discrimination sample, we find that the average change is higher: ₹862, which represents about 12% of the average amount lenders in this sample disbursed before accessing Auto Invest (₹7,091). Overall, the estimated cost of cultural biases appears to be sizable relative to the amounts lenders invested for both types of biases we study.

Note that the calculation proposed in equation (8) does not account for the possibility that the rupee amounts disbursed to each demographic group would have been different had cultural biases not influenced lenders' choices in the period before accessing Auto Invest.

To assess whether accounting for this difference could influence our estimates of the costs of cultural biases substantially, we thus propose a modified version of the definition in equation (8) for robustness purposes. We separate the amounts disbursed and returns earned in the pre- and post-periods by Hindu lenders coming from Hindu and Muslim borrowers for the case of in-group vs. out-group discrimination and from Shudra and

³⁷Note that the average size of loans on the platform is substantially larger (₹90,523), because, as we discussed when introducing our setting, each loan borrowers received is financed by multiple lenders.

non-Shudra borrowers for the case of stereotypical discrimination.

To obtain a counterfactual for the first period, we split the amounts Hindu lenders disbursed through the platform in the pre-period among the two borrower groups (Hindus vs. Muslims and Shudra vs. non-Shudra) based on the shares of the post-period funds lenders attribute to each group rather than the true shares they attributed to them in the pre-period. In this way, we keep the true total amounts lenders disbursed through the platform in the pre-period fixed, but we assume that, if lenders faced no cultural biases in the pre-period, they would have split such amounts between borrowers based on the post-period shares.

Economically, the change in the shares of Hindu borrowers in Hindu lenders' portfolios cannot be large given that the majority of borrowers on the platform are Hindu. Indeed, the share of Hindu borrowers in Hindu lenders' portfolio moves from 89.7% in the pre-period to 88.7% in the post-period. For this reason, we would not expect that the correction we propose in this second method will deliver estimates for the cost of in-group vs. out-group bias that are substantially different from those discussed above. And, indeed, we find that the average lender-level change in earnings for Hindu lenders based on this correction is ₹382, which is quite similar to the value estimated above.

The same correction, instead, is likely to imply a larger estimated cost of the bias for the case of stereotypical discrimination, because the share of Shudra borrowers in Hindu lenders' portfolios goes from 41.8% in the pre-period to 46.2% in the post-period. When accounting for this composition change in the pre-period, we obtain an average cost of bias of ₹2,254 at the lender level, which is more than twice as large as the estimate that does not use this correction.

8 Conclusions

We propose a field setting to test for and quantify the extent and effects of cultural biases in high-stake economic choices. We detect evidence that in-group vs. out-group discrimination and stereotypical discrimination are prevalent and economically sizable. These forms of discrimination make discriminating individuals—in our case, lenders—worse off in terms of consumption utility, because by discriminating they finance loans by borrowers who perform worse than other (discriminated) borrowers available on the lending platform.

Our test and results suggest a new role that robo-advising tools might have for future research in economics—they can provide a benchmark to assess the types and sizes of agents' biases in decision-making. For instance, by coding robo-advising tools that embed different forms of biases or rules of thumb detected in the literature and by comparing decision-makers' unassisted choices with those they make after accessing such tools, one could disentangle the role of alternative biases and quantify them.

Moreover, future research should study whether exposure to robo-advising suggestions lets decision-makers learn about optimal choices and develop rules of thumb that can assist them also when a robo-advisor is not

available. Interactive robo-advising tools might teach borrowers how to create goal-setting strategies (Gargano and Rossi (2020)) or provide just-in-time simple financial literacy contents (Burke et al. (2021)).

More broadly, our results beget additional work across several fields on understanding how human and machine-based decision-making interact and complement or substitute each other in a world in which the combination of the two forms of decision-making is becoming ubiquitous in all daily economic decision-making problems agents face.

References

- Adams, P. D., S. Hunt, C. Palmer, and R. Zaliauskas (2019). Testing the effectiveness of consumer financial disclosure: Experimental evidence from savings accounts. Technical report, National Bureau of Economic Research.
- Agarwal, S., S. Alok, P. Ghosh, S. Ghosh, T. Piskorski, and A. Seru (2017). Banking the unbanked: What do 255 million new bank accounts reveal about financial access? *Columbia Business School Research Paper* (17-12).
- Agarwal, S., S. Alok, P. Ghosh, and S. Gupta (2019). Fintech and credit scoring for the millennials: evidence using mobile and social footprints. *Available at SSRN 3507827*.
- Akerlof, G. A. and R. E. Kranton (2000). Economics and identity. *The quarterly journal of economics* 115(3), 715–753.
- Alesina, A., P. Giuliano, and N. Nunn (2013). On the origins of gender roles: Women and the plough. *Quarterly Journal of Economics* 2(128), 469–530.
- Ambedkar, B. R. (1947). *Who were the Shudras?*, Volume 1. Ssoft Group, INDIA.
- Avenancio-León, C. and T. Howard (2019). The assessment gap: Racial inequalities in property taxation. *Available at SSRN 3465010*.
- Balyuk, T. (2019). Financial innovation and borrowers: Evidence from peer-to-peer lending. *Rotman School of Management Working Paper* (2802220).
- Balyuk, T. and S. Davydenko (2019). Reintermediation in fintech: Evidence from online lending. *Michael J. Brennan Irish Finance Working Paper Series Research Paper* (18-17).
- Bapuji, H. and S. Chrispal (2020). Understanding economic inequality through the lens of caste. *Journal of Business Ethics* 162(3), 533–551.
- Bartlett, R., A. Morse, R. Stanton, and N. Wallace (2019). Consumer-lending discrimination in the fintech era. Technical report, National Bureau of Economic Research.
- Becker, G. (1957). The economics of discrimination. *University of Chicago Press*.
- Berglund, H. (2004). Religion and nationalism: Politics of bjp. *Economic and Political Weekly*, 1064–1070.
- Bertrand, M., D. Chugh, and S. Mullainathan (2005). Implicit discrimination. *American Economic Review* 95(2), 94–98.
- Bhagavatula, S., M. Bhalla, M. Goel, and B. Vissa (2017). The business of religion and caste in india. *Working Paper*.
- Bhagavatula, S., M. Bhalla, M. Goel, and B. Vissa (2018). Cultural homophily in corporate boards and firm. *Working Paper*.
- Bhat, M. M. A. (2020). Hate crimes in india.
- Bhattacharya, U., A. Hackethal, S. Kaesler, B. Loos, and S. Meyer (2012). Is unbiased financial advice to retail investors sufficient? answers from a large field study. *The Review of Financial Studies* 25(4), 975–1032.
- Bhavnani, R. R. (2014). India National and State Election Dataset.
- Bhutta, N., A. Hizmo, and D. Ringo (2021). How much does racial bias affect mortgage lending? evidence from human and algorithmic credit decisions. *Evidence from Human and Algorithmic Credit Decisions (July 15, 2021)*.
- Bohren, J. A., K. Haggag, A. Imas, and D. G. Pope (2019). Inaccurate statistical discrimination. Technical report, National Bureau of Economic Research.
- Borjas, G. J. and M. S. Goldberg (1978). Biased screening and discrimination in the labor market. *The american economic review* 68(5), 918–922.
- Brass, P. R. (2011). *The production of Hindu-Muslim violence in contemporary India*. University of Washington Press.

