How Costly Are Cultural Biases?

Francesco D'Acunto[*]
Boston College

Pulak Ghosh[†]
IIMB

Rajiv Jain[‡]
Faircent

Alberto G. Rossi[§]
Georgetown University

This version: January 2021

Abstract

To detect and quantify the effects of cultural biases in large-stake risky choices, we study a leading Indian FinTech peer-to-peer lending platform paired with an automated robo-advising tool. Comparing the choices lending consumers ("lenders") make with those made by the robo-advising tool on their behalf, we find both in-group vs. out-group discrimination and stereotypical discrimination are pervasive and sizable. Discrimination affects performance negatively: discriminating lenders face 32% higher default rates and about 11% lower returns on the loans they issue to borrowers who belong to favored demographic groups relative to available borrowers in discriminated groups. The extent of discrimination is higher in locations in which cultural biases are salient, due to historical inter-ethnic conflict and political polarization.

Keywords: Taste-based Discrimination, Statistical Discrimination, Cultural Finance, Robo-Advising, Lending, Disintermediation, Partisanship, Inter-ethnic Conflict, Consumer Finance.

[*]Carroll School of Management, Boston College, Chestnut Hill, MA, USA. e-Mail: dacuntof@bc.edu
[†]Indian Institute of Management, Bannerghatta Road- 560076, Bangalore, India. e-Mail: pulak.ghosh@iimb.ac.in
[‡]Fairassets Technologies India Private Ltd, Gurugram, Haryana, India.
[§]McDonough School of Business, Georgetown University, Washington, DC, USA. Fellow of the Luohan Academy. e-Mail: agr60@georgetown.edu. For very helpful comments, we thank Vicki Bogan, Paul Goldsmith-Pinkham, Apoorv Gupta, Filippo Mezzanotti, Anirban Mitra, Adair Morse, Chris Parsons, Antoinette Schoar, Kelly Shue, Huan Tang, David Thesmar, and Ansgar Walther, as well as seminar participants at the 2020 Bartunek Research Forum. All errors are our own.

1 Introduction

The cultural norms to which agents are exposed, especially when deep-rooted and based on centuries-long societal customs, can have long-lived influence on agents' beliefs and economic decision-making (Guiso, Sapienza, and Zingales (2006), Alesina, Giuliano, and Nunn (2013), D'Acunto (2019a); D'Acunto et al. (2019, 2020). Norms can also produce cultural biases—they can shape agents' beliefs about economic signals in a way that makes their choices deviate from those of a neoclassical economic agent (Guiso, Sapienza, and Zingales (2009), D'Acunto (2019b)).[1] Examples of cultural biases include taste-based discrimination against members of different social groups (*in-group vs. out-group discrimination*; see, e.g., Tajfel et al. (1979)), discrimination against members of one's own group due to cultural stereotypes (*stereotypical discrimination*, see e.g. (Becker (1957),Akerlof and Kranton (2000), Parsons et al. (2011)), and inaccurate statistical discrimination based on biased beliefs grounded in cultural norms when decision-makers lack information (Bohren et al. (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)). Moreover, homophily might help discriminators screen and/or monitor the members of their own social group more successfully than others, which could make 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 that disentangles these alternative channels.

In this paper, we exploit a FinTech peer-to-peer (P2P) lending setting paired with an automated robo-advising tool to ask whether cultural biases affect decision-making under risk and to estimate if and by how much cultural biases affect discriminators' performance. We find that the extent of cultural-bias-induced discrimination is sizable and reduces discriminators' investment performance: discriminating lenders face 32% higher default rates and about 11% lower returns on the loans they issue to borrowers who belong to favored demographics relative to available borrowers in the discriminated groups.

Our lab is a leading P2P platform in India, *Faircent.* As of October 2020, the platform hosted about 1.5 million borrowers and 140,000 lenders, both of which categories reside across all Indian states. The Indian setting allows us to study two forms of discrimination. We start with in-group vs. out-group discrimination, whereby the members of two conflicting social groups tend to favor members of their

[1] These 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.

own group (in-group) and disfavor members of the conflicting group (out-group). In the Indian context, as we discuss in detail below, this form of bilateral discrimination has been detected between Hindus and Muslims. We then study stereotypical discrimination, whereby everybody discriminates against one social group so much so that even the members of that group discriminate among each other. We discuss how, in the Indian context, this form of discrimination can arise due to the centuries-old caste-based system.

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

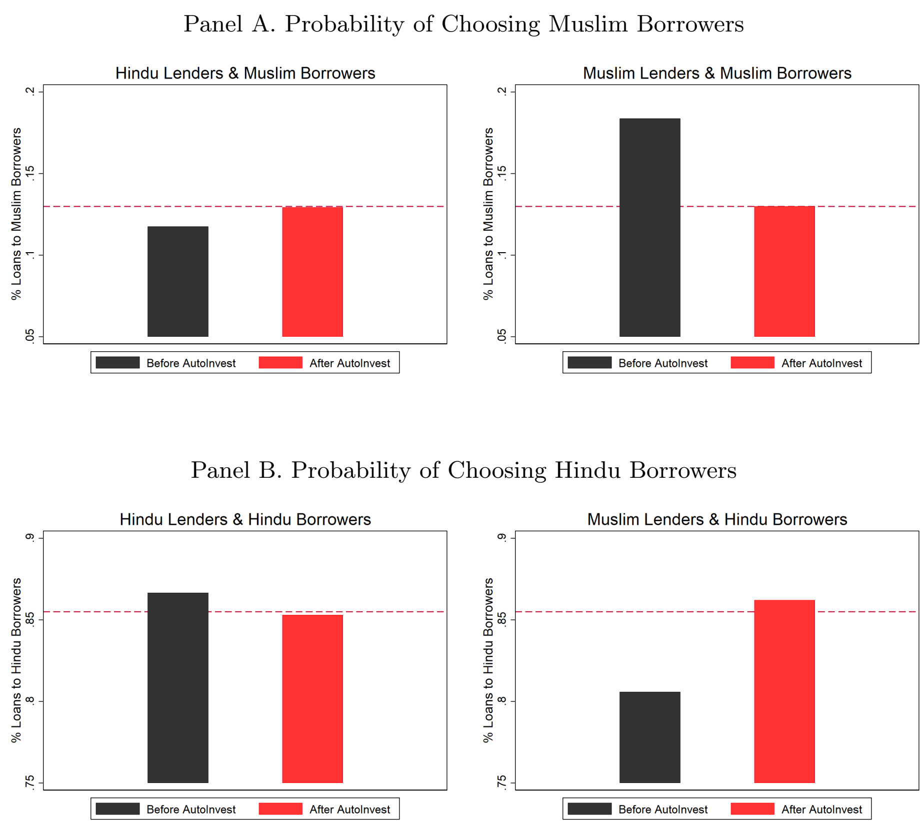

Figure 1 previews our baseline results in the raw data for the case of in-group vs. out-group discrimination. We compare the average share of Hindu and Muslim borrowers who were financed by Hindu and Muslim lenders before lenders adopted the automated robo-advising tool (black bars) and after lending choices were delegated to the tool (red bars). The raw data reveal three noticeable facts. First, consistent with the presence of an in-group vs. out-group bias for at least one of the two groups, both Hindu and Muslim lenders tend to favor borrowers of their same religion. Second, once the robo-advising tool makes choices, the shares of Muslim and Hindu borrowers change for both groups of lenders in opposite directions, which suggests all lenders discriminated against out-group borrowers when making choices autonomously. Third, the share of borrowers of different religions are virtually identical for Hindu and Muslim lenders after adoption of robo-advising, which suggests robo-advising appears to fully debias lenders by equating the share of borrowers of each religion in lenders' portfolios to the share of borrowers of each religion in the pool of available borrowers on the platform.[2]

A set of features make our field setting a viable laboratory in which to tackle the questions at stake. First, we observe many risky choices in a real-world high-stakes setting in which individual agents (henceforth, "lenders"), as opposed to bank loan officers, invest their own capital.[3] Because our lenders allocate their own capital across the borrowers they select on the platform, our setting includes no scope for misaligned incentives between the capital owner and the lending officer. Second, the platform, rather than lenders, screens prospective borrowers based on their risk profile *before* borrowers are visible to lenders. Decoupling the screening decision from the lending decision minimizes the scope for statistical discrimination by lenders, because lenders know the platform screens and assigns scores to the riskiness of borrower profiles. Third, within our sample period, the platform introduced an automated robo-advising tool that makes decisions on behalf of lenders who adopt it (D'Acunto et al. (2019), Rossi and Utkus (2020), D'Acunto and Rossi (2020)). When making investment decisions, the robo-advising tool does not consider borrowers' social groups (Morse and Pence (2020)). For this reason, we can compare the portfolios of borrowers that lenders picked before the adoption of robo-advising—which might have been shaped by cultural biases—and the portfolios after adoption, in which discrimination, if any, would be statistical in nature. Finally, because the platform provides loan servicing and the monitoring of borrowers (Iyer et al. (2016); Tang (2019); Balyuk (2019)), we can assess the sign and size of the economic effects of cultural biases by comparing the performance of loans issued by the same lenders based on their own choices or based on the robo-advising tool's choices.

We find the raw-data patterns in Figure 1 are robust features of the data: they do not change in multivariate analyses that control for the loan-level measures of risk we observe (interest rate, maturity, and loan amount) or when we restrict the variation within lenders, and hence absorb unobserved systematic time-invariant differences across lenders such as education levels or cultural background.

Moreover, we perform several heterogeneity tests to further assess the role of discrimination. For instance, we find that the extent of debiasing (change in shares of in-group borrowers after accessing robo-advising) is a monotonic function of the share of funds deposited in the platform each lender allocates to the robo-advising tool: the higher the share, the lower the favoritism toward in-group borrowers.[4] We also show that the extent of debiasing is economically and statistically larger in

[2]As we discuss below, the robo-advising tool as well as the selection of the borrower pool that accesses the platform might be subject to statistical discrimination. Fully debiasing in our context means that the two forms of discrimination we study do not appear to have a role in lending choices after adoption of the robo-advising tool.

[3]The P2P platform we study only includes individual investors and no institutional investors.

[4]As we discuss below, once they adopt the tool, lenders can decide the share of funds they want to allocate to it. In this way, we can study the intensive margin of debiasing by comparing choices lenders make at the same point in time, some of which are autonomous and some of which are automated.

settings in which the Hindu-Muslim conflict is salient—in cities with a higher presence of Hindu-Muslim riots, in states in which parties that foment this conflict obtain higher vote shares, and for lenders who have been more exposed to Hindu-Muslim riots during their formative years.

After detecting the presence of in-group vs. out-group discrimination, we move on to study how discrimination relates to loan-repayment behavior in our setting. If lenders discriminate in favor of their in-group and against their out-group, and this discrimination is not statistical, lenders should apply different standards when selecting borrowers: lenders should dig deeper into the pool of in-group borrowers—they should fund loans to lower-quality borrowers from their in-group rather than from their out-group (Agarwal et al. (2017)). Consistent with this intuition, we find that the loans lenders grant to in-group borrowers before adopting the robo-advising tool perform systematically worse than the loans they grant to out-group borrowers. In-group-borrower loans are about 9 percentage points (pp) more likely to default (32% of the average default rate). Also, among borrowers who pay at least a portion of the overall amount due, in-group borrowers are less likely to repay in full. The share of the overall amount paid is systematically lower for in-group borrowers at any quantile of the distribution. On average, lenders lose about 130,000 rupees ($1,770) out of an average investment of 1.2 million rupees due to in-group vs. out-group discrimination.

Note that in our setting, several channels that could predict the opposite direction of the association between discrimination and loan repayment are shut down by construction, because in our setting, no scope arises for the role of homophily in screening and monitoring borrowers (Schoar (2012); Drexler and Schoar (2014); Fisman et al. (2017); Fisman et al. (2020)), 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 between lenders and borrowers (Rao (2019)), preferences of physical appearance (Duarte, Siegel, and Young (2012); Ravina (2019)), or systematic ethnic differences in housing collateral value, given that loans are not backed by collateral (Avenancio-León and Howard (2019); Naaraayanan (2019)). Moreover, in our setting, the lender allocates her own financial capital to loans, and hence no scope exists for diverging incentives between loan officers and financial-capital owners to explain the emergence of discrimination (Paravisini, Rappoport, and Ravina (2017); Dobbie et al. (2020)), nor do we have intermediaries, such as real-estate agents, who might be the discriminating agent in lieu of the lender (Goldsmith-Pinkham and Shue (2020)). And, the fact that a robo-advising tool selects borrowers who perform better than those lenders select suggests that such tools might employ hard and soft information better than lenders (Iyer et al. (2016); Morse (2015)).[5]

[5]We cannot provide the details of the robo-advising algorithm, which is proprietary to the firm, but the tool uses more than the hard information studied in Iyer et al. (2016). Borrowers' performance is prima facie evidence against lenders' ability to infer borrowers' quality better than the robo-advising tool.

The second form of cultural bias we study is stereotypical discrimination, which, in India, as we discuss in detail below, arises against members of the lower traditional Hindu caste, namely, *Shudra*.[6] Even in this case, we find evidence consistent with discrimination: before lenders (including *Shudra* lenders) move to the robo-advising tool, *Shudra* borrowers are less likely to appear in their loan portfolios relative to the share of *Shudra* in the population of borrowers. Once lenders move to the robo-advising tool, the share of *Shudra* borrowers they finance increases substantially and the difference in defaults between *Shudra* borrowers and other borrowers drops. We also find that the discrimination against *Shudra* borrowers is higher for lenders who reside in states in which the share of crimes against lower-caste inhabitants is higher, which we interpret as settings in which discrimination against lower castes is more salient.

A unique feature of the Indian setting, which helps our analysis, is that the caste to which an individual belongs is not always easy to recognize. First, castes are not disclosed formally in any context our lenders can observe. Second, although certain individual characteristics might provide indications for the likelihood that the individual belongs to a certain caste—for instance, the surname, location, occupation, and complexion—in many cases, no certainty exists about an individual's caste unless the individual discloses it.