- Breza, E. (2019). Peer effects and loan repayment: Evidence from the krishna default crisis. *Working Paper*.
- Brownstein, M. (2015). Implicit bias.
- Burke, C., I. Chak, K. Croxson, F. D'Acunto, R. Jonathan, A. Rossi, and J. Shaw (2021). Improving households' debt management with robo-advising. *Working Paper*.
- Bursztyn, L., B. Ferman, S. Fiorin, M. Kanz, and G. Rao (2018). Status goods: Experimental evidence from platinum credit cards. *The Quarterly Journal of Economics* 133(3), 1561–1595.
- Bursztyn, L., S. Fiorin, D. Gottlieb, and M. Kanz (2019). Moral incentives in credit card debt repayment: evidence from a field experiment. *Journal of Political Economy* 127(4), 1641–1683.
- Chhibber, P. and R. Verma (2014). The bjp's 2014'modi wave': An ideological consolidation of the right. *Economic and Political Weekly*, 50–56.
- Chhibber, P. K. and R. Verma (2018). *Ideology and identity: The changing party systems of India*. Oxford University Press.
- Chidambaram, S. et al. (2020). The civil society roots of bjp's majoritarian nationalism.
- Chiu, L., B. Wolfe, and W. Yoo (2018). Do fintech lenders fairly allocate loans among investors? quid pro quo and regulatory scrutiny in marketplace lending. *Quid Pro Quo and Regulatory Scrutiny in Marketplace Lending (September 30, 2018)*.
- Cornell, B. and I. Welch (1996). Culture, information, and screening discrimination. *Journal of political Economy* 104(3), 542–571.
- Cowgill, B. and C. E. Tucker (2020). Algorithmic fairness and economics. *The Journal of Economic Perspectives*.
- Crouzet, N., A. Gupta, and F. Mezzanotti (2019). Shocks and technology adoption: Evidence from electronic payment systems.
- D'Acunto, F. (2019). Tear down this wall street: The effect of anti-market ideology on investment decisions. *Available at SSRN 2705545*.
- D'Acunto, F., D. Hoang, M. Paloviita, and M. Weber (2019a). Human frictions in the transmission of economic policy. *Working Paper* (2019-07).
- D'Acunto, F., D. Hoang, M. Paloviita, and M. Weber (2019b). Iq, expectations, and choice. Technical report, National Bureau of Economic Research.
- D'Acunto, F., N. Prabhala, and A. G. Rossi (2019). The promises and pitfalls of robo-advising. *The Review of Financial Studies* 32(5), 1983–2020.
- D'Acunto, F., M. Prokopczuk, and M. Weber (2019). Historical antisemitism, ethnic specialization, and financial development. *The Review of Economic Studies* 86(3), 1170–1206.
- D'Acunto, F., T. Rauter, C. K. Scheuch, and M. Weber (2020). Perceived precautionary savings motives: Evidence from fintech. Technical report, National Bureau of Economic Research.
- D'Acunto, F., A. Rossi, and M. Weber (2019). Crowdsourcing peer information to change spending behavior. *Chicago Booth Research Paper* (19-09).
- D'Acunto, F. and A. G. Rossi (2020). Robo-advising. *Handbook of Technological Finance*.
- D'Acunto, F. and A. G. Rossi (forthcoming). Regressive mortgage credit redistribution in the post-crisis era. *Review of Financial Studies*.
- D'Acunto, F., M. Weber, and J. Xie (2019). Punish one, teach a hundred: The sobering effect of punishment on the unpunished. *Fama-Miller Working Paper*, 19–06.
- D'Andrea, A. and N. Limodio (2019). High-speed internet, financial technology and banking in africa. *BAFFI CAREFIN Centre Research Paper* (2019-124).
- Dhattiwala, R. and M. Biggs (2012). The political logic of ethnic violence: The anti-muslim pogrom in gujarat, 2002. *Politics & Society* 40(4), 483–516.

- Diep-Nguyen, H. and H. Dang (2019). Social collateral. Technical report, Working Paper.
- Dobbie, W., A. Liberman, D. Paravisini, and V. S. Pathania (2020). Measuring bias in consumer lending. *Review of Economic Studies*.
- Drexler, A. and A. Schoar (2014). Do relationships matter? evidence from loan officer turnover. *Management Science* 60(11), 2722–2736.
- Duarte, J., S. Siegel, and L. Young (2012). Trust and credit: The role of appearance in peer-to-peer lending. *The Review of Financial Studies* 25(8), 2455–2484.
- Dumont, L. (1980). *Homo hierarchicus: The caste system and its implications*. University of Chicago Press.
- Engineer, A. (1997). *Communal riots in post-independence India*. Universities Press.
- Fisman, R., D. Paravisini, and V. Vig (2017). Cultural proximity and loan outcomes. *American Economic Review* 107(2), 457–92.
- Fisman, R., A. Sarkar, J. Skrastins, and V. Vig (2020). Experience of communal conflicts and intergroup lending. *Journal of Political Economy* 128(9), 3346–3375.
- Fox, R. G. (1969). Varna schemes and ideological integration in indian society. *Comparative Studies in Society and History* 11(1), 27–45.
- Fuster, A., P. Goldsmith-Pinkham, T. Ramadorai, and A. Walther (2017). Predictably unequal? the effects of machine learning on credit markets. Technical report, CEPR Discussion Papers.
- Ganguly, S. (2003). The crisis of indian secularism. *Journal of Democracy* 14(4), 11–25.
- Gargano, A. and A. G. Rossi (2020). There’s an app for that: Goal-setting and saving in the fintech era. *Available at SSRN*.
- Gargano, A. and A. G. Rossi (2021). Correcting present bias in saving decisions with fintech.
- Ghassem-Fachandi, P. (2012). *Pogrom in Gujarat: Hindu nationalism and anti-Muslim violence in India*. Princeton University Press.
- Graff, V., J. Galonnier, et al. (2012). Hindu-muslim communal riots in india ii (1986-2011). *Online Encyclopedia of Mass Violence, Paris, Sciences Po, CERI*.
- Guiso, L., P. Sapienza, and L. Zingales (2006). Does culture affect economic outcomes? *Journal of Economic perspectives* 20(2), 23–48.
- Guiso, L., P. Sapienza, and L. Zingales (2009). Cultural biases in economic exchange? *The quarterly journal of economics* 124(3), 1095–1131.
- Guiso, L., P. Sapienza, and L. Zingales (2013). The determinants of attitudes toward strategic default on mortgages. *The Journal of Finance* 68(4), 1473–1515.
- Hewstone, M., M. Rubin, and H. Willis (2002). Intergroup bias. *Annual review of psychology* 53(1), 575–604.
- Higgins, S. (2019). Financial technology adoption. *Working Paper*.
- Hjort, J. (2014). Ethnic divisions and production in firms. *The Quarterly Journal of Economics* 129(4), 1899–1946.
- Hjort, J., C. Song, and C. Yenkey (2021). Ethnic investing and the value of firms.
- Iyer, R., A. I. Khwaja, E. F. Luttmer, and K. Shue (2016). Screening peers softly: Inferring the quality of small borrowers. *Management Science* 62(6), 1554–1577.
- Jenkins, L. D. (2008). Legal limits on religious conversion in india. *Law and contemporary problems* 71(2), 109–127.
- Jenkins, R. (2014). *Social identity*. Routledge.

- Jha, S. (2014). ‘unfinished business’: Historic complementarities, political competition and ethnic violence in gujarat. *Journal of Economic Behavior & Organization* 104, 18–36.
- Jost, J. T. and M. R. Banaji (1994). The role of stereotyping in system-justification and the production of false consciousness. *British journal of social psychology* 33(1), 1–27.
- Karlan, D., M. Mobius, T. Rosenblat, and A. Szeidl (2009). Trust and social collateral. *The Quarterly Journal of Economics* 124(3), 1307–1361.
- Kaul, N. (2017). Rise of the political right in india: Hindutva-development mix, modi myth, and dualities. *Journal of Labor and Society* 20(4), 523–548.
- Kleinberg, J., H. Lakkaraju, J. Leskovec, J. Ludwig, and S. Mullainathan (2018). Human decisions and machine predictions. *The quarterly journal of economics* 133(1), 237–293.
- Lorenzen, D. N. (1999). Who invented hinduism? *Comparative studies in society and history* 41(4), 630–659.
- Maikkēl, E. E. (1999). *Untouchable: Dalits in Modern India*. Lynne Rienner Publishers.
- Malmendier, U. and S. Nagel (2011). Depression babies: do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics* 126(1), 373–416.
- Menon, N. and A. Nigam (2007). *Power and contestation: India since 1989*. Zed Books.
- Mitra, A. and D. Ray (2014). Implications of an economic theory of conflict: Hindu-muslim violence in india. *Journal of Political Economy* 122(4), 719–765.
- Mobius, M., T. Rosenblat, and Q. Wang (2016). Ethnic discrimination: Evidence from china. *European Economic Review* 90, 165–177.
- Muthukumar, A. (2020). Casteism camouflaged as culture. *Harvard Political Review*.
- Naaraayanan, S. L. (2019). Women’s inheritance rights and entrepreneurship gender gap. Technical report.
- NCRB (2019). Crime in india—statistics, volume ii. *National Crime Records Bureau of India*.
- Nosek, B. A., M. R. Banaji, and A. G. Greenwald (2002). Harvesting implicit group attitudes and beliefs from a demonstration web site. *Group Dynamics: Theory, Research, and Practice* 6(1), 101.
- Oza, R. (2007). The geography of hindu right-wing violence in india. *Violent geographies: Fear, terror, and political violence*, 153–74.
- Paravisini, D., V. Rappoport, and E. Ravina (2017). Risk aversion and wealth: Evidence from person-to-person lending portfolios. *Management Science* 63(2), 279–297.
- Parsons, C. A., J. Sulaeman, M. C. Yates, and D. S. Hamermesh (2011). Strike three: Discrimination, incentives, and evaluation. *American Economic Review* 101(4), 1410–35.
- Payne, B. K., H. A. Vuletich, and J. L. Brown-Iannuzzi (2019). Historical roots of implicit bias in slavery. *Proceedings of the National Academy of Sciences* 116(24), 11693–11698.
- Phelps, E. S. (1972). The statistical theory of racism and sexism. *The american economic review* 62(4), 659–661.
- Prakash, G. (2007). Secular nationalism, hindutva, and the minority. *The crisis of secularism in India*, 177–188.
- Pritlove, C., C. Juando-Prats, K. Ala-Leppilampi, and J. A. Parsons (2019). The good, the bad, and the ugly of implicit bias. *The Lancet* 393(10171), 502–504.
- Pursiainen, V. (2020). Cultural biases in equity analysis. *Journal of Finance, Forthcoming*.
- Rambachan, A., J. Kleinberg, J. Ludwig, and S. Mullainathan (2020). An economic perspective on algorithmic fairness. In *AEA Papers and Proceedings*, Volume 110, pp. 91–95.
- Rao, G. (2019). Familiarity does not breed contempt: Generosity, discrimination, and diversity in delhi schools. *American Economic Review* 109(3), 774–809.

- Ravina, E. (2019). Love & loans: The effect of beauty and personal characteristics in credit markets. *Available at SSRN 1107307*.
- Rossi, A. G. and S. P. Utkus (2020). Who benefits from robo-advising? evidence from machine learning. *Evidence from Machine Learning (March 10, 2020)*.
- Schoar, A. (2012). The personal side of relationship banking. *Available at SSRN 2024653*.
- Sharma, S. (2015). Caste-based crimes and economic status: Evidence from india. *Journal of comparative economics* 43(1), 204–226.
- Simon, H. A. (1993). Altruism and economics. *The American Economic Review* 83(2), 156–161.
- Sinha, G. S. and R. C. Sinha (1967). Exploration in caste stereotypes. *Social Forces* 46(1), 42–47.
- Tajfel, H., J. C. Turner, W. G. Austin, and S. Worchel (1979). An integrative theory of intergroup conflict. *Organizational identity: A reader* 56, 65.
- Tang, H. (2019). Peer-to-peer lenders versus banks: substitutes or complements? *The Review of Financial Studies* 32(5), 1900–1938.
- Tantri, P. (2021). Fintech for the poor: Financial intermediation without discrimination. *Review of Finance* 25(2), 561–593.
- Ticku, R. (2015). Riot rewards? study of bjp’s electoral performance and hindu muslim riots. Technical report, The Graduate Institute of International and Development Studies.
- Vallee, B. and Y. Zeng (2019). Marketplace lending: A new banking paradigm? *The Review of Financial Studies* 32(5), 1939–1982.
- Verma, V. (2016). Secularism in india. *The Oxford Handbook of Secularism*, 214.
- Ziegfeld, A. (2020). A new dominant party in india? putting the 2019 bjp victory into comparative and historical perspective. *India Review* 19(2), 136–152.

**Figure 2: Lending to In-Group vs. Out-Group Borrowers:
Extent of Debiasing (Intensive Margin)**

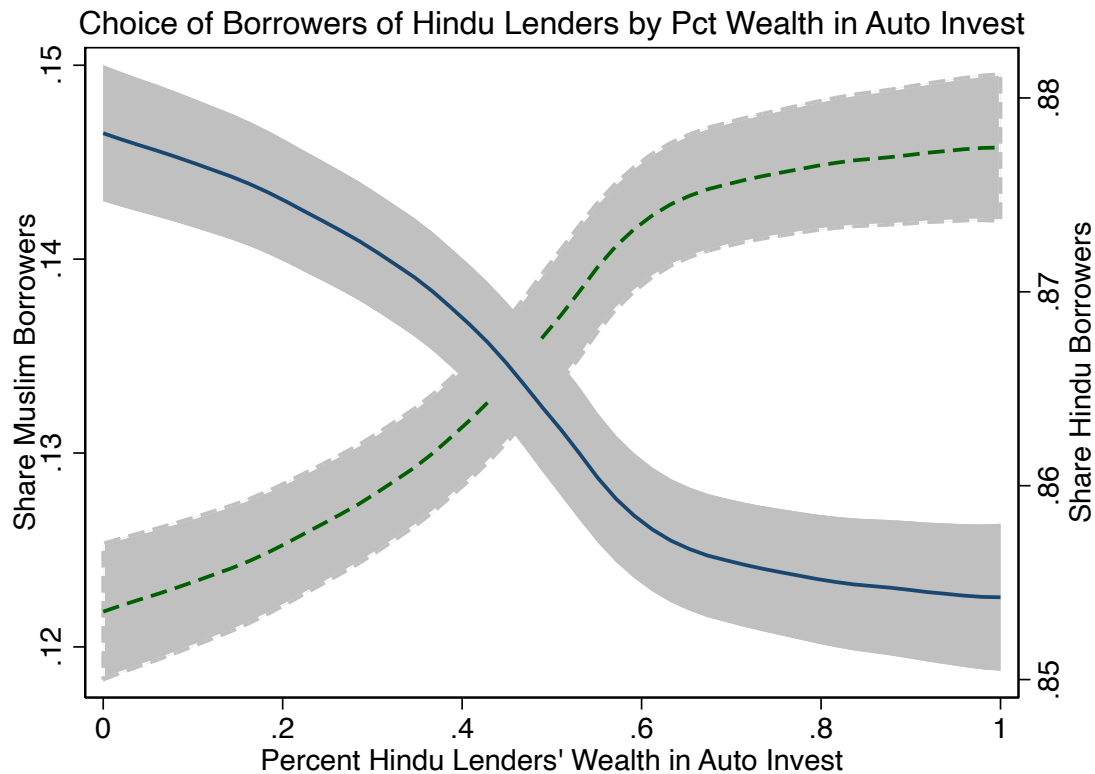


Figure 2 plots the coefficient estimates of kernel-weighted local mean smoothing regressions for whether borrowers are Hindu (blue, solid line, measured on the right y-axis) and whether borrowers are Muslim (green, dashed line, measured on the left y-axis) on the share of their available funds Hindu borrowers who moved to the robo-advising lending tool (Auto Invest) allocated to such tool. This share represents the intensive margin of usage of Auto Invest by Hindu borrowers. Grey bandwidths correspond to 95% confidence intervals around the estimated coefficients. We use an Epanechnikov kernel and evaluate the relation at 50 grid points.