We exploit this unique feature to design an intensive-margin field test of stereotypical discrimination inspired by the experimental literature (Mobius, Rosenblat, and Wang (2016)). We build on Bhagavatula et al. (2017) and Bhagavatula et al. (2018), who propose an algorithm to replicate a human's assessment of the probabilities that individuals whose characteristics are observed belong to each of the four Indian castes. In this way, we obtain variation in the likelihood that lenders can recognize borrowers as *Shudra*. We find that discrimination against *Shudra* borrowers increases with the likelihood that the borrower is recognizable as a *Shudra*, whereas discrimination virtually disappears for *Shudra* borrowers whose caste lenders have difficulty assessing.

These intensive-margin results on stereotypical discrimination are important also because they allow us to disentangle our discrimination-based interpretation from an interpretation based on kin altruism—the tendency of individuals to take costly action to favor those who share similar ethnic or other demographic characteristics relative to others (e.g., see Simon (1993)).[7] Our results on in-group vs. out-group discrimination are consistent with both potential explanations because even in the case of kin altruism, Hindu lenders would be willing to pay a cost to provide funding to Hindu borrowers over Muslim borrowers, and vice versa for Muslim lenders. The fact that Shudra lenders are less willing to fund Shudra borrowers relative to other groups, though, although consistent with

[6] Less than 0.5% of borrowers and virtually no lenders are *Dalit* on Faircent, and hence, unfortunately, we cannot test for discrimination against this group.

[7] We thank David Thesmar for raising this point.

stereotypical discrimination, is the opposite of what kin altruism would predict.

Overall, our setting allows us to propose tests to detect and quantify the extent and effects of cultural biases in a consumer-finance field setting, in which the scope for statistical discrimination by lenders, which was hard to disentangle in earlier research, is minimal. We detect evidence that both in-group vs. out-group discrimination and stereotypical discrimination are prevalent and economically sizable in high-stakes lending decisions. These forms of discrimination make discriminating agents worse off in terms of consumption utility because discriminating lenders finance borrowers who perform worse than discriminated borrowers, who are also available on the lending platform.

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 (perhaps implicit) culturally biased choices. Interestingly, robo-advising tools might be a viable substitute for financial disclosure in improving the outcomes of consumers and individual investors, because they do not require them to understand all the aspects of the investment problems they face (Adams et al. (2019); D'Acunto et al. (2019b); D'Acunto et al. (2019a)).

2 Institutional Setting: Borrowers' Screening and Automated Investment Tool

The setting for our analysis is Faircent, a large FinTech platform that specializes in P2P lending in India, Faircent. As of October 2020, the platform hosts about 1.5 million borrowers and 140,000 individual lenders,[8] both of which groups reside across all Indian states. This 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)).

Several features of the platform are crucial to the design and interpretation of our tests. First and foremost, the screening of borrowers follows a two-step approach. The platform engages in an in-depth assessment of prospective borrowers' risk profile and screens borrowers based on this assessment *before* admitting them to the borrower pool that the lenders can access.

Once a prospective borrower signs up, he/she submits a loan application that includes the proposed amount of the loan, the motivation for the loan, the borrower's credit score, occupation, geographic location, and whether the borrower has dependents (children or elderly). The first (platform-based) screening step starts with an automated algorithmic-based assessment of the borrower's viability and

[8]This platform only hosts individual lenders and no institutional lenders.

ability to repay, which is largely based on the borrower's credit score, proposed loan amount and maturity, and occupation.[9] Based on this step, prospective borrowers that fall below a set threshold in terms of credit viability are dismissed from the platform.

This step captures the potential role for statistical discrimination due to machine-learning algorithms screening prospective borrowers (Bartlett, Morse, Stanton, and Wallace (2019), Fuster et al. (2017), Cowgill and Tucker (2020), Rambachan et al. (2020)), which minimizes the scope for statistical discrimination at the time lenders make their choices. In the Online Appendix, we provide raw-data evidence about the outcomes of the automated screening procedure in terms of the distributions of credit scores of borrowers who are accepted and those who are rejected (see Figure A.1), for which we find that rejected borrowers have systematically lower credit scores not only in the overall sample of applications, but also when we focus only on applicants who request loans of the same size. We also provide evidence on the stark relationships between borrowers' credit scores and the annual interest rates, maturity, and loan amounts that are assigned by the platform as a direct function of credit scores (see Figure A.2). The prospective borrowers who are approved and accept the parameters the platform assigns to them proceed to the second screening step. Borrowers who pass this step are assigned a risk category as well as a proposal for the interest rate and maturity of the loan based on the amount the borrower wants to raise.

Faircent's (human) employees perform the second screening step. This step consists of the in-person verification of several borrowers' characteristics. This step aims to eliminate lenders' potential concerns regarding the viability and accuracy of the automated risk profiling process on the platform. The characteristics Faircent employees verify personally include borrowers' identity (through scanned identity cards), an in-person connection via video call, a personal picture to compare with the identity-card picture, the proof of two income paystubs or incoming transactions in a bank account under the borrower's name, the proof of utility payments and addresses, the picture of the borrower's housing location, and the picture of the borrower's work location. Borrowers who fail these verification steps are dismissed from the platform. Borrowers who pass these verification steps are admitted to the borrower pool that lenders can browse.

The two-step screening procedure ensures that in our setting, lenders make choices *after* a substantive risk assessment has already been performed, whose outcomes lenders observe. Decoupling borrowers' risk screening from lending decisions, which departs from earlier research that studied the choices of loan officers, is important to tackle the question we propose: if certain demographic groups of borrowers were more represented among high-risk prospective borrowers, they would be rejected more often than other borrowers by the platform, and hence, the scope for statistical discrimination—

[9]Due to confidentiality, we cannot report the details about how the proprietary algorithm screens borrowers.

the possibility that lenders are more likely to finance certain demographic groups than others because such groups are less likely to belong to high-risk categories in the population—on the part of lenders is substantially lower in our setting than in loan-officer-based settings.

Finally, once viable borrowers are vetted they access the borrower pool, from which individual lenders access borrowers' demographic characteristics and the in-depth qualitative and quantitative risk assessment from the screening process. At each point in time, lenders observe borrowers' demographic characteristics (including names and surnames, location of residence, occupation, and picture), the characteristics of the loan required (loan amount, interest rate, maturity, and motivation for the loan if the borrower provided any), and the risk categories assigned during the platform's screening process. In this step, lenders decide whom they want to fund. Because lenders observe borrowers' characteristics, their lending choices might be influenced by taste-based discrimination based on such characteristics (e.g., see Ravina (2019)).

The second important feature of the platform for our purposes is the fact that lenders can make their choices in two modes—manually or automatically. Under the manual mode, lenders browse the borrower pool and decide who they want to fund and by how much. Lenders can fund up to one fifth of each loan amount; that is, the platform imposes that each loan is financed by at least five lenders.[10] Moreover, lenders need to have the funds they want to commit deposited on the platform before their proposals are posted.

The second mode of investment is an automated robo-advising tool, Auto Invest. Lenders can adopt Auto Invest at any time. Upon adoption, lenders decide the amount of funds they want to allocate to Auto Invest and the amount of funds, if any, they want to keep investing manually. Moreover, lenders can choose the share of funds they want to allocate across six risk-based categories of borrowers. Based on these risk-based directions, Auto Invest allocates the lenders' funds automatically to existing borrowers subject to the same restrictions as in the manual mode. The intent of choosing risk-based categories of borrowers is to mimic lenders' manual choices, because the six risk categories among which lenders allocate their funds on Auto Invest are the same risk categories they see as attached to borrowers once they appear in the pool.

Because of Auto Invest, we can compare manual lending decisions with automatic investment decisions. Moreover, because lenders can choose the share of deposited funds they want to commit to Auto Invest or the manual investment, we can study the difference in lending choices between humans and machines at the intensive margin, that is, comparing the choices of lenders who mostly lend through Auto Invest and the choices of lenders who mostly invest manually at the same point in time.[11]

[10] Faircent imposes this restriction to enhance the diversification of lenders' loan portfolios.
[11] Unfortunately, we do not observe whether each individual loan contribution was made by the lender manually or through

Finally, our data include not only loan applications and lending decisions but also loans' performance over time. Once the full amount of the loan is funded, 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 forms that lenders and borrowers have to sign. No lenders enjoy any form of seniority. Upon execution, Faircent provides borrowers with an equated monthly installment (EMI)—the monthly payment—and services the loan in house. Faircent monitors the status of loans each month. It declares the status as "closed" if the amount due, including principal and interest, is fully repaid or if the loan has been delinquent for more than 360 days. Borrowers whose loans are closed while delinquent are dismissed from the platform.

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.

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 the individual identifier Faircent assigns at signup, name and surname, the city and state location of residence, and the date of birth. On top of these characteristics, the borrowers' sample also includes information on 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 either lenders or borrowers for their religion and/or caste. Our source for information about these dimensions, which are crucial in our analysis, is the *Marriage registry data set* (see Bhagavatula et al. (2017) for an example of an earlier use of these data). These data include demographic information about religion and caste elicited at the time of marriage for a random sample of 2,481,158 Indians. Specifically, it includes names and surnames, date of marriage, state of birth, current city of residence, height (in centimeters), and religion and caste.

We find 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 date and shares the same surname) belong to individuals who are assigned the same religion. When

Auto Invest. For this reason, we cannot study how choices made manually or through Auto Invest at the same point in time vary within individual lenders.

we only consider the two religions on which our analysis focuses—Hindu and Muslim—96% of the surname-date-of-birth pairs are matched to only one religion. For these reasons, we assign religions to lenders and borrowers based on surname and date of birth, which we observe in the Faircent data.

Assignment of castes to borrowers and lenders is less straightforward. First, the caste information in the marriage registry is quite dispersed and includes about 540 narrowly-defined partitions, which sometimes merely correspond to the individual's surname. As we discuss in more detail below, to make our analysis of castes meaningful, we need a reliable way to assess to which of the four main *varnas* borrowers and lenders belong so as to split borrowers into those belonging to the Shudra varna—which we argue is subject to stereotypical discrimination—or other varnas.[12] Moreover, we do not find that any of the combinations of characteristics we observe in the Faircent data or the marriage registry, that is, names, surnames, and dates of birth,[13] restrict the possible set of castes enough to proceed in the same way as with religions. To infer borrower and lenders' varnas, we rely on earlier research in computer science that has tackled the same issue of assigning castes to individuals in cross-sectional data sets based on marriage-registry information and individuals' surnames. Specifically, we use the methodology developed by Bhagavatula et al. (2017) and Bhagavatula et al. (2018), which we discuss in more detail below.

The fourth data set we use is the *Lender-Borrowing Mapping*, which is a cross-sectional data set at the level of unique lender-borrower-loan triads. This information is critical to our ability 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. For example, if a borrower requested $10,000 and the money was lent by five different lenders, we have information on the amount lent by each individual lender to this borrower and the data set includes five lender-borrower-loan triads for this loan.

Our information about loans comes from a cross-sectional and a panel data set. The *Loan characteristics data set* is a cross-sectional data set at the loan level. For each loan, the data report the borrower to whom the loan was extended as well as the total amount lent, the interest rate, the maturity of the loan, and the proposed monthly payment. Moreover, we observe the loan's status as of March 31, 2020 (active or closed), as well as the whether the last payment happened within the previous 31 days (i.e., whether the loan is in good standing).

The *Loan performance data set* is a panel data set at the loan-month level. In this unbalanced panel, we observe the monthly payments that borrowers provided during the life of the loan each

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

[13] Note Faircent observes the borrower's/lender's location at the time they are active on the platform, whereas the location in the marriage registry is the individual's location of residence at the time of marriage.

month in which the loan's status was active (including zero if the borrower missed a payment in a month), for both active and close loans as of March 31, 2020.

Finally, the *Auto Invest data set* is a cross-sectional data set at the lender level, which provides us with information on whether lenders have ever activated the robo-advising tool called (Auto Invest) to make automatic lending decisions, and if yes, the activation date and the share of the funds available on the platform that the lender allocated to Auto Invest instead of keeping for manual lending choices. This data set allows us to compare the lending decisions and performance of financed loans at the lender level before versus after their activation of Auto Invest, as well as across lenders who allocated a higher or lower amount of funds to the robo-advising tool.

3.1. Summary Statistics

Armed with these seven data sets, we construct the two working samples we use in our analysis.

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. We select the sample by only including lenders who have activated Auto Invest at any point in time between its introduction on the platform in early 2018 and the end of the sample period (March 31, 2020). This selection implies all the lenders in the sample are adopters, and we do not compare the choices of lenders who never adopted Auto Invest with those of lenders who adopted it at some point in time.[14] However, at each point in time, the sample includes lenders who have already activated the tool and others who have not yet activated it. For the intensive-margin analysis of lending choices, instead, we only consider lenders at the time they have activated Auto Invest, across lenders who allocate a higher or lower share of their funds to the robo-advising tool, and hence who make a higher or lower fraction of choices autonomously.

The size of the sample selected in this way includes 113,283 unique lender-borrower-loan triads. Panel A of Table 1 reports summary statistics for this sample. Borrowers' religion, consistent with the split between Hindu and Muslim individuals in the general Indian population, is tilted toward Hindus, because only 13% of the borrowers are Muslim. The religious imbalance is even higher on the lender side, for which we find 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 meaningful analysis of Muslim lenders' behavior. About 45% of the loans in the data set were issued at a time when lenders had activated the robo-advising tool. The average share of funds allocated to the tool is about 60%, but substantial cross-sectional variation exists across lenders. As far as loan characteristics are

[14]All our results and patterns are substantially more pronounced if we also include lenders who have not adopted Auto Invest during the sample period.

concerned, the average maturity (tenure) is 22 months and the modal maturity is 24 months. The average loan amount is slightly above 130,000 rupees, which corresponds to about $1,770,[15] 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.[16]

The second working sample we create allows us to study the role of stereotypical discrimination in lending decisions. We select the sample using the same steps discussed above, and we further restrict it to the subsample of (Hindu) borrowers for whom we can retrace the caste in the form of one of the four varnas based on the marriage registry data. Overall, the final sample includes 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.