Figure 3: Spatial Heterogeneity of In-group vs. Out-group Conflict

Panel A. Hindu-Muslim Riots, 1980-2000

Panel B. Average Vote Shares for the Bharatiya Janata Party (BJP) 1977-2015

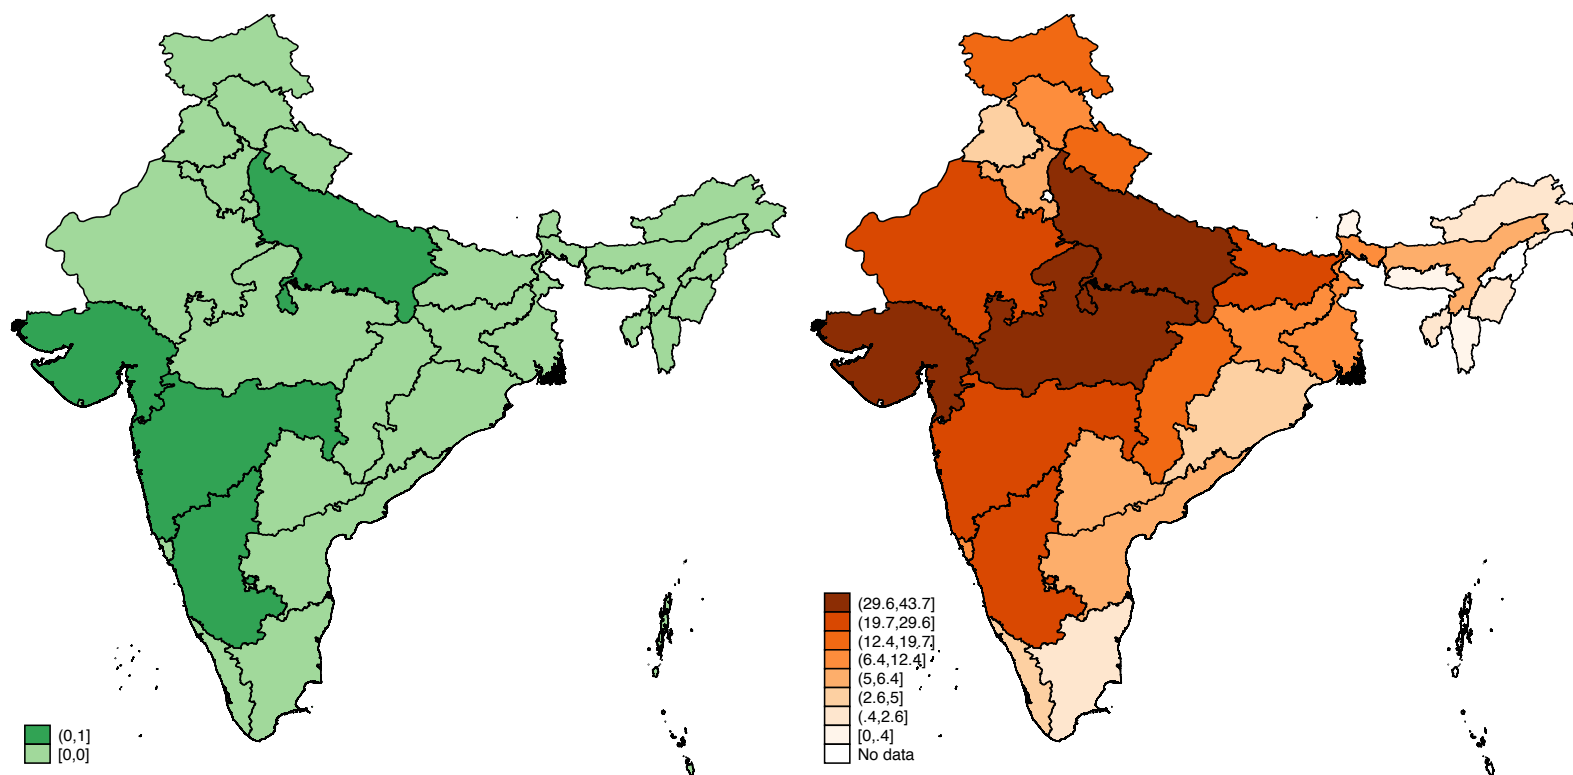


Figure 3 depicts the spatial variation of proxies for the vividness of Hindu-Muslim conflict across Indian states. Panel A compares states in which large-scale riots between Hindus and Muslims and/or pogroms against the Muslim minority happened between 1980 and 2000. Dark green states (Gujarat, Uttar Pradesh, Delhi, Maharashtra, and Karnataka) are states in which such events happened based on Ticku (2015). Panel B compares states based on the average vote share of BJP candidates to national and local elections between 1977 and 2015. We obtain candidate-level election results from 1977 to 2015 from Bhavnani (2014). We first compute the voting shares for each election in each state and then average these shares within states. The darker is a state, the higher is the average BJP vote share.

**Figure 4: Change in Lending to Out-group Borrowers:
Salience of Hindu-Muslim Conflict**





Figure 4 reports the results of estimating the following specification by ordinary least squares across different subsamples reported on top of each column:

$$\begin{aligned} Muslim\ Borrower_{i,j} = & \alpha + \beta\ Auto\ Invest_j + \gamma\ Hindu\ Lender_j \\ & + \delta\ Hindu\ Lender_j \times Auto\ Invest_j + \zeta\ \mathbf{x}_i + \epsilon_{i,j} \end{aligned}$$

where $Muslim\ Borrower_{i,j}$ is equal to 1 if the borrower i who receives funding from lender j is Muslim, and zero otherwise; $Auto\ Invest_j$ is equal to 1 if the lender made the loans after activating Auto Invest and 0 otherwise; $Hindu\ Lender_j$ is equal to 1 if lender j is Hindu; and \mathbf{x}_i is a vector of loan-level characteristics that are direct proxies for the risk profiles of the loans lenders extend to borrowers—loan maturity (measured in months), loan amount, and the annual interest rate associated with the loan. These loan-level characteristics are assigned to borrowers by the platform's algorithm when the loan requests are vetted before borrowers access the borrower pool. We cluster standard errors at the lender level. Panel A plots estimated coefficient $\hat{\gamma}$, which captures the extent of lender bias before accessing Auto Invest. Panel B plots estimated coefficient $\hat{\delta}$, which captures the de-biasing effect of Auto Invest.

**Figure 5: Lending to Discriminated Borrowers:
Shudra Caste Borrowers Before and After Debiasing**

Panel A. All Lenders



Panel B. Only Shudra Lenders (Easier to Recognize Own Caste)





Figure 5 plots the average share of borrowers in Hindu lenders' portfolios who are Shudra before lenders moved to the robo-advising tool (Auto Invest, black bars) and after lending decisions are made by Auto Invest (red bars). Panel A consider all Hindu lenders on the platform whereas Panel B only includes Shudra Hindu lenders, for whom recognizing the caste of Shudra borrowers might be weakly easier. In each panel, the left graph considers all borrowers in lenders' portfolios; the middle graph only considers borrowers whose caste can be recognized by a human with a probability above 50% as defined by the algorithm designed by Bhagavatula et al. (2018); the right graph only considers borrowers whose caste can be recognized with a probability above 70% based on the same algorithm.

Figure 6: Change in Lending to Discriminated Borrowers—Shudra Caste Borrowers

Panel A. De-Biasing After AutoInvest by Recognizability of Shudra Borrowers



Panel B. Heterogeneous De-biasing After AutoInvest by Salience of Shudra Discrimination




Figure 6 reports the results of estimating the following specification by ordinary least squares:

$$Shudra\ Borrower_{i,j} = \alpha + \beta\ Auto\ Invest_j + \zeta\ \mathbf{x}_i + \epsilon_{i,j}$$

where $Shudra\ Borrower_{i,j}$ is equal to 1 if the borrower i who receives funding from lender j is a Shudra, and zero otherwise; $Auto\ Invest_j$ is equal to 1 if the lender made the loans after activating Auto Invest and 0 otherwise; and \mathbf{x}_i is a vector of loan-level characteristics that are direct proxies for the risk profiles of the loans lenders extend to borrowers—loan maturity (measured in months), loan amount, and the annual interest rate associated with the loan. These loan-level characteristics are assigned to borrowers by the platform’s algorithm when the loan requests are vetted before borrowers access the borrower pool. Panel A reports the $\hat{\beta}$ coefficients across Shudra borrowers with different recognizability. Panel B further divides the estimates across states with high and low crimes against Shudra. We cluster standard errors at the lender level.

Figure 7: Spatial Heterogeneity of Salience of Stereotypical Discrimination

Crimes Against Scheduled Castes per inhabitant (2018)

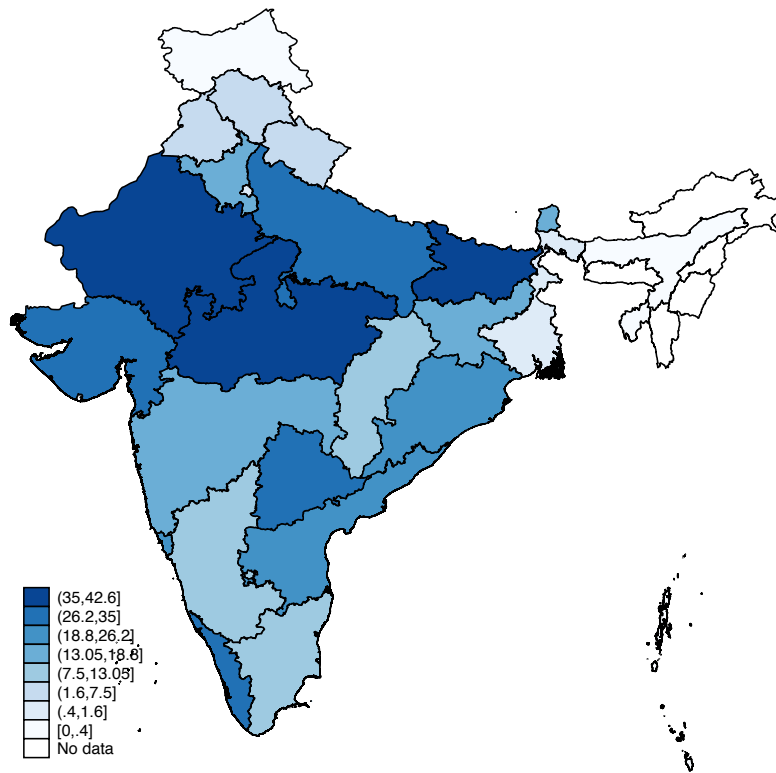


Figure 7 depicts the spatial variation of a proxy for the salience of discrimination against lower castes by Hindus across Indian states, that is, the number of crimes against Scheduled Castes (which includes members of the Shudra varna as well as those belonging to lower castes) per 100,000 inhabitants in 2018 based on the official data from the Indian National Crime Records Bureau (NCRB (2019)). The darker is a state, the higher is the number of crimes against Schedules Classes per inhabitant in the state.