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

The first form of cultural bias we consider is 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-identified cleavages (Tajfel et al. (1979); Hewstone et al. (2002); Jenkins (2014)).[17]

A fundamental implication of this theory in terms of economic decision-making, which we can test directly in our setting, is that if favoring in-group vs. out-group members were driven by a cultural bias, lenders would be more willing to choose in-group borrowers even if they had to face a monetary cost from favoring them (D'Acunto et al. (2020)). In our setting, 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. Moreover, once lending decisions are automated through the robo-advising tool, 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 stark contrast to what we should find

[15] This conversion does not consider the varying exchange rate between the US dollar and Indian rupee in each month our sample covers.

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

[17] Unfortunately, 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).

if lenders' favoritism toward in-group borrowers were due to lenders' ability to screen and monitor in-group members better than out-group members, because in this case, in-group borrowers should perform better, on average, when lenders make decisions autonomously.

The Indian setting provides an ideal laboratory to study in-group vs. out-group discrimination in the context of religious conflict, and especially the conflict between Hindus—the 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)). This conflict has been vivid throughout India's history and has intensified since independence, that is, when the present-day territory of India was separated from the present-day territories of Pakistan and Bangladesh,[18] both of which hosted a stronger Muslim presence.

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, that propose Hindu nationalism as a key part of their political platforms, raised to power at the national level (see Kaul (2017) among others). For instance, the approval of the *Citizenship Amendment Act* in 2019 has produced a recent vivid wave of riots and violence between Hindus and Muslims covered by the local and international media (Bhat (2020)). Changes in relative incomes between Hindus and Muslims over the last two decades have also been proposed as a factor contributing to the recent increase in Hindu-Muslim violence (Mitra and Ray, 2014). As we discuss in more detail below, over the last two decades, violent riots and localized conflicts between Hindus and Muslims have been relatively more common in certain Indian states than others.

We exploit the religious conflict between Hindu and Muslim communities to 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 lenders were making all their decisions autonomously. Then, we compute the change in the propensities to lend to Hindus and Muslims once lenders moved to using the robo-advising tool (Auto Invest), which made decisions automatically and based on characteristics that were unrelated to borrowers' religion. Third, we assess whether the extent of lenders' in-group vs. out-group bias and debiasing were stronger in the areas of India in which the salience of the Hindu-Muslim conflict was higher, where we use the salience of the conflict as a proxy for the extent

[18]Present-day Bangladesh was part of Pakistan until its independence in 1971.

of cultural bias in this context.

4.1. Debiasing In-group vs. Out-group Lending with Robo-advising

We start by considering the extent of in-group vs. out-group bias in lending by Hindu and Muslim lenders in the raw data. Specifically, in Figure 1, we report the average share of Hindu and Muslim borrowers within the pool of borrowers for Hindu and Muslim lenders separately. We report these averages for the same lenders at two points in time; that is, the average shares for the period in which the lender was making loan decisions on his/her own and the average shares after the lender started to use the robo-advising tool.

The top-left graph of Figure 1 reports the share of lending to Hindu borrowers by Hindu lenders before (black bar) and after (red bar) using Auto Invest. The top-right subfigure, instead, reports the percentage of lending to Hindu borrowers by Muslim lenders before and after using Auto Invest.

Three broad patterns are worth noticing here. First, consistent with the presence of an in-group vs. out-group bias for at least one of the two religious groups, we find that for each religion, lenders tend to choose a higher share of borrowers of the same religion relative to the share of borrowers affiliated with that religion in the pool of non-homophilic lenders. The share of Hindu lenders' borrowers who are Hindu is 86%, whereas the share of Muslim lenders' borrowers who are Hindu is only 80%. Conversely, the share of Hindu lenders' borrowers who are Muslim is 12%, whereas the share of Muslim lenders' borrowers who is Muslim is 18%.

A second pattern is that, after they start to use the robo-advising tool, which does not consider borrowers' religion in the lending choice, the shares of Muslim and Hindu borrowers change for all lenders, and the changes are in opposite directions: Hindu lenders' borrowers who are Hindu decrease from 86% to 84%, whereas the share of Muslim borrowers increases from 12% to 13%. At the same time, the share of Muslim lenders' borrowers who are Hindu increases from 80% to 84%, whereas the share of Muslim borrowers decreases from 18% to 13%.

The fact that after using the robo-advising tool, lenders of different religions change their borrowing choices in *opposite* directions is important to 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, which would be hard to absorb without a proper identification strategy. 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, that is, 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.

The third fact that Figure 1 emphasizes is that the share of borrowers of different religions is equalized for Hindu and Muslim lenders after they use Auto Invest. This fact is emphasized by the red dashed horizontal lines crossing the four graphs of Figure 1. This fact reassures us that unobserved channels that might make in-group borrowers more profitable to lenders than out-group borrowers are unlikely to exist in our setting. Ultimately, the robo-advising 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, thus avoiding any favoritism toward in-group 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 the funds they have with Faircent that they want to allocate to Auto Invest. If lenders choose a share lower than 1, they can still make choices autonomously for the share of resources they do not allocate to Auto Invest. This institutional detail allows us to consider the extent of debiasing not only at the extensive margin (the choice of borrowers before and after lenders adopt Auto Invest) but also at the intensive margin (the choice of borrowers lenders who allocate a higher or lower share of their funds to Auto Invest make).

We report the analysis of this intensive margin of debiasing in Figure 2. In this figure, we only consider Hindu lenders who use the robo-advising tool.[19] We sort lenders based on the share of their funds with the platform that they decide to 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. Consistent with the results at the extensive margin discussed above, even in terms of the intensive margin, a greater use of Auto Invest—which, by construction, coincides with a lower share of choices that lenders make autonomously—is associated with a lower share of Hindu borrowers, and the negative relationship is monotonic. The larger the extent of the use of Auto Invest, that is, the lower the share of choices made autonomously, the lower the share of in-group borrowers that appear in lenders' portfolios. We detect a mirroring pattern for Muslim borrowers (green dashed line), which is mechanical given that the sample includes only Hindu and Muslim borrowers. Overall, even in terms of the intensive margin, the data reveal that a higher use of automated lending choices reduces the in-group vs. out-group bias we detected for lenders when they chose borrowers autonomously.

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

[19]We 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.

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 our setting, in Table 2, we report the results for estimating the following multivariate specification estimated at the level of borrower-lender couples:

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

where $Muslim\ Borrower_{i,j}$ is equal to 1 if borrower i who 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 $\boldsymbol{xi}$ 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. We cluster standard errors at the lender level.

The coefficient of interest, δ, measures the change in the probability of Hindu lenders lending to a Muslim borrower after activating Auto Invest.

Column (1) of Table 2 reports the baseline correlation without adding any control variables. We report this specification, which corresponds to the raw-data results of Figure 1, to assess the statistical significance of those results.

Indeed, the estimated coefficient $\hat{\gamma}$ is negative and statistically as well as economically significant, indicating Hindu lenders were about 5.4 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 estimate of the coefficient β on Auto Invest shows a non-significant effect of Auto Invest on the overall likelihood that Muslim borrowers appear in the data, on average, in the periods in which all lenders use the robo-advising tool, relative to before they started to use it. The lack of significance is due to the fact that the effect is positive for Hindu lenders and negative for Muslim Borrowers, as we show in the bottom subfigures of Figure 1, so that, overall, the share of Muslim borrowers in the sample does not change but is reallocated among Hindu and Muslim lenders.

The main coefficient of interest, $\hat{\delta}$, is positive and significant, indicating Auto Invest increases the likelihood that Hindu lenders lend to Muslim borrowers by about 3.4 pp. This coefficient denotes a drop in the in-group vs. out-group bias, but not a full debiasing, which would have implied an estimated interaction coefficient of around 5.4 pp. This lack of full debiasing is likely explained by the fact that, even after moving to Auto invest, lenders can still choose some loans on their on, above and beyond what the Auto Invest chooses for them. We verify this conjecture below and show the extent of debiasing monotonically increases with the ratio of loans made by Auto Invest relative to those chosen directly by lenders in the Auto Invest period.

In the second column of Table 2, we move on to the multivariate specifications. We find the tenure and the loan amount are significantly related to the probability of the borrower being Muslim. The first is positive and significant, whereas the second is negative and significant. However, 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 to the baseline specification (column (3)). In this way, we assess the scope for discrimination by Hindu lenders against Muslim borrowers after absorbing systematic time-invariant differences across lenders, including unobserved time-invariant dimensions such as financial literacy, cognitive skills, and education levels. This specification absorbs the lender's religion, but we can see that the estimated coefficient on the likelihood of choosing Muslim borrowers after using Auto Invest stays almost identical to the results in columns (1) and (2).

In the last two columns of Table 2, we focus on the intensive margin of usage of the robo-advising tool. Column (4) reports the multivariate estimates for lenders who allocate less that 40% of their funds to Auto Invest, whereas column (5) reports the results for those who allocate more than 40% of their funds to Auto Invest.[20] Based on the univariate results in Figure 2, we expect that the strength of the effect of debiasing increases with the percentage of funds that lenders allocate to Auto Invest. Consistently, the $\hat{\delta}$ coefficient is small and insignificant for lenders who do not use Auto Invest intensively and positive and significant for other lenders. Comparing the estimates with the corresponding ones in the full sample with lender fixed effects and loan-level characteristics, the positive estimated increase in lending to Muslim borrowers by Hindu lenders after using Auto Invest is about 20% larger for the set of lenders who use Auto Invest more intensively.

Overall, both our univariate and multivariate estimates suggest that, once lenders move to using

[20]The results are similar irrespective of the choice of threshold.

the robo-advising tool, they debias in the sense that they do not display a favoritism for choosing in-group borrowers and avoiding out-group borrowers, and we detect this form of debiasing for both Hindu and Muslim lenders, whose lending choices move in opposite directions once they move to Auto Invest in terms of the religion of the borrowers who enter their portfolios.

4.2. Heterogeneity: Salience of the Hindu-Muslim Conflict

The baseline results so far report evidence of in-group vs. out-group bias in lending as well as evidence of debiasing for the average investor in our P2P platform after they start using a robo-advising tool that creates loan portfolios without considering borrowers' religion as a choice variable. Because the bias goes in opposite directions for Hindu and Muslim borrowers before they access robo-advising, and hence, the reaction of their portfolios also goes in opposite direction after they all access robo-advising, most concerns about the endogeneity of adoption of Auto Invest have little relevance in our context. In particular, standard concerns about unobserved time-varying shocks that might simultaneously explain the decision to adopt robo-advising and the reaction of lenders' portfolios are barely relevant here, because the reactions are very different and the lender's religion determines the direction of the reaction.

At the same time, the Indian setting allows us to dig deeper into our understanding of the role of in-group vs. out-group bias by providing scope for cross-sectional heterogeneity tests. Specifically, based on earlier research in psychology, sociology, and economics, we can isolate sources of heterogeneity in the extent to which the Hindu-Muslim conflict should be salient to lenders, based on the geographic location of lenders and/or their demographic characteristics. For each of these sources of variation, we can thus test the hypothesis that the extent of bias in lending is larger for lenders for whom the Hindu-Muslim conflict is more salient, because for this group, the scope for in-group vs. out-group bias is higher. Therefore, under each of the dimensions we introduce below, based on our interpretation of the baseline results, we would expect the following: (i) Before accessing robo-advising, lenders for whom the conflict is more salient tend to bias their lending against the out-group borrowers and in favor of the in-group borrowers more than other lenders, and (ii) after accessing robo-advising, everybody's shares of lending to each specific religious group are equalized, and hence, the portfolios of lenders for whom the conflict is more salient react more after adopting robo-advising than the portfolios of other lenders.

_4.2.1_ Hindu-Muslim Riots The first dimension that makes the Hindu-Muslim conflict more salient is localized Hindu-Muslim riots, which, although recorded in India for decades, have substantially increased in frequency and the extent of violence over the last two decades (Oza (2007)).

The violence of riots, the aggressive positioning and rhetoric of local politicians on both fronts, and the local media coverage are likely to make the Hindu-Muslim conflict more salient to lenders who are more exposed to such dimensions (D'Acunto et al. (2019)), that is, lenders who reside in locations with a high incidence of riots. Based on this rationale, we expect that the extent of bias against out-group borrowers by in-group lenders is higher for lenders residing in locations that have faced more Hindu-Muslim riots, whereas it is lower for lenders residing in locations in which riots were more limited in number and scope.

For this test, we build on Ticku (2015), who categorizes and collects the occurrence of Hindu-Muslim riots at the local level from 1980 to 2000. The geographic aggregate we consider is the state level. This choice is driven by a few considerations. First, all lenders in our sample are mapped to an Indian state, which they have to choose through a pre-specified menu of available states of residence at the time of signup and is verified personally by Faircent employees based on the lenders' documentation. At the same time, not all lenders report the city in which they reside: If we were considering the city level, we would need to drop all lenders who do not report their city of residence, which would restrict the sample and select it in a way that might potentially correlate with lending choices. Second, regulation and policies that influence the relationship between religious groups are often implemented at the state level; for example, see the case of "anti-conversion laws" (Jenkins (2008); Dhattiwala and Biggs (2012)). Moreover, deep-rooted cultural norms shaped at the local level persistently relate to present-day violence across religious groups in India (Jha (2014)).

Panel A of Figure 3 depicts the cross-state geographic variation we employ based on the incidence of riots as documented by Ticku (2015). In the map, dark-green states (Gujarat, Marahashtra, Karnataka, and Uttar Pradesh) are those in which Hindu-Muslim riots have been most prevalent, whereas riots in other states were less prevalent.