Figure 8: Lower Default of Discriminated Borrowers Relative to Others Before Auto Invest

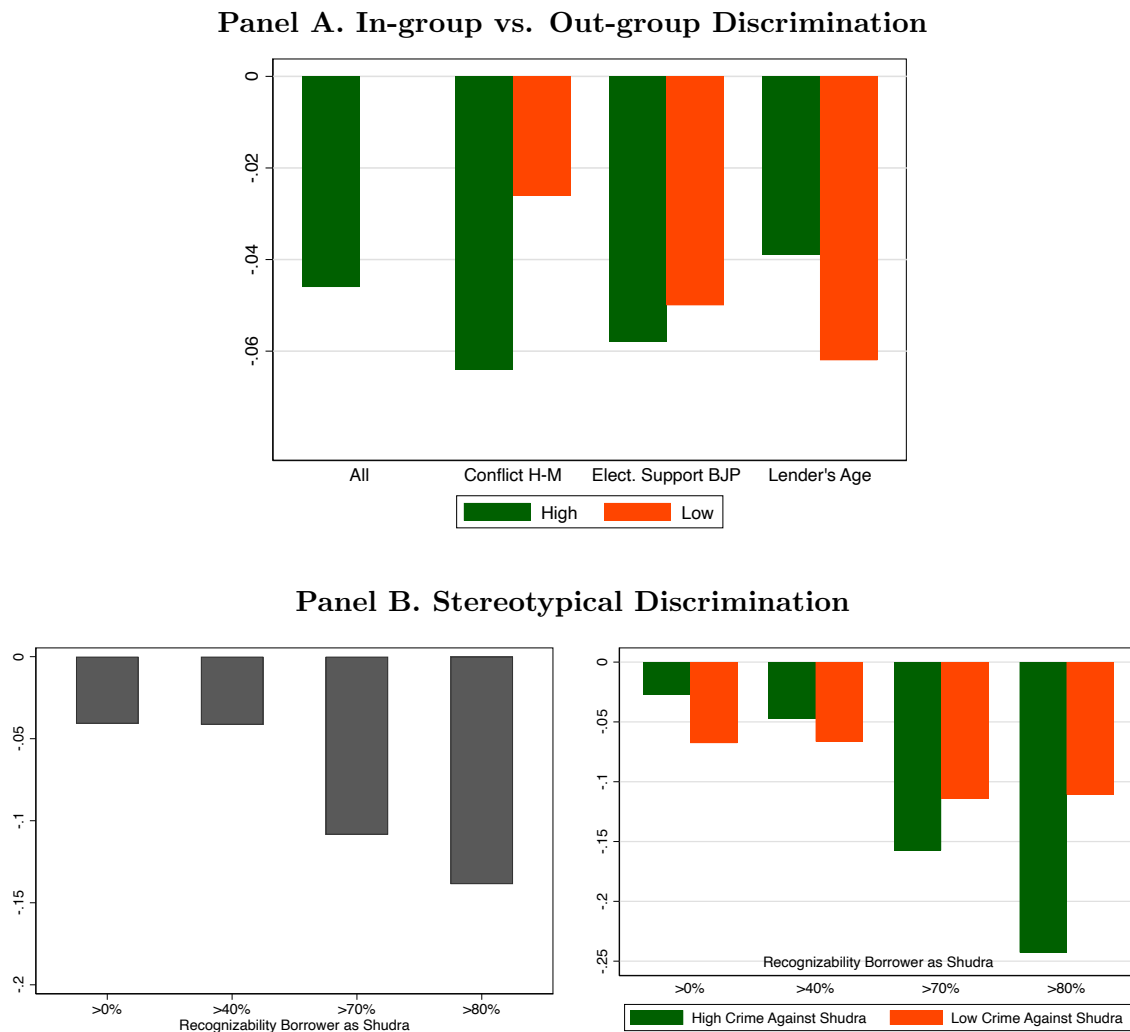



Figure 8 plots the relative default of Muslim borrowers relative to Hindu borrowers in Hindu lenders' portfolios (Panel A) and of Shudra borrowers relative to other Hindu borrowers in all Hindu lenders' portfolios before lenders moved to the robo-advising tool (Auto Invest) and across different subsamples. Panel A includes borrowers in Hindu lenders' portfolios. In Panel B, the probability of caste recognition of Shudra borrowers is based on the algorithm developed by Bhagavatula et al. (2018).

Figure 9: Fraction of Loan Repaid Before and After Debiasing

Panel A. In-group vs. Out-group Discrimination



Panel B. Stereotypical Discrimination


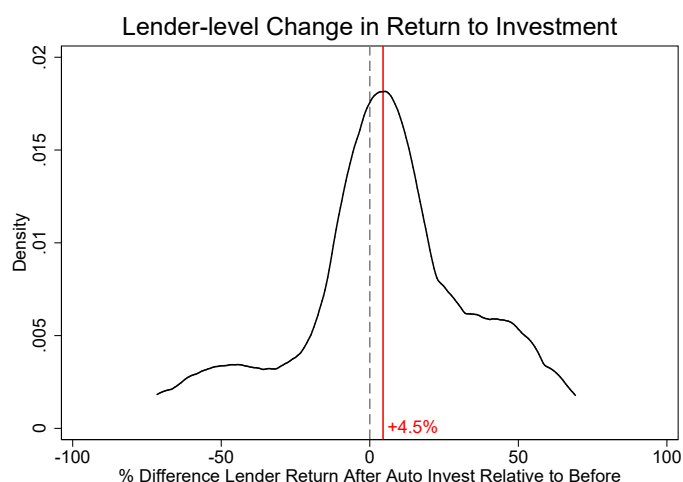


Figure 9 plots a set of cumulative distribution functions (CDFs) for different groups of borrowers over the share of the overall amounts due (principal plus interest) that borrowers repaid by the time their loan account was closed. In all Panels, the left graph refers to the CDFs of borrowers in lenders' portfolios before moving to the robo-advising tool (Auto Invest), whereas the right graph refers to the CDFs after moving to Auto Invest. Panel A includes borrowers in Hindu lenders' portfolios. The green solid lines are the CDFs for Hindu borrowers and the orange dashed lines for Muslim borrowers. Panel B only includes Shudra borrowers in Hindu lenders' portfolios. The green solid lines are the CDFs for Shudra borrowers whose probability of caste recognition is below 15% based on the algorithm developed by Bhagavatula et al. (2018), and the orange dashed lines for other Shudra borrowers.

Figure 10: Lender-level Change in Returns After Debiasing

Panel A. In-group vs. Out-group Discrimination



Panel B. Stereotypical Discrimination




Figure 10 plots the density of the lender-level average change in the returns to loans originated on the platform after using Auto Invest relative to before. To compute the change, we first compute the overall return to the amounts invested by each lender separately before and after using Auto Invest across all loans they originated. We then subtract the two returns and average these changes across lenders. Loan-level returns are thus value weighted within lenders and lender-level returns are equally weighted across lenders. Panel A refers to Hindu lenders' loans to Hindu borrowers (in-group vs. out-group discrimination), whereas Panel B refers to all Hindu lenders' loans to Shudra borrowers (stereotypical discrimination).

Table 1. Summary Statistics

Panel A. In-group vs. Out-group Discrimination Sample						
	<u>N. obs.</u>	<u>Mean</u>	<u>St. dev.</u>	<u>25th perc.</u>	<u>Median</u>	<u>75th perc.</u>
Muslim Borrower	113,283	0.13	0.34	0.00	0.00	0.00
Hindu Lender	113,283	0.99	0.11	1.00	1.00	1.00
Auto Invest	113,283	0.45	0.50	0.00	0.00	1.00
Auto Invest allocation (%)	113,283	0.59	0.37	0.22	0.60	1.00
Tenure (months)	113,283	22.08	8.98	15.00	24.00	24.00
Loan Amount (rupees)	113,283	131,074	102,575	50,000	100,000	188,000
Interest Rate	113,283	0.24	0.07	0.20	0.23	0.27
Delinquent	113,127	0.29	0.45	0.00	0.00	1.00

Panel B. Stereotypical Discrimination Sample						
	<u>N. obs.</u>	<u>Mean</u>	<u>St. dev.</u>	<u>25th perc.</u>	<u>Median</u>	<u>75th perc.</u>
Shudra Borrower	62,831	0.39	0.49	0.00	0.00	1.00
Auto Invest	62,831	0.43	0.49	0.00	0.00	1.00
Auto Invest allocation (%)	62,831	0.58	0.37	0.22	0.57	1.00
Tenure (months)	62,831	22.03	9.02	12.00	24.00	30.00
Loan Amount (rupees)	62,831	131,797	105,994	50,000	100,000	200,000
Interest Rate	62,831	0.24	0.07	0.20	0.23	0.26
Delinquent	62,736	0.28	0.45	0.00	0.00	1.00

Table 1 reports summary statistics for the main variables in the analysis across the two datasets used in the analysis of in-group vs. out-group discrimination (Panel A) and stereotypical discrimination (Panel B). In both panels, the unit of observation is a lender-borrower-loan triad. Borrower-lender characteristics include the religion/caste of borrowers and lenders. *Auto Invest* is a dummy variable that equals 1 if the lender uses the robo-advising lending tool, whereas *Auto Invest allocation* is the share of funds lenders have available on the P2P platform that they allocate to the robo-advising tool. Loan-level characteristics include the loans' tenure, size, and annual interest rate, as well as a dummy variable that equals 1 if the loan was delinquent at the time it was closed and zero otherwise.

**Table 2. Change in Lending to Out-group Borrowers:
Hindu vs. Muslim**

<i>Dependent variable:</i> Muslim Borrower					Low Use Auto Invest	High Use Auto Invest
	(1)	(2)	(3)	(4)	(5)	(6)
Hindu Lender \times Auto Invest	0.045** (2.51)	0.046** (2.51)	0.045** (2.07)	0.043** (1.96)	0.009 (0.23)	0.052* (1.94)
Hindu Lender	-0.058*** (-3.52)	-0.058*** (-3.54)				
Auto Invest	-0.026 (-1.45)	-0.025 (-1.40)	-0.030 (-1.43)	-0.033 (-1.53)	0.011 (0.29)	-0.048* (-1.83)
Maturity		0.012*** (4.78)	0.009*** (3.20)	0.010*** (3.35)	0.010** (2.31)	0.009** (2.31)
Loan Amount (₹000)		-0.001*** (-10.34)	-0.001*** (-10.81)	-0.001*** (-10.92)	-0.001*** (-5.47)	-0.001*** (-9.27)
Interest Rate		-0.009 (-0.47)	-0.015 (-0.74)	-0.012 (-0.62)	0.013 (0.47)	-0.035 (-1.25)
Constant	0.181*** (11.14)	0.163*** (8.74)	0.119*** (11.72)	0.120*** (11.77)	0.099*** (6.77)	0.137*** (9.60)
Lender Fixed Effects			✓	✓	✓	✓
Year Fixed Effects				✓	✓	✓
N. obs.	113,284	113,283	113,283	113,283	39,366	72,104

Table 2 reports the results of estimating the following specification by ordinary least squares:

$$Muslim\ Borrower_{i,j} = \alpha + \beta\ Auto\ Invest_j + \gamma\ Hindu\ Lender_j + \delta\ Hindu\ Lender_j \times Auto\ Invest_j + \zeta\ \mathbf{x}_i + \epsilon_{i,j}$$

where $Muslim\ Borrower_{i,j}$ is equal to 1 if the borrower i who receives funding from lender j is Muslim, and zero otherwise; $Auto\ Invest_j$ is equal to 1 if the lender made the loans after activating Auto Invest and 0 otherwise; $Hindu\ Lender_j$ is equal to 1 if lender j is Hindu; and \mathbf{x}_i is a vector of loan-level characteristics that are direct proxies for the risk profiles of the loans lenders extend to borrowers—loan maturity (measured in months), loan amount, and the annual interest rate associated with the loan. These loan-level characteristics are assigned to borrowers by the platform's algorithm when the loan requests are vetted before borrowers access the borrower pool. We cluster standard errors at the lender level.

Table 3. Loan Defaults Before and After Debiasing

<i>Dependent variable:</i>	In-group vs. Out-group Discrimination			Stereotypical Discrimination		
Delinquent Loan		Hindu Lenders			All Lenders	
	(1)	(2)	(3)	(4)	(5)	(6)
Auto Invest	-0.108*** (-5.08)			-0.157*** (-6.25)		
Hindu Borrower × Auto Invest		-0.112*** (-5.21)	-0.072*** (-4.55)			
Muslim Borrower × Auto Invest		-0.073** (-2.49)	-0.070*** (-2.82)			
Muslim Borrower	-0.024** (-2.02)	-0.046** (-2.58)	-0.002 (-0.10)			
Non-Shudra Borrower × Auto Invest					-0.160*** (-6.04)	-0.100*** (-4.39)
Shudra Borrower × Auto Invest					-0.148*** (-4.54)	-0.104*** (-3.32)
Shudra Borrower				-0.038*** (-3.03)	-0.041*** (-2.89)	-0.030** (-2.14)
Constant	0.499*** (30.95)	0.501*** (30.84)	0.403*** (17.69)	0.490*** (32.13)	0.491*** (32.25)	0.421*** (7.77)
Loan Risk Characteristics			✓			✓
Lender Fixed Effects			✓			✓
Time Fixed Effects			✓			✓
N. obs.	16,985	16,985	16,985	6,821	6,821	6,821
R-Square	0.012	0.012	0.263	0.020	0.020	0.357

Table 3 reports the results of estimating variations of the following specification by ordinary least squares:

$$\begin{aligned}
 \text{Delinquent Loan}_{ij} = & \alpha + \gamma \text{Muslim Borrower}_j \\
 & + \delta \text{Muslim Borrower}_j \times \text{Auto Invest}_j \\
 & + \theta \text{Hindu Borrower}_j \times \text{Auto Invest}_j + \zeta \mathbf{x}_i + \epsilon_{i,j},
 \end{aligned}$$

where $\text{Delinquent Loan}_{ij}$ is equal to 1 if the loan associated with borrower i and lender j is closed as delinquent; Auto Invest_j is equal to 1 if the lender made the loans after activating Auto Invest and 0 otherwise; $\text{Muslim Borrower}_{i,j}$ ($\text{Hindu Borrower}_{i,j}$) is equal to 1 if the borrower i who receives funding from lender j is Muslim (Hindu), and zero otherwise; and \mathbf{x}_i is a vector of loan-level risk characteristics—loan maturity (measured in months), loan amount, and the annual interest rate associated with the loan. These loan-level characteristics are assigned to borrowers by the platform’s algorithm when the loan requests are vetted before borrowers access the borrower pool. We cluster standard errors at the lender level.

Table 4. Loan Returns Before and After Debiasing

<i>Dependent variable:</i>	In-group vs. Out-group Discrimination			Stereotypical Discrimination		
Loan Return		Hindu Lenders			All Lenders	
	(1)	(2)	(3)	(4)	(5)	(6)
Auto Invest	0.204* (1.84)			0.297** (2.14)		
Hindu Borrower × Auto Invest		0.217* (1.83)	0.097 (0.83)			
Muslim Borrower × Auto Invest		0.065 (-1.57)	-0.007 (-0.05)			
Muslim Borrower	0.516*** (11.58)	0.615** (7.59)	0.305*** (4.50)			
Non-Shudra Borrower × Auto Invest					0.272* (1.73)	0.270* (1.73)
Shudra Borrower × Auto Invest					0.359** (2.55)	0.262* (1.71)
Shudra Borrower				0.223*** (3.34)	0.170 (1.49)	-0.231** (2.45)
Constant	-0.173* (-1.77)	-0.181* (-1.77)	-0.485*** (-4.01)	-0.411 (-3.33)	-0.396*** (-2.98)	-1.436*** (-7.52)
Loan Risk Characteristics			✓			✓
Lender Fixed Effects			✓			✓
Time Fixed Effects			✓			✓
N. obs.	2,546	2,546	2,546	1,128	1,128	1,128
R-Square	0.031	0.031	0.553	0.026	0.026	0.632

Table 4 reports the results of estimating variations of the following specification by ordinary least squares:

$$\begin{aligned}
 \text{Loan Return}_{ij} = & \alpha + \gamma \text{Muslim Borrower}_j \\
 & + \delta \text{Muslim Borrower}_j \times \text{Auto Invest}_j \\
 & + \theta \text{Hindu Borrower}_j \times \text{Auto Invest}_j + \zeta \mathbf{x}_i + \epsilon_{i,j},
 \end{aligned}$$

where Loan Return_{ij} is the standardized return of the loan associated with borrower i and lender j at loan closure; Auto Invest_j is equal to 1 if the lender made the loans after activating Auto Invest and 0 otherwise; $\text{Muslim Borrower}_{i,j}$ ($\text{Hindu Borrower}_{i,j}$) is equal to 1 if the borrower i who receives funding from lender j is Muslim (Hindu), and zero otherwise; and \mathbf{x}_i is a vector of loan-level risk characteristics—loan maturity (measured in months), loan amount, and the annual interest rate associated with the loan. These loan-level characteristics are assigned to borrowers by the platform’s algorithm when the loan requests are vetted before borrowers access the borrower pool. We cluster standard errors at the lender level.