In columns (1)-(2) of Table 3, we report the results for estimating equation (1) separately for the group of lenders in states with a low incidence of Hindu-Muslim riots over the period 1980-2000 and those in states with a high incidence of riots. Consistent with our conjecture, we find the extent of bias by Hindu Lenders against Muslim borrowers was about twice as large for Hindu lenders in the latter group of states (5.9 pp), and statistically significant, whereas the size of the bias for lenders in other states was about 3 pp, and we cannot reject the null of no bias statistically. Conversely, the reaction of Hindu lenders' portfolios after the adoption of robo-advising was substantially higher for lenders in states with a high incidence of Hindu-Muslim riots (3.6 pp, $p<0.1$) relative to the reaction of Hindu lenders' portfolios in states with a lower incidence of riots (2 pp, $p>0.1$). Our results are similar if we change the threshold of the number of local riots based on which we split lenders into one group or the other.

4.2.2 Electoral Support for the BJP

We move on to consider a second dimension that varies in the cross-section of lenders and that earlier research has associated with the salience of the Hindu-Muslim conflict: the local vote share for the *Bharatiya Janata Party* (BJP). A defining feature of the BJP is that the cultural and ideological roots of its platform have always been based on traditional Hindu values (e.g., see Berglund (2004) and Chhibber and Verma (2018)). In particular, the BJP's value and ideological platform includes fostering the notion of *hindutva*, which implies a coincidence between the spheres of Indian culture and traditional Hindu values (e.g., see Prakash (2007) and Chidambaram et al. (2020)). The BJP is the result of a set of mergers of post-independence parties in India and has shared a leading role in Indian national and state-level politics with the *Indian National Congress* since independence (e.g., see Ziegfeld (2020)). In contrast, the BJP, the Indian National Congress has been proposing instances of secularization and, even when supporting traditional Hindu values as defining Indian public life, has been less supportive of conflict between the Hindu majority and Muslim minority (e.g., see Ganguly (2003) and Verma (2016)).

Based on these considerations, we exploit the state-level variation in vote shares of the BJP to capture variation in the extent to which local lenders might see traditional Hindu values as foundational, and hence might be more likely to display an in-group vs. out-group bias against the Muslim population. Note we are not arguing 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 local lenders.

We 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. For our test, we compute the vote share of BJP candidates in each election cycle and state and then compute the average BJP vote shares across elections within each state. Panel B of Figure 3 reports the state-level distribution of the average vote share for the BJP between 1977 and 2015.

We compute the average of shares within states, because different states vote for state-level elections in different years and across different cycles. All our results are virtually identical—that is, the allocation of states to those in the high-BJP-vote-share group and low-BJP-vote-share group does not change—if instead we only consider votes cast for national elections as well as if we only consider votes cast over the last two decades instead of since 1977. As is clear from Panel B of Figure 3, using vote shares for the BJP provides us with more variation and a different variation across states than the occurrence of Hindu-Muslim riots depicted in Panel A.

We report the results for the heterogeneity analysis based on the vote share for the BJP in columns (3)-(4) of Table 3. When we split our lenders into two group based on the vote share for the BJP, we find the heterogeneity of the effects is even starker than for the variation in Hindu-Muslim riots: the extent of bias against Muslim borrowers by Hindu lenders was 8.7 pp before accessing robo-advising in high-BJP-vote share states, whereas it was only 2.9 pp—a size that hinders us from rejecting the null of no bias statistically—for lenders in low-BJP-vote-share states. Conversely, after lenders start using the roboadvising tool, those in high-BJP-vote-share states debias fully both economically and statistically, whereas the "debiasing" of other lenders (who did not initially display a substantial anti-Muslim bias in lending) is economically and statistically negligible.

4.2.3 Exposure to Hindu-Muslim Conflict during Formative Years

The last heterogeneity dimension we consider is not based on variation across space, but on variation across lender cohorts. Here, 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 during his term as Chief Minister of the state of Gujarat from 2001 to 2014, and especially since he became Prime Minister of India in 2014 (Chhibber and Verma (2014)).

The stronger support for the BJP and especially its rise to national power has pushed the issue of Hindu-Muslim relations toward the top of the list of topics in Indian political discourse. To design a cross-sectional test, we build on the recent literature in economics and finance that documents the long-run effect of beliefs and convictions that agents develop during their formative years (late childhood and early adolescence) on their beliefs and convictions later in life (Malmendier and Nagel (2011)). 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 afterwards, during their adulthood, at a time when their political beliefs and convictions were likely already cemented. Specifically, in our context, we compare lenders who were born after 1990 (i.e., they were 24 or younger when Narendra Modi became Prime Minister) and lenders born before 1990.[21]

We report the results for this third heterogeneity test in columns (5)-(6) of Table 3. Column (5) includes lenders born after 1990, and column (6) includes those born before 1990. The extent of pre-robo-advising bias by Hindu lenders against Muslim borrowers is substantially larger for lenders born after 1990 (6.6 pp) and statistically significant, whereas the bias is smaller (1.8 pp) for lender whose convictions were formed before the BJP's rise. It is so small that we fail to reject the null of no bias. Similar to the other tests, we also find the debiasing effect is large for lenders born after 1990 (5.9 pp), who in fact are almost completely debiased after they access robo-advising, whereas the effect is economically and statistically zero for older lenders, who did not initially display a substantial bias against Muslim borrowers.

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

Overall, our heterogeneity tests provide evidence that collectively points in the same direction: the lenders for whom the Hindu-Muslim conflict is likely to be more salient, and for whom the in-group vs. out-group bias based on social conditioning is therefore likely to be stronger, are indeed the lenders who display the stronger bias against out-group borrowers before accessing the robo-advising tool. Consistently, they are also the lenders for whom the bias decreases substantially after using the robo-advising tool, relative to before using it.

On top of providing evidence consistent with lending choices being driven by in-group vs. out-group bias, these heterogeneity results further reduce concerns related to the endogeneity of the adoption of Auto Invest in our context: if the endogenous adoption were explained by unobserved factors that simultaneously explain the change in outcomes, not only should these factor predict changes in outcomes of opposite signs for Hindu and Muslim lenders, ceteris paribus, but also within each group, the unobserved factors should exist for lenders for whom the Hindu-Muslim conflict is more salient but not for other lenders, whether we proxy for varying salience through geographic location or through lenders' cohorts.

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 certain social groups, because society attaches negative stereotypes to the members of such groups (Becker (1957),Akerlof and Kranton (2000)). Stereotypical discrimination differs from the in-group vs. out-group bias we discussed above in at least two relevant dimensions: First, everybody shares the negative stereotypes associated with members of the discriminated group, *including* the members of the discriminated group, who might thus engage in actions that discriminate against members of their own group (Jost and Banaji (1994); Nosek et al. (2002); Pritlove et al. (2019)). For 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, also known as *implicit bias* (Bertrand et al. (2005); Brownstein (2015)).

Moreover, stereotypical discrimination is often based on deep-rooted cultural norms and beliefs that have developed in society and tend to be highly stable over time (e.g., see D'Acunto et al. (2019); Payne et al. (2019)). Whereas, as we saw in the case of Hindu–Muslim conflict, events and shocks might modify the extent of in-group vs. out-group bias, the extent of stereotypical discrimination tends to persist in the long run.

The Indian setting is well suited to the study of stereotypical discrimination, because of the centuries-long and persistent stereotypes attached to members of lower castes. Castes, as we discuss in more detail below, represent social divisions that are deeply ingrained in India's Hindu heritage and that create a hierarchical ranking between groups (Dumont (1980)), which produces a set of positive stereotypes associated with members of the highest-ranking castes and negative stereotypes associated with members of the lowest-ranking castes, which have been studied in sociology for decades (see, e.g., Sinha and Sinha (1967)).

A second unique feature of the Indian setting is that variation exists in the extent to which somebody's caste can be recognized based on observational characteristics such as names and surnames, physical appearance, and occupation, especially without being part of the same caste (Muthukumar (2020)). Certain last names are more recognizable than others as belonging to a specific caste, and certain last names may belong to one caste in a certain community but another caste in other communities. This variation in caste recognizability is why India provides a natural setting in which to test for the extent to which the salience of one's caste based on observational characteristics, rather than the (unobserved to lenders in our setting) actual belonging to a discriminated caste, affects lenders' choices of borrowers. This feature would not exist in other settings such as, for instance, the US, in which instead researchers have exploited the fact that certain names, surnames, and physical traits make one's social group immediately recognizable to any other member of society.

Before motivating the empirical tests in more detail and discussing the results, we provide a concise summary of the institutional feature of the caste system in India and especially of the aspects that are most relevant to our tests.

5.1. A Concise Primer on the Indian Caste System

Based on a set of traditional and foundational Hindu writings, 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)).[22] In the traditional interpretation, these social groups have a strict hierarchical relation to one other. *Brahmins* were the highest caste, and traditionally included the Hindu clerics cast, as well as teachers and researchers, that is, all those who dedicated their lives to contemplative activities rather than engaging directly in administration or manual work. The second caste (*Kshatriyas*) traditionally encompassed warriors and rulers. Members of this caste historically covered governmental and military positions. The third caste, the *Vaishyas*, included farmers, traders, and merchants. Historians have emphasized similarities between the notion of *Vaishyas* varna and the "Third state" in pre-revolutionary France, for instance.

[22] Here, 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.

Against the three top varnas stands the (*Shudra*) caste, which has historically included laborers such as peasants and servants in various roles (Ambedkar (1947)). This caste was explicitly lower ranked than other castes, and its members were employed in roles of service to the benefit of higher castes. The centuries-long strict implementation of societal roles based on the caste system has created a set of cultural stereotypes attached to members of each caste and consequent outright discrimination towards the members of the *Shudra* caste.

Note that members of the outcast group, the *Dalits*, have faced even stronger discrimination and segregation over the centuries, both physically and socially (Maikkēl (1999)). Perhaps as a consequence of this strong discrimination and segregation, unfortunately, less than 2% of the borrowers in our platform are *Dalits*, which makes meaningfully studying the effects of lending between them and members of the four varnas impossible for us. For this reason, our empirical tests and analysis focus on comparing borrowers' lending decisions toward *Shudra* borrowers relative to borrowers who belong to the three higher castes. Importantly, as discussed above, a tenet of stereotypical discrimination is that the members of the discriminated group themselves might have an implicit bias that is so ingrained that they discriminate against other members of their own caste. In our context, this feature suggests *Shudra* borrowers are likely to be discriminated against by both *Shudra* lenders and other lenders.

5.2. Proxying for the Recognizability of Borrowers' Caste

As we discussed when introducing the raw-data evidence on stereotypical discrimination, the Indian setting is unique in that substantial variation exists in the extent to which somebody's caste can be inferred from baseline information. In our P2P application, lenders observe borrowers' names and surnames, see their face picture, and some other limited set of demographics. The most salient information on the platform relates to the borrowers' risk characteristics, for which we control directly in our multivariate analyses, such as the loans' interest rates, maturity, and amounts, which are ultimately set by the platform's algorithm.

In other cultural environments, such as the US, the recognizability of somebody's ethnicity based on name, surname, and face picture is very high, and in fact, earlier research in psychology, sociology, and economics has exploited the salience of ethnic features observed from these three pieces of evidence to assess how choices depends on agents' ethnicity. In the case of India, instead, these three features do not always provide immediate recognition of one's caste.

Indeed, features such as names, surnames, occupations, and complexion sometimes help agents identify each others' caste, but they do not. Ultimately, whether a potential borrower belongs to the *Shudra* varna is easier or harder to assess on a case-by-case basis.[23]

[23]Recall that, due to the lack of Dalit borrowers and lenders on the platform, we do not have this fifth group in our

We exploit these two features—variation in caste recognizability based on demographic information and based on whether lenders and borrowers belong to the same community—to design heterogeneity tests for the effects of stereotypical discrimination in lending.

We first focus on proxying for the extent to which a borrowers' caste is easily or minimally recognizable based on demographic characteristics. Ideally, we would have measured directly the extent to which lenders recognized the borrowers' caste, for instance, by asking them to guess the caste when facing a potential borrower and then comparing this lender-reported caste with the actual caste of the borrower. Unfortunately, had we asked explicitly, we would have invalidated our test: making borrowers' caste salient by asking would have interfered with the lenders' decisions, potentially in both directions. For instance, lenders could have willingly changed their discriminating behavior relative to what they would have done without our intervention, to pretend they were not discriminating against lower castes in daily life. For this reason, we could not directly elicit the extent of lenders' ability to recognize castes.

To overcome this issue, we instead implemented an off-the-shelf algorithm 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. Specifically, we rely on the methodology developed and detailed in Bhagavatula et al. (2017), of which we provide a brief summary here.[24] The procedure relies on two unique aspects of the caste system in India. First, castes are endogamous—marriages occur mainly between individuals that belong to the same caste. Second, in India, 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. Varnas are self-reported, and the possibility of misreporting is virtually non-existent, because varnas are a fundamental variable that matrimonial agencies use to match potential couples, due to the fact that castes tend to be endogamous. In this step, the procedure also groups together all the different variations of the same last name.

In the second step, the procedure assigns one or more castes to each of the last names. The relation between caste and last names is not unique, because the same last name can be associated with different castes especially in different cities or regions. The algorithm assigns the probability of a last name belonging to a given caste to equal the proportion of times the matrimonial website users with that last name self-identify as belonging to that caste in every state and city.

In the third step, we assign a caste (and its probability) to each borrower and lender on the platform

sample.

[24] We thank Manaswini Bhalla for graciously running the algorithm developed in Bhagavatula et al. (2017) and Bhagavatula et al. (2018) on our data.

on the basis of the user's last name, state, and city.

Figure A.3 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, which includes about 80% of the full sample of borrowers on the platform. Except for a fat right tail of 17% of borrowers whose probability of being *Shudra* is close to 1, the probability is distributed rather homogeneously throughout the support.

5.3. Debiasing Stereotypical Lending with Robo-advising

We start by showing in Figure 4 the motivating plots on stereotypical discrimination in lending based on borrowers' castes. The top three graphs report the average of *Shudra* borrowers within lenders' portfolios for lenders belonging to any castes. The leftmost graph considers the full set of borrowers in a lender's portfolio of loans; that is, it includes borrowers whose caste is easy to recognize based on names, surnames, and face pictures (which our lenders observe on the platform), as well as borrowers whose caste is not recognizable. In this case, we detect no difference in the share of *Shudra* borrowers before and after lenders use the robo-advising tool: the share of loans going to *Shudra* borrowers is virtually the same before and after access to the tool and is equal to 31%, which is close to the share of *Shudra* in the overall borrower population on the platform. This result is consistent with the possibility that, when castes are barely identifiable, no discrimination exists against *Shudra* borrowers.

Because we instead restrict the sample to subgroups in which the borrowers' caste is more and more recognizable, whatever the case, as we do while moving towards the graphs on the right, we observe a very different pattern. In subfigure (b), the pre-Auto Invest lending to *Shudra* borrowers was only 27%, which was further reduced to 25% in subfigure (c), where we restrict the sample to borrowers whose caste is highly recognizable. Crucially for our interpretation, in terms of stereotypical discrimination, we find the lending after the use of Auto Invest is virtually identical across all three subpopulations and equal to 31%.

Overall, the three graphs on top of Figure 4, which are based on raw data, seem consistent with the presence of stereotypical discrimination in our sample: if castes are not recognizable, *Shudra* borrowers are treated similarly to other borrowers. But as castes become more and more recognizable, whenever lenders make decisions on their own, they discriminate more and more against *Shudra* borrowers.

As discussed above, a defining feature of stereotypical discrimination relative to in-group vs. out-group discrimination is that even the members of the same discriminated category are often subject to the implicit stereotypical bias that is ingrained through social conditioning. For this reason, even *Shudra* lenders should discriminate in favor of higher-caste borrowers and against *Shudra* borrowers, ceteris paribus. The Indian setting, in addition, allows for an additional unique proof of concept: as the extent of caste recognizability varies, one might expect that *Shudra* lenders more easily recognize members of the same caste in cases in which members of other castes might have a harder time, because certain features that a *Shudra* might recognize as familiar might not resonate with other castes' lenders. For this reason, if anything, the extent of discrimination of *Shudra* lenders against *Shudra* borrowers might be *higher* for the same levels of caste recognizability by third parties.

And, indeed, the raw-data plots in the bottom three subfigures of Figure 4, in which we restrict the sample to include only *Shudra* lenders, provide evidence consistent with this conjecture. Not only do *Shudra* lenders behave similarly to *Brahmin*, *Kshatriyas*, and *Vaishyas* lenders to *Shudra* borrowers, but if anything *Shudra* lenders discriminate against *Shudra* borrowers more than other lenders. For all levels of recognizability, in the raw data, *Shudra* lenders tended to provide about 1 pp lower loans to *Shudra* borrowers than to lenders of other castes.

These plots represent interesting motivating evidence but are based on raw data and do not keep constant borrowers' characteristics, which might vary systematically across lenders' portfolio. Moreover, we need to assess the statistical significance of the raw-data results. For these reasons, in Table 4, we report regression results based on the following specification:

$$Shudra\ Borrower_{i,j} = \alpha + \beta\ Auto\ Invest_j + \zeta\ \boldsymbol{x}_i + \epsilon_{i,j}, \tag{2}$$

where the regressors are defined as in equation 1.

Table 4 contains six columns. The first two columns do not impose constraints on caste recognizability. Consistent with the results in subfigure (a), we find no relation between Auto Invest and the probability of lending to *Shudra* in this case. In the second column, we add borrowers characteristics, and we find the coefficient becomes positive and significant, indicating that—once we control for borrower characteristics—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. Finally, the coefficients on the control variables *tenure in months* and *loan amount* are negative and significant, suggesting Shudra borrowers have been borrowing for less time and request lower loan amounts.

Column (3) and (4) impose that the probability of recognizing the caste by way of the algorithm is at least 50%. As expected, the results becomes more significant for the baseline results, even though they just miss the 10% significance level. They remain positive and significant for the results that include covariates. Finally, the last two columns constrain the analysis to the castes that are easily recognizable—the classification probability is greater than 70%. When we impose this restriction, we find the relation between lending to Shudra borrowers and adopting AutoInvest is positive and strongly significant, suggesting once again that before using Auto Invest lenders were actively discriminating against *Shudra* borrowers.

5.4. Heterogeneity: Variation in the Salience of Negative Stereotypes Attached to Lower Castes

The results discussed so far refer to all lenders and borrowers. Similar to our analysis of in-group vs. out-group discrimination, in this section, we exploit pre-determined variation in the extent of the scope for discriminating members of the Shudra caste on the part of lenders. Specifically, we consider two heterogeneity dimensions: the extent of inter-caste hatred across geography (Indian states), which, based on earlier research, we capture with the local incidence of crimes against members of lower castes, as well as the extent to which borrowers' castes are systematically more or less recognizable based on whether borrowers belong to the same community as the lenders.

_5.4.1_ Hatred Crimes against Lower Castes Recent research documents a substantial increase in hatred crimes against members of lower castes over the last decade, which has been heterogeneous across space. Researchers argue that several social factors might have driven this phenomenon, in particular, the fact that, even though at a very slow pace, lower castes have increased their economic well-being and have been the target of policies aimed at guaranteeing quotas of public employees in both national and local institutions from lower castes (Sharma (2015); Bapuji and Chrispal (2020)). Ultimately, the slow decrease in disparities between the economic and social conditions of higher castes and lower castes might have triggered the increase in acts of hatred against lower castes.

We conjecture that the stereotypical discrimination of *Shudra* borrowers might be higher in areas in which the conflict between higher and lower castes might be 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)). We report the results when using crime rates based on the 2018 annual report, but all our results are similar when using other recent years, because the ranking of states based on the rate of crime against lower-caste victims is quite stable over the recent years.

Figure 5 plots the cross-sectional variation in crimes against lower-caste victims, and we estimate equation (2) separately for lenders in Indian states above and below the median rate of crimes against lower castes per 100 inhabitants in the state-level sample (18.8). We report the results in columns (1)-(2) of Table 5. Consistent with our conjecture, we find that after the adoption of robo-advising, the share of Shudra borrowers in any lenders' portfolio increases by 2.7 pp, which is about 12% more of the average share of Shudra borrowers in lenders' portfolios before accessing the tool as captured by the constant term. When focusing on the subsample of lenders who live in states below the median, we find that the share of *Shudra* borrowers in portfolios only increases by 1.3 pp, and we fail to reject the null that this estimated increase is equal to zero statistically.

Moreover, we estimate the baseline specification separately for the largest urban areas in the sample—the city of Delhi and the city of Mumbai, which is the largest city in Maharashtra. We consider these conglomerates separately, because even if violence against low castes is frequent in these cities and often reported in the local media, and hence salient, because of the large and dense population, the crime rate is perhaps not a good proxy for the salience of inter-caste hatred. And, indeed, we find that lenders who live in these large urban conglomerates behave similarly to lenders in states with a high crime rate against lower castes, with an estimated higher ratio of Shudra borrowers after robo-advising of 2.2 pp out of an average of 19% before robo-advising.

5.4.2 Borrowers and Lenders in the Same Communities Another reason to focus on the large urban conglomerates of Delhi and Mumbai is to study the variation in the second dimension that might make a caste systematically more recognizable, namely, whether the lender and borrower belong to the same community, which in this case we capture with them being in the same large city, that is, at the finest geographic location we observe in our data.

We limit this test to the two largest urban conglomerates, because if we consider other cities, the share of borrowers and lenders in the same city drops dramatically; as we discussed in the data section, the P2P platform connects lenders to borrowers from all over India, and hence, when we choose lenders from smaller cities, we do not have meaningful variation in the share of their borrowers who are local or not—almost all borrowers are not local. This variation, instead, is meaningful for Delhi and Mumbai.

In columns (4)-(5) of Table 5, we estimate equation (2) for the subsample of lenders who live in Delhi and separately for borrowers who also live in Delhi (same-state borrowers) and other borrowers (different-state borrowers). We find the size of the increase in the share of *Shudra* borrowers' in Delhi's lenders portfolio is about three times higher after accessing robo-advising for Delhi borrowers than for non-Delhi borrowers. That is, after accessing robo-advising, the Auto Invest algorithm assigns more Shudra borrowers to all Delhi lenders, but it picks disproportionally more *Shudra* borrowers from Delhi relative to Sh *Shudra*udra borrowers from outside Delhi. This pattern is consistent with the possibility that Delhi lenders were discriminating against Delhi *Shudra* borrowers—who might have been easier to recognize—relative to *Shudra* borrowers from other cities, who they might have more difficulty recognizing. Even in this case, the fact that the constant is higher for same-state *Shudra* borrowers does not contradict the possibility of Delhi lenders' higher discrimination of Delhi *Shudra* borrowers, because this share has to be compared with the share of overall available borrowers in Delhi who are *Shudra*, and not with the share of Shudra borrowers among those outside Delhi, where the *Shudra* represent a lower share of the overall population.

We propose a similar split in columns (6)-(7) of Table 5, in which we limit the analysis to lenders living in Maharashtra, which hosts Mumbai. The sample size here is substantially smaller, but we still find that the average increase in the share of *Shudra* borrowers among the local borrowers increases by about 0.5 pp more after robo-advising relative to the increase in the share of *Shudra* borrowers from outside Maharashtra, whom Mumbai lenders might have had more difficulty recognizing. This difference, though (as well as the baseline coefficient on Auto Invest for same-state borrowers), is just below the standard level of statistical significance, which might not appear surprising given the limited sample size in this test.

6 Cultural Debiasing and Lenders' Performance

Sections 4 and 5 have studied the effects of discrimination in two different domains. First, we considered in-group vs. out-group discrimination: Muslim lenders tend to prefer lending to Muslim borrowers, Hindu lenders tend to prefer lending to Hindu borrowers. As a result, when individuals adopt Auto Invest, Muslim lenders lend more to Hindu borrowers (and less to Muslim borrowers), and vice versa. Second, we studied stereotypical discrimination by considering the patterns of lending by caste, and found that *all* lenders, including Shudra lenders, tended to discriminate against Shudra borrowers before their lending decisions were automated through an algorithm that does not consider castes a dimension for assigning credit. In this section, we move on to assess the financial implications of cultural biases in lending; that is, we ask the following question: Does the cultural debiasing associated with the use of a robo-advising tool improve or worsen lenders' performance?

6.1. In Which Direction Should Performance Change?

Intuitively, if lenders have biases in favor of or against a certain religion or caste, they might need to 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.

Our setting displays two features that make it appropriate for testing for the presence and effects of cultural biases. First, it abstracts from the potential screening and monitoring roles of in-group lending, which are instead documented and studied extensively in Fisman, Paravisini, and Vig (2017) and Fisman, Sarkar, Skrastins, and Vig (2020). Our platform connects lenders and borrowers all over India and does not involve a form of localized relationship lending, in which the loan officer of a local branch of a commercial bank screens and monitors local borrowers on behalf of the lender (the bank) when making her lending decisions. This screening/monitoring channel is a defining feature of local-branch lending by loan officers to local borrowers, who live and operate in the same small social environment as the loan officers. In that setting, this channel is so compelling that, quantitatively, Fisman, Paravisini, and Vig (2017) and Fisman, Sarkar, Skrastins, and Vig (2020) show it prevails relative to the negative effects of discrimination in terms of borrower quality in the context of localized relationship lending. The fact that, in our setting, the screening/monitoring channel is shut down makes it an ideal setting in which to test for the effects and value of cultural discrimination in the field. More broadly, our design is reminiscent of other work that uses online applications to abstract from the social-network role of personal connections in financial decisions, by studying the choices of agents who obtain information about others they do not know and with whom they do not interact in their daily lives.

The second unique feature of our setting is that the lending decision does not simultaneously include an assessment of borrowers' risk profile and a potential discriminatory choice: the risk levels of borrowers and their loan requests are assessed ex ante by the Faircent algorithm when the borrowers are admitted to the P2P platform. When making decisions, lenders see this risk assessment of potential borrowers, which is unrelated to borrowers' religion or caste. By studying lending decisions under the robo-advising tool, we can ex-post verify directly that the Faircent algorithms does not use religion or caste when assessing borrowers' risk, by checking that the shares of borrowers by religion or caste are equalized across lenders of different types (see Figure 1 and Figure 4). Decoupling the risk assessment from the lending decision, which are conjoint decisions for loan officers of a commercial bank's local branch, allows us to consider deviations from the robo-advising-determined outcomes as driven by active discrimination on the part of lenders rather than by systematic heterogeneity in risk levels that might correlate with borrowers' religion or caste.

Because the screening/monitoring channel is inactive in our case, our prediction is the *opposite* relative to the result in Fisman, Paravisini, and Vig (2017) and Fisman, Paravisini, and Vig (2017). We conjecture that cultural debiasing improves rather than worsens lending performance.

6.2. Performance after Debiasing In-group vs. Out-group Discrimination

6.2.1 **Extensive Margin of Performance** To assess the evolution of the extensive margin of lending performance—whether borrowers default on their loans—we estimate the following baseline specification:

$$Delinquent\ Loan_{ij} = \alpha + \beta\ Auto\ Invest_j + \zeta\ \boldsymbol{x}_i + \epsilon_{i,j}, \tag{3}$$

where $Delinquent\_Loan_{ij}$ is equal to 1 if the loan associated with borrower i and lender j is closed as delinquent. We report the results for estimating the baseline version of this specification for in-group vs. out-group discrimination in columns (1)-(2) of Table 6.

In column (1), we find Hindu lenders improve the extensive margin of their performance after accessing the robo-advising tool, in the sense that their portfolios of borrowers become less likely to default on loans. The estimated $\hat{\beta}$ coefficient is negative and significant, and suggests that loan delinquency decreases by about 9 pp for Hindu lenders. This effect corresponds to a 40% drop in the likelihood of delinquency after Hindu lenders access robo-advising relative to before. In column (2), we estimate a similarly large negative coefficient for Muslim lenders, which implies that the likelihood of default in their loan portfolios drops by as much as 65% (0.326/0.505).