**Table 5. Loan Returns Before and After Debiasing:
Quantile Regressions**

Panel A. In-group vs. Out-group Discrimination								
<i>Dependent variable:</i>	OLS		25 th percentile		Median		75 th percentile	
Loan's Return	Hindu (1)	Muslim (2)	Hindu (3)	Muslim (4)	Hindu (5)	Muslim (6)	Hindu (7)	Muslim (8)
Auto-Invest (Without Risk Controls)	0.217*** (4.92)	0.065** (2.25)	0.250* (1.82)	0.173** (2.23)	0.274*** (11.22)	0.121*** (3.84)	0.109*** (11.66)	-0.004 (-0.29)
Auto-Invest (With Risk Controls)	-0.063** (-2.20)	-0.006 (-0.30)	-0.008 (-1.01)	-0.002 (-0.38)	-0.035 (-0.04)	-0.000 (-0.04)	0.000 (0.02)	0.008 (1.04)
N. obs.	2,326	220	2,326	220	2,326	220	2,326	220

Panel B. Stereotypical Discrimination								
<i>Dependent variable:</i>	OLS		25 th percentile		Median		75 th percentile	
Loan's Return	Low Prob Shudra (1)	High Prob Shudra (2)	Low Prob Shudra (3)	High Prob Shudra (4)	Low Prob Shudra (5)	High Prob Shudra (6)	Low Prob Shudra (7)	High Prob Shudra (8)
Auto-Invest (Without Risk Controls)	0.079*** (2.89)	0.068*** (4.54)	0.485*** (4.61)	0.080* (1.73)	0.088*** (3.02)	0.072*** (8.13)	0.030*** (5.64)	0.025*** (9.29)
Auto-Invest (With Risk Controls)	-0.053 (-0.82)	-0.039 (-1.04)	-0.001 (-0.07)	0.000 (0.05)	-0.011 (-0.55)	0.014** (1.99)	0.020 (0.55)	0.020*** (4.83)
N. obs.	462	1,158	462	1,158	462	1,158	462	1,158

Table 5 reports the results of estimating the following set of quantile regressions:

$$Q_{\tau}(\text{Loan's Return}_{i,j}) = \alpha(\tau) + \beta(\tau) \text{Auto Invest}_j + X'_{i,j} \zeta(\tau) + \epsilon_{i,j},$$

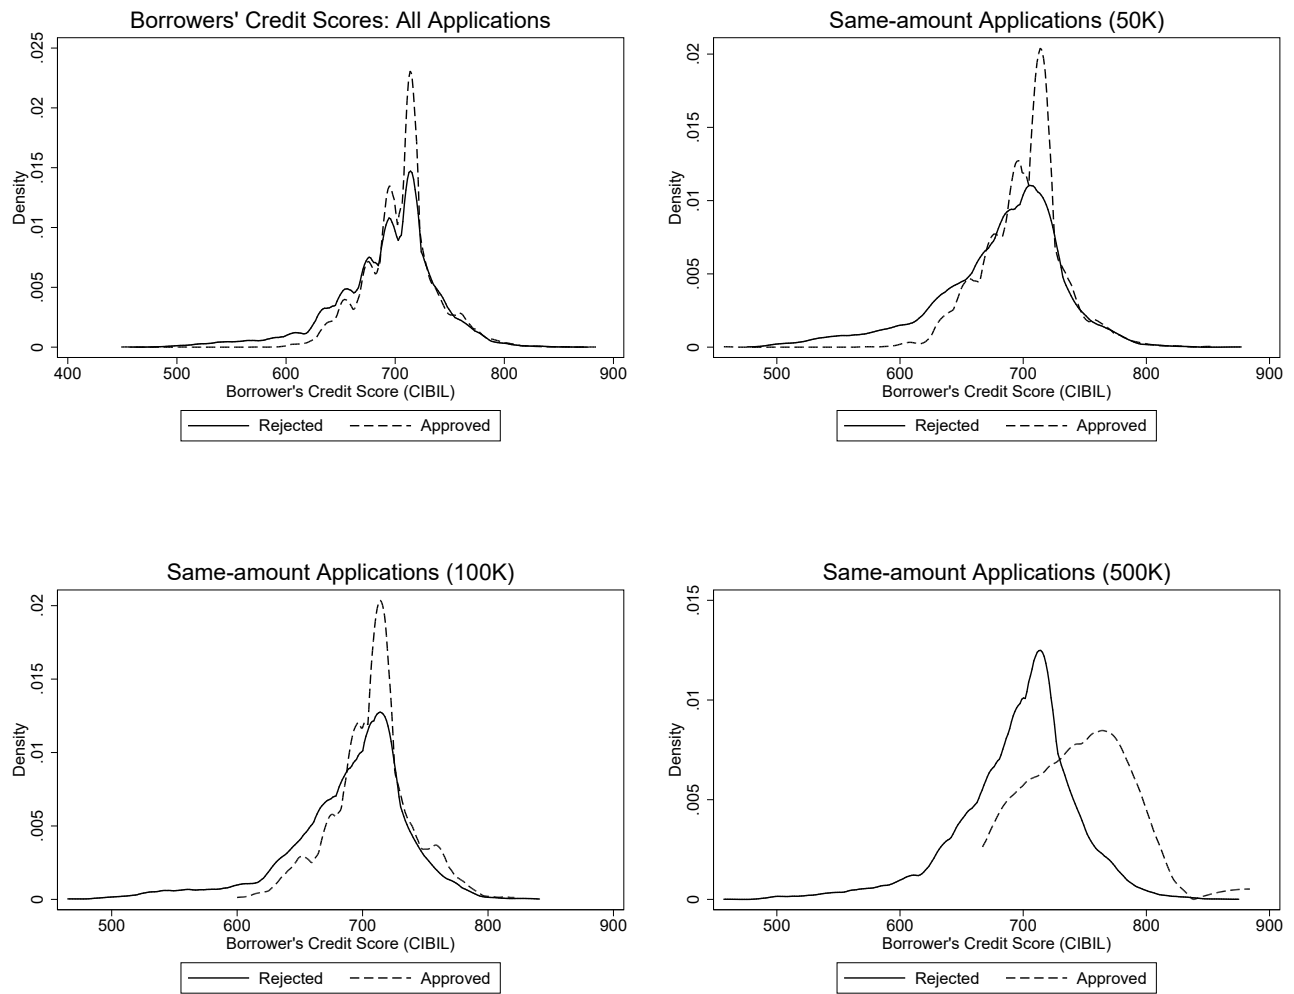
whose outcome variable is quantile Q_{τ} of the distribution of the (standardized) loan return associated with borrower i and lender j throughout the sample period; Auto Invest_j is equal to 1 if the lender made the loans after activating Auto Invest and 0 otherwise; and $X_{i,j}$ is a vector of loan-level characteristics that are direct proxies for the risk profiles of the loans lenders extend to borrowers—loan amount, and the annual interest rate associated with the loan. These loan-level characteristics are assigned to borrowers by the platform's algorithm when the loan requests are vetted before borrowers access the borrower pool. In each panel, the first row reports the estimates of $\beta(\tau)$ *without* controlling for loan's risk characteristics. The second row reports the estimates of the same specifications when risk controls are included. We cluster standard errors at the lender level.