These results suggest that lenders' performance improves after using Auto Invest relative to before. This drop could be due to cultural debiasing or by other changes in lenders' portfolios after adoption of Auto Invest. Unfortunately, we could not set up a randomized-controlled trial at the time the tool was introduced to lenders, and hence, we cannot answer this question directly with a causal test. Instead, we tackle this question in two steps as follows. Consider Hindu lenders first. If our hypothesis were correct, we should observe that, on average, Muslim borrowers are less likely than Hindu borrowers to default before Auto Invest is used, because Hindu lenders dig deeper into the Hindu pool to discriminate in favor of choosing Hindu borrowers. By contrast, for Muslim lenders, we should see that Muslim borrowers are more likely to default than Hindu borrowers, because Muslim lenders are digging deeper into the pool of Muslim borrowers to discriminate in their favor when giving out loans autonomously and not according to Auto Invest.

Moreover, our conjecture has implications for which of the two subgroups of borrowers should reduce their likelihood of default after Auto Invest is used. For Hindu lenders, the decrease in defaults should be mainly driven by Hindu borrowers, because Auto Invest makes them choose Hindu borrowers who are, on average, better creditors than the Hindu borrowers they were choosing on their own. At the same time, the likelihood of default of Muslim borrowers did not change. Because Hindu lenders could have dug deeper into the pool of Muslim borrowers to find borrowers who were creditworthy, Hindu lenders were discriminating against such borrowers before. The opposite should be true for Muslim lenders. Muslim lenders' improvement in performance should be mainly driven by the fact that their Muslim borrowers were less likely to default, because Auto Invest should have eliminated Muslim borrowers who were not creditworthy. Instead, the quality of the pool of Hindu borrowers for Muslim lenders should not worsen, because Muslim lenders were discriminating against creditworthy Hindu borrowers.

We bring these hypotheses to the data by estimating versions of the following linear specification in columns (3)-(6) of Table 6:

$$\begin{aligned}
Delinquent\ Loan_{ij} = \alpha + \beta\ Auto\ Invest_j + \gamma\ Muslim\ Borrower_j \\
+ \delta\ Muslim\ Borrower_j \times Auto\ Invest_j \\
+ \theta\ Hindu\ Borrower_j \times Auto\ Invest_j + \zeta\ \boldsymbol{x}_i + \epsilon_{i,j}, \qquad (4)
\end{aligned}$$

In column (3) of Table 6, we find that, on average, the Muslim borrowers are less delinquent in Hindu lenders' portfolios than Hindu borrowers ($\gamma < 0$), which is consistent with Hindu lenders discriminating against Muslim borrowers. Column (4) reveals that the improvement in performance that Hindu lenders enjoyed when moving to Auto Invest was indeed largely driven by a drop in the average delinquency of their Hindu borrower subgroup, whereas the delinquency rates of Muslim borrowers did not change after accessing automated loan issuance relative to before. This pattern is consistent with the hypotheses we laid out above.

Column (5)-(6) of Table 6 repeat this analysis for Muslim lenders. On average, Muslim borrowers were more likely to default (although column (5) shows statistically insignificant coefficients, similar to the baseline effect of Auto Invest), and in column (6), we see that, indeed, the improvement is attributable to Muslim borrowers: After accessing Auto Invest, the likelihood of default of Muslim borrowers drops substantially for Muslim lenders, whereas the performance of the Hindu borrower subgroup barely changes.

6.2.2 Intensive Margin of Performance

Whereas defaults capture the extensive margin of performance—whether borrowers repay their loan in full or not—our data also allow us to study the intensive margin of performance—the share of the overall amount due by the borrower (including principal and interest) that the borrower pays back to the lender. This share is capped at 1 for borrowers who repay their loan in full, because no borrowers pay more than what is due to the lenders. In principle, the share can be as low as 0 if a borrower does not repay anything, but in our data, we find the number of borrowers who repay less than 20% of their amount due is minimal, because the platform automatically expels borrowers who pay less than 20% of the amount due on any of their loans.

We first assess the intensive margin of performance in the raw data. In Panel A of Figure 6, we consider the intensive margin of performance before and after in-group vs. out-group debiasing. The graph to the left plots the cumulative distribution functions (CDFs) for the share of the amount due paid by Hindu borrowers (solid green line) and by Muslim borrowers (orange dashed line) before Hindu lenders use Auto Invest, and hence at a time when Hindu lenders made all their lending choices autonomously. The cumulative distributions display evidence consistent with our conjecture that Hindu lenders might dig deeper into the pool of Hindu borrowers than into the pool of Muslim borrowers. Muslim borrowers who receive loans from Hindu lenders are substantially more likely to repay larger shares of their amount due before the loan servicing is closed. In fact, no Muslim borrowers picked by Hindu lenders pay less than 80% of the amount due. To the contrary, the repayment behavior of Hindu lenders is more volatile: about 20% of them repay less than 40% of the amounts due, and even when considering those who pay at least 80%, the CDF of Hindu borrowers is flatter than that of Muslim borrowers.[25]

The right graph of Figure 6 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. Three patterns are apparent when comparing this graph with the one to the left. First, Hindu borrowers at the bottom of the repayment distribution do not improve substantially. Second, both Hindu and Muslim borrowers at the top of the distribution improve under Auto Invest. Third, because the share of amount paid back is capped at 1, Hindu borrowers improve disproportionally more than Muslim lenders at the very top of the distribution. In particular, 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.

Ultimately, the raw-data analysis does not allow us to conclude definitively whether Hindu and Muslim borrowers' likelihood of repayment changed differently after their lenders moved to Auto Invest. To quantify the relative changes and absorb proxies for borrowers' risk, we need a framework that allows us to test how the share of amount due repaid changed over time and across borrowers in a multivariate setting (D'Acunto and Rossi (ming)). We thus estimate a set of quantile regressions of the following form:

$$Q_\tau(Share\ Repaid_{i,j,t}) = \alpha(\tau) + \beta(\tau)\ Auto\ Invest_{j,t} +\ X'_{i,j,t}\ \zeta(\tau) + \epsilon_{i,j,t}, \qquad (5)$$

whose outcome variable is quantile Q_τ of the distribution of the share of amount due repaid by borrower i to lender j throughout the sample period. All other variables are defined as in equation (4).

To interpret the estimates of equation (5), consider the special case of the median, which is the 50^{th} percentile of the distribution. The coefficient $\hat{\beta}(50)$ estimates that the median share of repaid loans was $\hat{\beta}(50)$ units 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 equation (5) in Panel A of Table 7. Specifically, we estimate

[25]Note also that the share of Hindu borrowers who repay intermediate amounts between 50% and 80% is negligible.

the effects of moving to Auto Invest at the 25^{th} percentile, the median, and the 75^{th} percentile. For the sake of comparison, we also report the results for estimating an OLS specification (which refers to the conditional mean of the distribution). Following D'Acunto and Rossi (ming), we estimate these specifications separately for Hindu borrowers (odd columns) and Muslim borrowers (even columns). Based on the conjecture that Hindu lenders favored Hindu borrowers before moving to Auto Invest, we should find that, throughout the distribution of the percent of loans repaid, the positive effect of moving to Auto Invest is larger for the Hindu borrowers than for the Muslim borrowers of the same lender, because Hindu lenders should have dug deeper into the pool of Hindu borrowers when choosing autonomously. And, indeed, the estimate sizes of the difference between the distribution of the share of loans repaid under Auto Invest are larger for Hindu borrowers than Muslim borrowers across the distribution.

6.3. Performance after Debiasing Stereotypical Discrimination

Overall, the patterns of performance of different groups of borrowers before and after accessing the robo-advising tool align with the conjectures implied by in-group vs. out-group discrimination on the part of Hindu and Muslim lenders. In this section, we assess if the performance of *Shudra* and non-*Shudra* lenders aligns with the conjecture of stereotypical discrimination by lenders.

6.3.1 Extensive Margin of Performance In columns (7)-(9) of Table 6, we consider the extensive margin of performance—whether borrowers default on their loans—by caste. For the case of stereotypical discrimination, our conjecture implies (i) the likelihood of defaults in lenders' portfolios decreases after accessing Auto Invest, (ii) *Shudra* borrowers were less likely to default than borrowers belonging to the other varnas before accessing Auto Invest; and (iii) the performance improvement is mainly driven by a decrease in the delinquencies of the non-*Shudra* subgroup of borrowers, rather than the *Shudra* . Columns (7)-(9) show results consistent with all these implications. First, the likelihood of default drops by 3.3 pp after lenders access Auto Invest. Second, on average, *Shudra* borrowers were about 2.9-pp less likely to default than other borrowers, suggesting lenders were imposing stricter standards on borrowers that are likely to be *Shudra* than they did on other borrowers. Third, most of the performance improvement due to Auto Invest is driven by an improvement of the non-*Shudra* borrower pool rather than of *Shudra* borrowers, because the subgroup of *Shudra* borrowers did not improve after lenders moved to Auto Invest relative to before, whereas non-*Shudra* borrowers picked by the robo-advising tool were 4.3 pp less likely to default than those lenders picked before the lenders moved to Auto Invest.[26]

[26]Statistical significance is sparser for one coefficient of interest, namely, the one attached to *Shudra* borrowers in column (8), but the same coefficient is statistically different from zero in column (9), and we should recall from our previous analysis

6.3.2 Intensive Margin of Performance

Finally, we consider the change in the intensive margin of performance—the share of the payment due by borrowers that is fully paid before the loan is closed—for lenders catering to borrowers with varied recognizability as Shudra and non-Shudra borrowers. We first analyze the raw data in Panel B of Figure 6. The left graph reports the CDFs of the percentage of amounts due repaid by borrowers whose probability of being Shudra is below the median (solid green line) or above the median (orange dashed line) before lenders accessed the robo-advising tool. Throughout the support of the probability of being Shudra, we find borrowers who are likely to be Shudra tend to repay a higher fraction of their amount due, relative to other borrowers.

Once lenders move to Auto Invest in the right graph of Panel B of Figure 6, the distance between the CDFs of the two types of borrowers decreases, and the drop is largely driven by an improvement in unlikely Shudra borrowers' payment behavior—the solid green line shifts to the right. Even for the case of stereotypical discrimination, the results for the pre-Auto Invest period are consistent with lenders imposing higher standards for the selection of likely Shudra borrowers than for the selection of other borrowers. Once the robo-advising tool debiases lenders' stereotypical discrimination, the standards applied to Shudra and non-Shudra borrowers converge.

As we did for the case of in-group vs. out-group discrimination, we propose a multivariate test for the relative change in these intensive-margin distributions before and after Auto Invest, by estimating the following set of quantile regressions as in equation (5). We report the results in Panel B of Table 7. Again, we estimate the effects of moving to Auto Invest at the 25^{th} percentile, the median, the 75^{th} percentile, and the conditional mean based on estimating an OLS specification. We estimate these specifications separately for borrowers whose probability of being Shudra is high (even columns) or low (odd columns). Based on the conjecture that lenders discriminated against Shudra borrowers before moving to Auto Invest, we should find the positive effect of moving to Auto Invest on borrowers' percent of loan repayment is larger for borrowers whose probability of being Shudra is high: lenders of all castes should have dug deeper into the pool of non-Shudra borrowers when choosing autonomously. We confirm that estimate sizes of the difference between the distribution of the share of loans repaid under Auto Invest are larger for borrowers who are unlikely to be Shudra than for other borrowers, even though we cannot always reject the null that the corresponding coefficients for the two groups are the same at standard levels of significance.

that the extent of discrimination against *Shudra* depends on the extent to which their caste is recognizable by lenders; lenders do *not* discriminate against *Shudra* borrowers in cases in which recognizing caste is harder.

7 Conclusions

We propose a unique setting in which to test for and quantify the extent and effects of cultural biases in an economically relevant consumer-finance setting, in which the potential for stereotypical discrimination, which has plagued similar studies in earlier research, is minimal.

We detect evidence that both in-group vs. out-group discrimination and stereotypical discrimination are prevalent and economically sizable in a large-stakes lending-decision setting. 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.

Note our evidence is not definitive about whether the presence of economically-sizable cultural biases is welfare reducing. Biases produce worse financial outcomes for discriminating lenders and misallocation of capital from deserving borrowers to undeserving borrowers in our context. At the same time, though, whether these effects are welfare reducing for lenders depends on whether lenders enjoy negative psychological utility from the act of financing borrowers from a group they discriminate against. If such negative psychological utility were measurable and large enough to overcome the higher consumption utility from selecting better-quality borrowers, imposing debiasing in lending choices would make lenders worse off. Providing theoretical and empirical settings in which to measure and compare these forms of utility, and hence, answer the welfare question, is an important challenge that awaits future research endeavors. Moreover, future research should understand whether setting concrete goals related to loans affects borrowers' repayment behavior, similar to how they affect saving behavior (Gargano and Rossi (2020)).

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).

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.

Bhavnani, R. R. (2014). India National and State Election Dataset.

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.

Breza, E. (2019). Peer effects and loan repayment: Evidence from the krishna default crisis. *Working Paper*.

Brownstein, M. (2015). Implicit bias.

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.

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. (2019a). Identity and choice under risk. *Available at SSRN 3466626*.

D'Acunto, F. (2019b). Tear down this wall street: The effect of anti-market ideology on investment decisions. *Available at SSRN 2705545*.

D'Acunto, F., A. Fuster, and M. Weber (2020). Diverse policy committees are more effective. *Working Paper*.

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., 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'Acunto, F., J. Xie, and J. Yao (2020). Trust and contracts: Empirical evidence. *Working Paper*.

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 inter-group 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*.

Ghassem-Fachandi, P. (2012). *Pogrom in Gujarat: Hindu nationalism and anti-Muslim violence in India*. Princeton University Press.

Goldsmith-Pinkham, P. and K. Shue (2020). The gender gap in housing returns. Technical report, National Bureau of Economic Research.

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.

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*.

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.

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.

Morse, A. (2015). Peer-to-peer crowdfunding: Information and the potential for disruption in consumer lending. *Annual Review of Financial Economics 7*, 463–482.

Morse, A. and K. Pence (2020). Technological innovation and discrimination in household finance. Technical report, National Bureau of Economic Research.

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.

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.

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.

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)

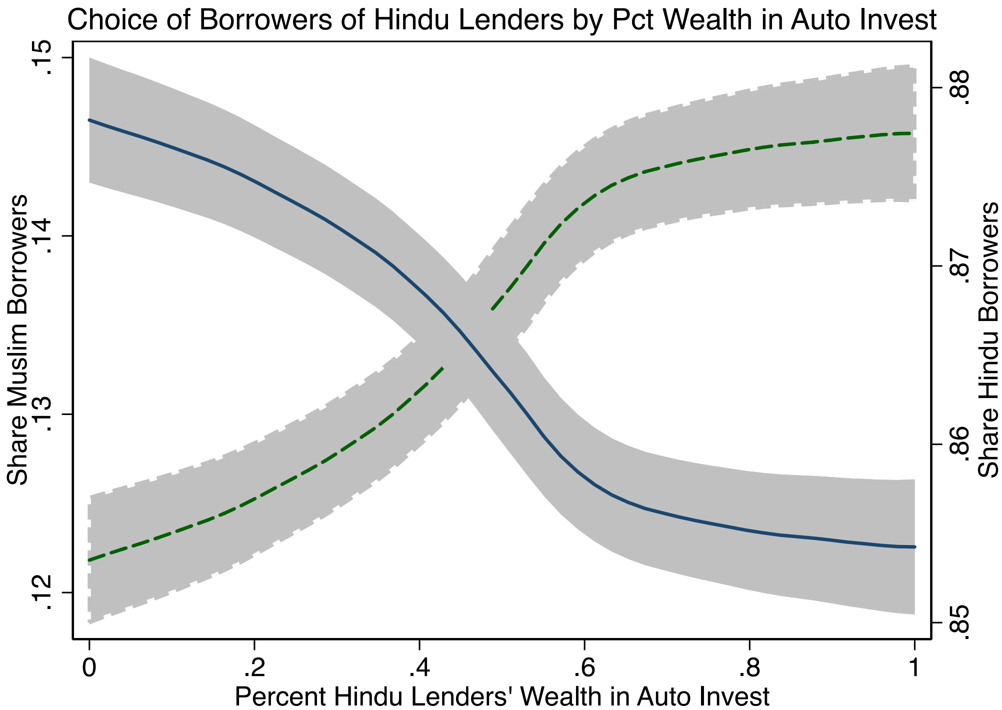

Figure 2 plots the coefficient estimates of kernel-weighted local mean smoothing regressions for whether borrowers are Hindu (blue, solid line, measured on the right y-axis) and whether borrowers are Muslim (green, dashed line, measured on the left y-axis) on the share of their available funds Hindu 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 smooth at 50 points.

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

Panel A. Hindu-Muslim Riots, 1980-2000

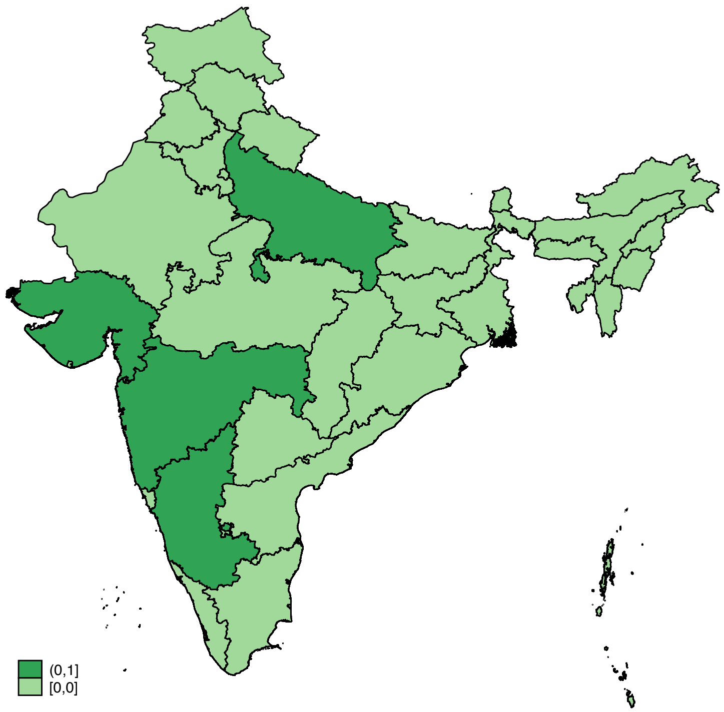

Panel B. Average Vote Shares for the Bharatiya Janata Party (BJP) 1977-2015

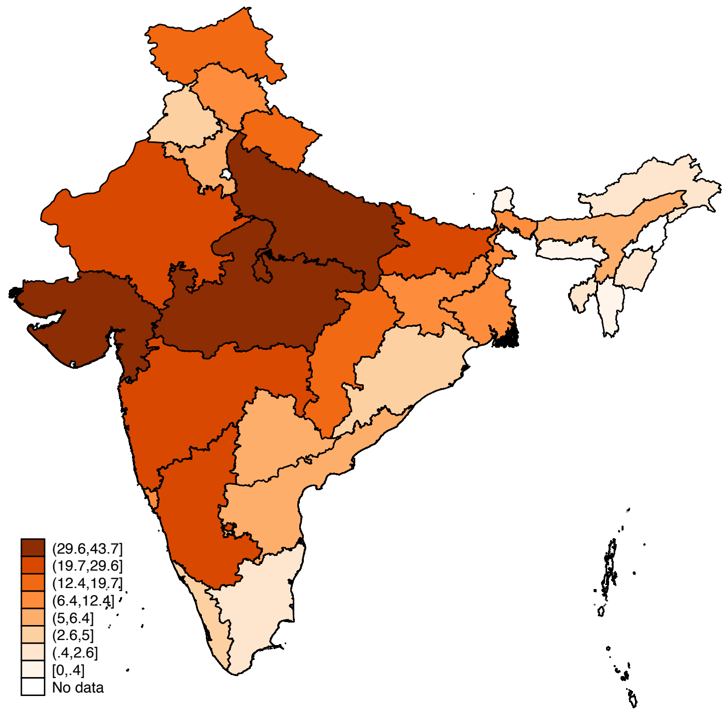

Figure 3 depicts the spatial variation of proxies for the vividness of Hindu-Muslim conflict across Indian states. Panel A compares states in which large-scale riots between Hindus and Muslims and/or pogroms against the Muslim minority happened between 1980 and 2000. Dark green states (Gujarat, Uttar Pradesh, Delhi, Marahashtra, 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: Lending to Discriminated Borrowers:
Shudra Caste Borrowers Before and After Debiasing

Panel A. All Lenders

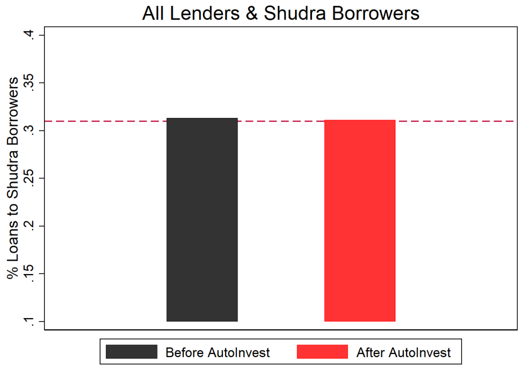

(a) Caste Barely Recognizable (Prob≥0%)

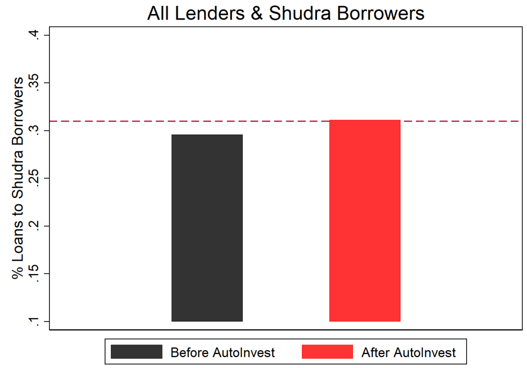

(b) Caste Somewhat Recognizable (Prob>50%)

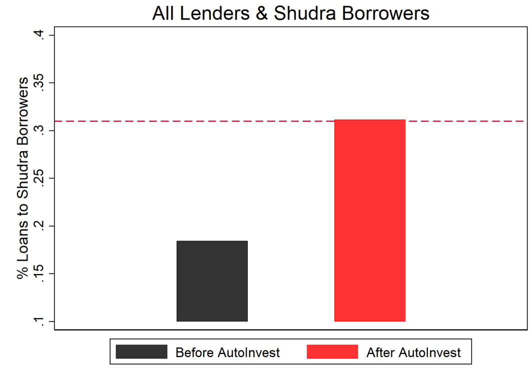

(c) Caste Easily Recognizable (Prob>70%)

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

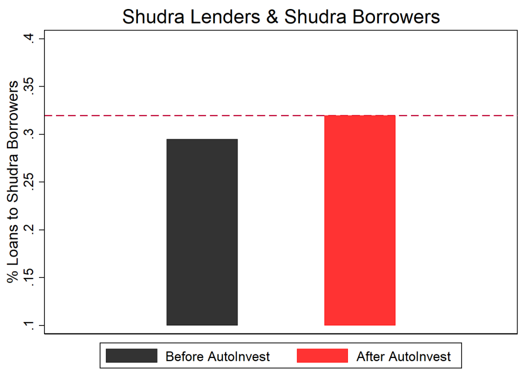

(d) Caste Barely Recognizable (Prob≥0%)

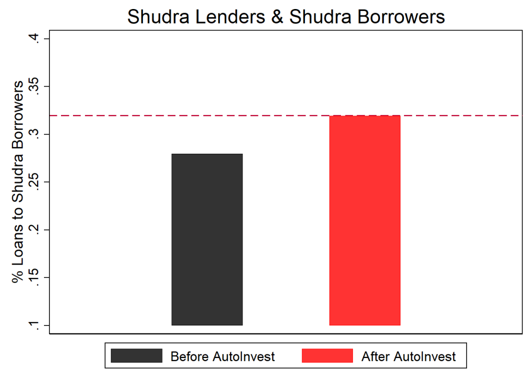

(e) Caste Somewhat Recognizable (Prob>50%)

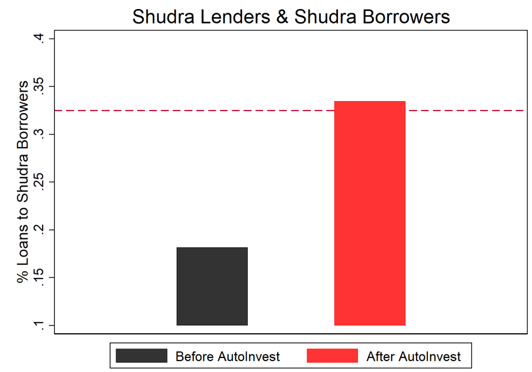

(f) Caste Easily Recognizable (Prob>70%)

Figure 4 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 5: Spatial Heterogeneity of Salience of Stereotypical Discrimination

Crimes Against Scheduled Castes per inhabitant (2018)

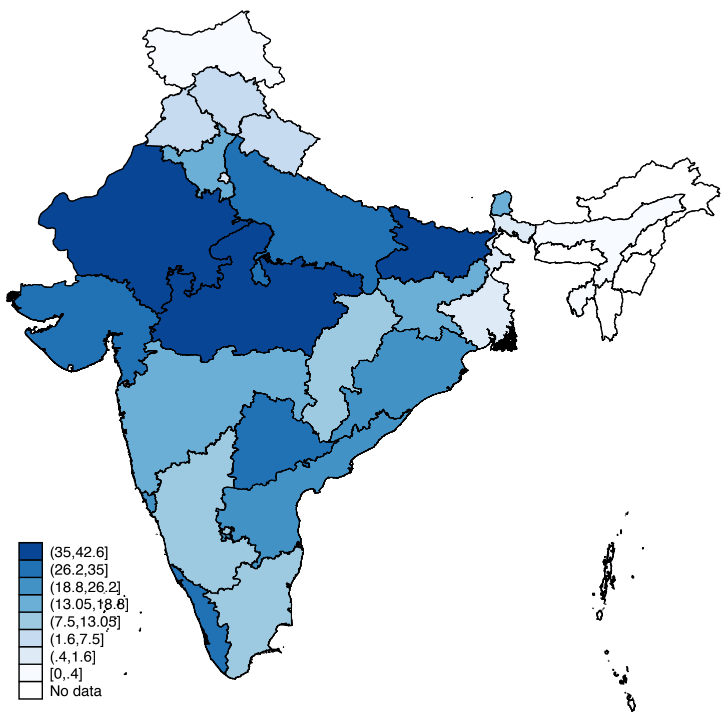

Figure 5 depicts the spatial variation of a proxy for the salience of discrimination against lower castes by Hindus across Indian states, that is, the number of crimes against Scheduled Castes (which includes members of the Shudra varna as well as those belonging to lower castes) per 100,000 inhabitants 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 6: Intensive Margin of Performance Before and After Debiasing