Online Appendix:
How Costly Are Cultural Biases?
Evidence from FinTech

Francesco D'Acunto, Pulak Ghosh, Rajiv Jain, Alberto G. Rossi

Not for Publication

**Figure A.1: Platform's Screening of Borrowers 1:
Credit Scores of Approved and Rejected Borrowers**



**Figure A.2: Platform's Screening of Borrowers 2:
Interest Rates, Maturities, and Loan Amounts by Credit Score**

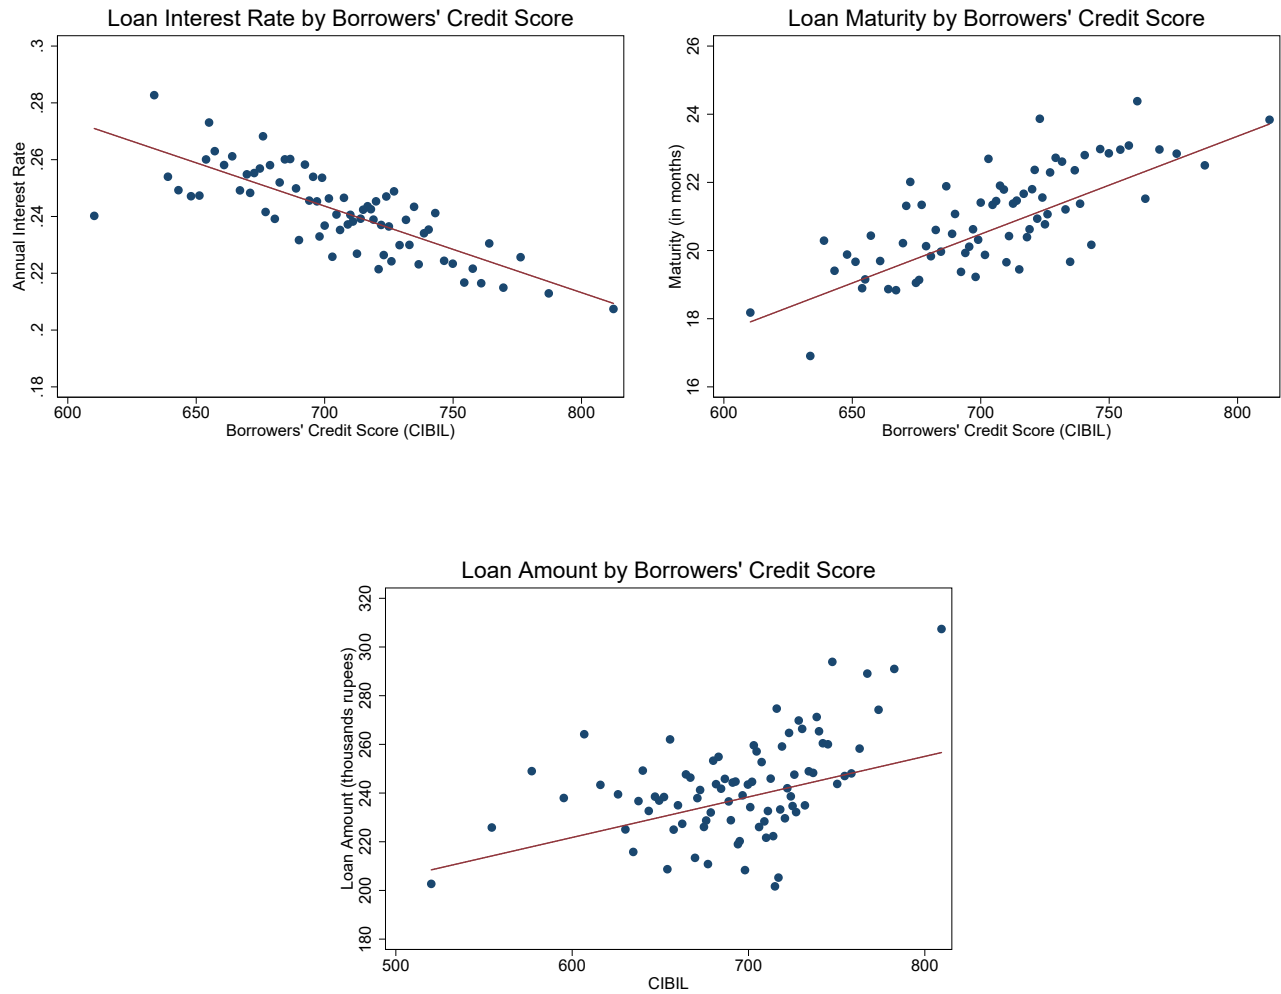


Figure A.3: Robo-Advising Tool—Auto Invest

My Auto Invest Allocation:


Setup your Auto Invest Allocation here

Total amount to allocate: ₹ 560,465.00

CATEGORIES	ALREADY DEPLOYED	MAX PROPOSAL AMOUNT (₹)	ALLOCATION (%)	ALLOCATION AMOUNT (₹)
High Range (>26%) Very High, Instant Min Proposal Amount: ₹ 500	₹ 8,500.00	500.00 ▼	20 ▼	112093
Mid Range (18% - 26%) Medium, High Min Proposal Amount: ₹ 1000	₹ 21,600.00	1,000.00 ▼	35 ▼	196162.75
Low Range (<18%) Prime, Minimal, Low Min Proposal Amount: ₹ 2000	₹ 38,235.00	2,000.00 ▼	45 ▼	252209.25

**Figure A.4: Geographic Distribution of Lending:
Number of Indian States in which Each Lender Disburses Funds**




Figure A.5: Ex-post Likelihood of Default and Interest Rates




Figure A.6: Distribution of Probabilities that Borrowers are Shudra




Figure A.7: Intensive Margin Performance Before and After Auto Invest: Full Sample




Table A.1. Determinants of Robo-Advice Adoption

<i>Dependent variable:</i> Ever Adopted Robo-Advising Tool		
	(1)	(2)
In-group Bias	-0.060 (-0.88)	0.007 (0.02)
Hindu Lender		-0.094 (-0.30)
In-group Bias × Hindu Lender		-0.076 (-0.24)
Constant	0.850*** (14.01)	0.950*** (3.13)
N. obs.	1,567	1,567
R-square	0.000	0.003

Table A.1 reports the results of estimating the following specification—estimated at the level of lender j —by ordinary least squares:

$$\begin{aligned} \text{Ever Adopted Robo}_j = & \alpha + \beta \text{In-group Bias}_j + \gamma \text{Hindu Lender}_j \\ & + \delta \text{In-group Bias}_j \times \text{Hindu Lender}_j + \epsilon_j \end{aligned}$$

This specification includes the full set of Faircent lenders rather than only the lenders who adopted the robo-advising tool at some point during our sample period, which constitutes the population of all our other empirical analysis. *Ever Adopted Robo_j* is a dummy variable that equals 1 if lender j has ever adopted the robo-advising tool at any point during our sample period, and zero otherwise; *In-group Bias_j* is the difference between the share of out-group borrowers in the overall pool and the share of that group's population in lender j 's loan portfolio. For instance, for Hindu lenders, *In-group Bias_j* = *Share Muslim Borrower Pool* – *Share Loans to Muslim Borrowers*.

**Table A.2. Change in Lending to Out-group Borrowers:
Robustness, Statistical Inference**

<i>Dependent variable:</i> Muslim Borrower					Low Use Auto Invest	High Use Auto Invest
	(1)	(2)	(3)	(4)	(5)	(6)
Hindu Lender \times Auto Invest	0.045	0.046	0.045	0.043	0.009	0.052
<i>By Lender</i>	(2.51)**	(2.51)**	(2.07)**	(1.96)**	(0.23)	(1.94)*
<i>By Lender and Borrower</i>	(2.24)**	(2.24)**	(2.00)**	(1.90)*	(0.23)	(1.93)*
<i>By Lender, Borrower, and Month</i>	(2.13)**	(2.14)**	(2.51)**	(2.19)**	(0.21)	(2.71)***
<i>By Lender Surname, Borrower Surname, and Month</i>	(2.16)**	(2.18)**	(2.01)**	(2.03)**	(0.22)	(1.94)*
Lender Fixed Effects			✓	✓	✓	✓
Year Fixed Effects				✓	✓	✓
N. obs.	113,284	113,283	113,283	113,283	39,366	72,104

Table A.2 reports the results of estimating the following specification by ordinary least squares:

$$\begin{aligned} \text{Muslim Borrower}_{i,j} = & \alpha + \beta \text{Auto Invest}_j + \gamma \text{Hindu Lender}_j \\ & + \delta \text{Hindu Lender}_j \times \text{Auto Invest}_j + \zeta \mathbf{x}_i + \epsilon_{i,j} \end{aligned}$$

where $\text{Muslim Borrower}_{i,j}$ is equal to 1 if the borrower i who receives funding from lender j is Muslim, and zero otherwise; Auto Invest_j is equal to 1 if the lender made the loans after activating Auto Invest and 0 otherwise; Hindu Lender_j is equal to 1 if lender j is Hindu; and \mathbf{x}_i is a vector of loan-level characteristics that are direct proxies for the risk profiles of the loans lenders extend to borrowers—loan maturity (measured in months), loan amount, and the annual interest rate associated with the loan. These loan-level characteristics are assigned to borrowers by the platform's algorithm when the loan requests are vetted before borrowers access the borrower pool. In each line, we report the t-statistics estimated with the indicated level of clustering of standard errors.