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

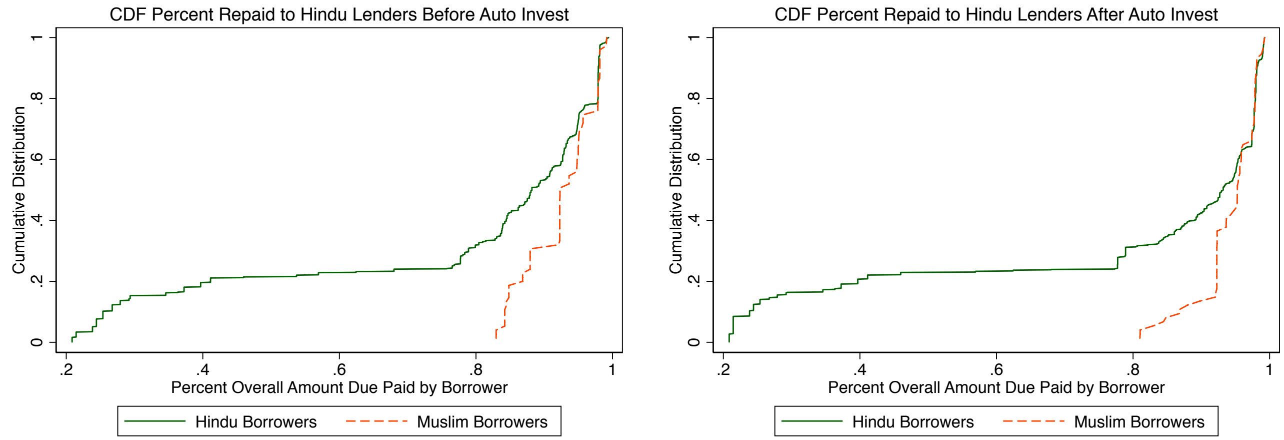

Panel B. Stereotypical Discrimination

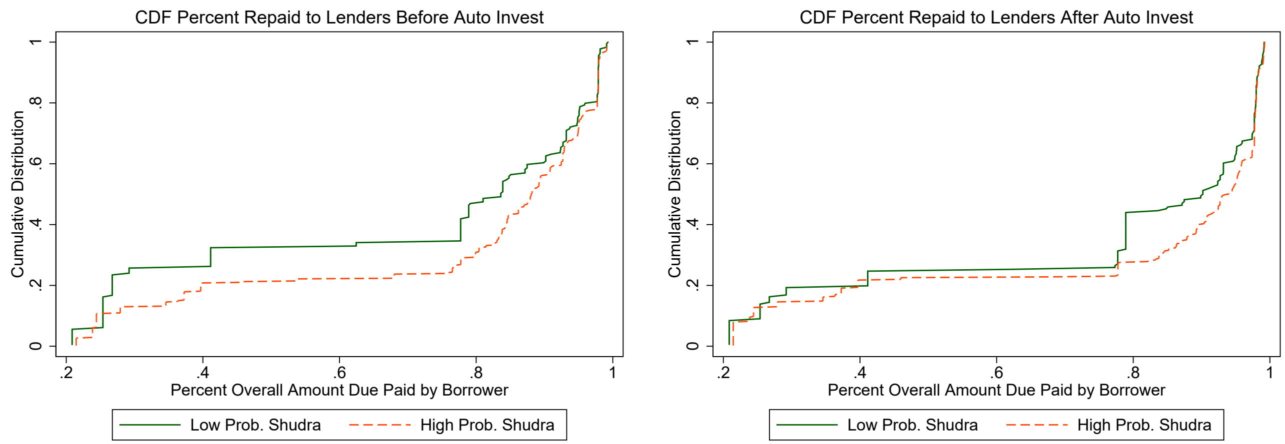

Figure 6 plots a set of cumulative distribution functions (CDFs) for different groups of borrowers over the share of the overall amounts due (principal plus interest) that borrowers repaid by the time their loan account was closed. In all Panels, the left graph refers to the CDFs of borrowers in lenders' 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.

Table 1. Summary Statistics

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

	N. obs.	Mean	St. dev.	25*th* perc.	Median	75*th* perc.
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

	N. obs.	Mean	St. dev.	25*th* perc.	Median	75*th* perc.
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	62,831	0.55	0.50	0.00	1.00	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

Dependent variable: Muslim Borrower	**Baseline**	**Borrower Char.**	**Lender FE**	**Low Use Auto Invest**	**High Use Auto Invest**
	(1)	(2)	(3)	(4)	(5)
Hindu Lender × Auto Invest	0.045***	0.046***	0.044***	0.009	0.053***
	(2.51)	(2.51)	(2.02)	(0.22)	(2.05)
Hindu Lender	-0.058***	-0.058***			
	(-3.52)	(-3.57)			
Auto Invest	-0.026	-0.025	-0.030	0.011	-0.041
	(-1.45)	(-1.40)	(-1.41)	(0.28)	(-1.59)
Tenure Months		0.000***	-0.000	0.001***	-0.000*
		(2.06)	(-0.15)	(2.24)	(-1.73)
Loan Amount		-0.000***	-0.000***	-0.000***	-0.000***
		(-10.73)	(-10.37)	(-6.47)	(-7.56)
Interest Rate		-0.025	-0.032	0.013	-0.066***
		(-1.22)	(-1.57)	(0.48)	(-2.27)
Constant	0.181***	0.194***	0.149***	0.117***	0.170***
	(11.14)	(11.02)	(20.62)	(11.82)	(17.07)
Lender FE			X	X	X
N. obs.	113,284	113,283	113,283	39,366	72,105

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

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

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 $\boldsymbol{x}_i$ 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. Change in Lending to Out-group Borrowers: Salience of Hindu-Muslim Conflict

Dependent variable: Muslim Borrower	**Hindu-Muslim Riots**		**BJP Vote Share**		**Lender Cohort**	
	Low	High	Low	High	Young	Senior
	(1)	(2)	(3)	(4)	(5)	(6)
Hindu Lender × Auto Invest	0.029	0.048***	0.015	0.137***	0.072***	0.004
	(0.75)	(2.62)	(0.88)	(4.05)	(3.19)	(0.18)
Hindu Lender	-0.031	-0.064***	-0.035*	-0.094***	-0.071***	-0.026
	(-1.28)	(-3.86)	(-1.94)	(-7.76)	(-4.37)	(-1.29)
Auto Invest	-0.002	-0.032*	0.005	-0.107***	-0.051***	0.016
	(-0.04)	(-1.79)	(0.28)	(-3.22)	(-2.31)	(0.72)
Tenure Months	0.000	0.000***	0.000***	-0.000	0.000	0.000
	(0.50)	(2.27)	(2.35)	(-0.83)	(1.50)	(1.36)
Loan Amount	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***	-0.000***
	(-7.41)	(-7.89)	(-10.02)	(-2.94)	(-5.53)	(-9.23)
Interest Rate	-0.064	-0.003	-0.004	-0.146***	-0.018	-0.028
	(-1.63)	(-0.13)	(-0.19)	(-2.18)	(-0.54)	(-1.08)
Constant	0.180***	0.193***	0.165***	0.264***	0.205***	0.163***
	(6.42)	(10.67)	(8.60)	(11.02)	(10.61)	(7.33)
N. obs.	46,079	67,204	94,909	15,251	44,689	68,594

Table 3 reports the results of estimating the following specification by ordinary least squares across different subsamples reported on top of each column:

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

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 $\boldsymbol{x}_i$ 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 4. Change in Lending to Discriminated Borrowers: Shudra Caste Borrowers

Dependent variable: Shudra Borrower	**Barely Recognizable** (Prob≥0%)		**Somewhat Recognizable** (Prob>50%)		**Easily Recognizable** (Prob>70%)	
	(1)	(2)	(3)	(4)	(5)	(6)
Auto-Invest	0.003	0.011***	0.007	0.010***	0.020***	0.020***
	(0.54)	(2.54)	(1.54)	(2.30)	(3.69)	(3.82)
Tenure in Months		-0.003***		-0.001***		0.002***
		(-10.06)		(-3.11)		(6.26)
Loan Amount		-0.000***		-0.000***		-0.000***
		(-4.09)		(-6.34)		(-16.10)
Interest Rate		0.102*		-0.037		-0.068
		(1.95)		(-0.92)		(-1.50)
Constant	0.385***	0.439***	0.213***	0.264***	0.177***	0.190***
	(113.36)	(28.11)	(77.75)	(18.80)	(50.66)	(11.76)
N. obs.	62,832	62,831	33,589	33,589	21,239	21,239

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

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

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 $\boldsymbol{x}_i$ 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 5. Change in Lending to Discriminated Borrowers: Salience of Caste Discrimination

Dependent variable: Shudra Borrower	**Inter-caste Hatred** (Prob>70%)		**Delhi and Maharashtra** (Prob>70%)	**Delhi Lenders**		**Maharashtra Lenders**	
	Low	High		Same-State Borrowers	Different-State Borrowers	Same-State Borrowers	Different-State Borrowers
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Auto-Invest	0.013	0.027**	0.022***	0.095***	0.035***	0.023	0.017**
	(1.49)	(2.23)	(2.91)	(3.33)	(2.77)	(1.64)	(2.13)
Tenure in Months	0.003***	0.002*	0.002***	-0.003	-0.003***	-0.001	-0.004***
	(3.81)	(1.84)	(4.59)	(-1.18)	(-2.79)	(-0.39)	(-8.56)
Loan Amount	-0.000***	-0.000***	-0.000***	-0.000***	-0.000	-0.000	-0.000
	(-10.18)	(-5.62)	(-11.02)	(-2.17)	(-0.06)	(-1.22)	(-1.11)
Interest Rate	0.035	-0.222***	-0.070	-0.615***	0.417***	-0.082	0.071
	(0.46)	(-2.08)	(-1.08)	(-2.21)	(2.41)	(-0.63)	(0.87)
Constant	0.165***	0.221***	0.193***	0.600***	0.329***	0.341***	0.480***
	(5.63)	(5.33)	(8.83)	(5.63)	(7.26)	(7.40)	(18.53)
N. obs.	7,327	3,700	10,212	1,187	7,296	3,763	18,565

Table 5 reports the results of estimating the following specification by ordinary least squares across different subsamples reported on top of each column:

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

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 $\boldsymbol{x}_i$ 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 6. Extensive Margin of Performance After Debiasing: In-group vs. Out-group and Stereotypical Discrimination

	In-group vs. Out-group Discrimination						Stereotypical Discrimination		
Dependent variable:	**Lender**		**Lender**		**Lender**		**All Lenders**		
Delinquent Loan	Hindu	Muslim	Hindu	Hindu	Muslim	Muslim		(Prob>70%)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Auto Invest	-0.090***	-0.326***	-0.089***		-0.342***		-0.033*	-0.033*	
	(-5.17)	(-2.99)	(-5.16)		(-3.28)		(-1.72)	(-1.72)	
Hindu Borrower × Auto				-0.094***		-0.330***			
				(-5.38)		(-2.79)			
Muslim Borrower × Auto				-0.048*		-0.527***			
				(-1.79)		(-2.40)			
Muslim Borrower			-0.045***	-0.072***	0.293	0.444***			
			(-4.08)	(-4.02)	(1.57)	(5.63)			
Shudra borrower × Auto									0.019
									(0.42)
Non-Shudra borrower × Auto									-0.043**
									(-2.05)
Shudra Borrower								-0.029	-0.044*
								(-1.35)	(-1.80)
Tenure Months	0.012***	0.007	0.012***	0.012***	0.007	0.007	0.020***	0.020***	0.020***
	(11.82)	(1.32)	(11.83)	(11.84)	(1.40)	(1.38)	(12.86)	(12.84)	(12.87)
Loan Amount	-0.000***	-0.000	-0.000***	-0.000***	-0.000	-0.000	-0.000***	-0.000***	-0.000***
	(-8.63)	(-0.85)	(-8.84)	(-8.83)	(-0.62)	(-0.62)	(-5.81)	(-5.84)	(-5.87)
Interest Rate	1.273***	0.722	1.277***	1.278***	0.961	0.958	1.368***	1.360***	1.360***
	(12.93)	(0.67)	(12.96)	(12.98)	(0.99)	(0.95)	(8.60)	(8.54)	(8.56)
Constant	0.034	0.266	0.038	0.041	0.162	0.157	0.018	0.043	0.046
	(0.71)	(0.68)	(0.80)	(0.86)	(0.45)	(0.41)	(0.32)	(0.74)	(0.80)
N. obs.	16,985	100	16,985	16,985	100	100	3,457	3,457	3,457

Table 6 reports the results of estimating variations of the following specification by ordinary least squares across different subsamples reported on top of each column:

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

where $\textit{Delinquent Loan}_{ij}$ is equal to 1 if the loan associated with borrower i and lender j is closed as delinquent; $\textit{Auto Invest}_j$ is equal to 1 if the lender made the loans after activating Auto Invest and 0 otherwise; $\textit{Muslim Borrower}_{i,j}$ ($\textit{Hindu Borrower}_{i,j}$) is equal to 1 if the borrower i who receives funding from lender j is Muslim (Hindu), and zero otherwise; and $\boldsymbol{x}_i$ 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 7. Intensive Margin of Performance After Debiasing: In-group vs. Out-group and Stereotypical Discrimination

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

Dependent variable:	**Mean**		**25^{th}** *percentile*		**Median**		**75^{th}** *percentile*	
Pct Loan Repaid	Hindu (1)	Muslim (2)	Hindu (3)	Muslim (4)	Hindu (5)	Muslim (6)	Hindu (7)	Muslim (8)
Auto-Invest	0.053***	0.014*	0.062*	0.043***	0.059***	0.030***	0.021***	-0.001
	(4.65)	(1.84)	(1.75)	(2.21)	(8.97)	(3.77)	(9.35)	(-0.27)
Constant	0.772***	0.924***	0.770***	0.879***	0.892***	0.923***	0.957***	0.978***
	(84.62)	(149.22)	(27.14)	(54.72)	(169.92)	(140.74)	(523.90)	(289.89)
R-Square	2189	211	2189	211	2189	211	2189	211

Panel B. Stereotypical Discrimination

Dependent variable:	**Mean**		**25^{th}** *percentile*		**Median**		**75^{th}** *percentile*	
Pct Loan Repaid	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	0.079***	0.068***	0.485***	0.080*	0.088***	0.072***	0.030***	0.025***
	(2.89)	(4.54)	(4.61)	(1.73)	(3.02)	(8.13)	(5.64)	(9.29)
Constant	0.697***	0.771***	0.292***	0.765***	0.836***	0.881***	0.949***	0.954***
	(32.66)	(65.35)	(3.55)	(21.04)	(36.84)	(127.16)	(227.86)	(456.68)
R-Square R-Square	462	1,158	462	1,158	462	1,158	462	1,158

Table 7 reports the results of estimating the following set of quantile regressions:

$$Q_\tau(Share\ Repaid_{i,j,t}) = \alpha(\tau) + \beta(\tau)\ Auto\ Invest_{j,t} +\ X'_{i,j,t}\ \zeta(\tau) + \epsilon_{i,j,t},$$

whose outcome variable is quantile Q_τ of the distribution of the share of amount due repaid by borrower i to 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 $X_{i,j,t}$ 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.

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**

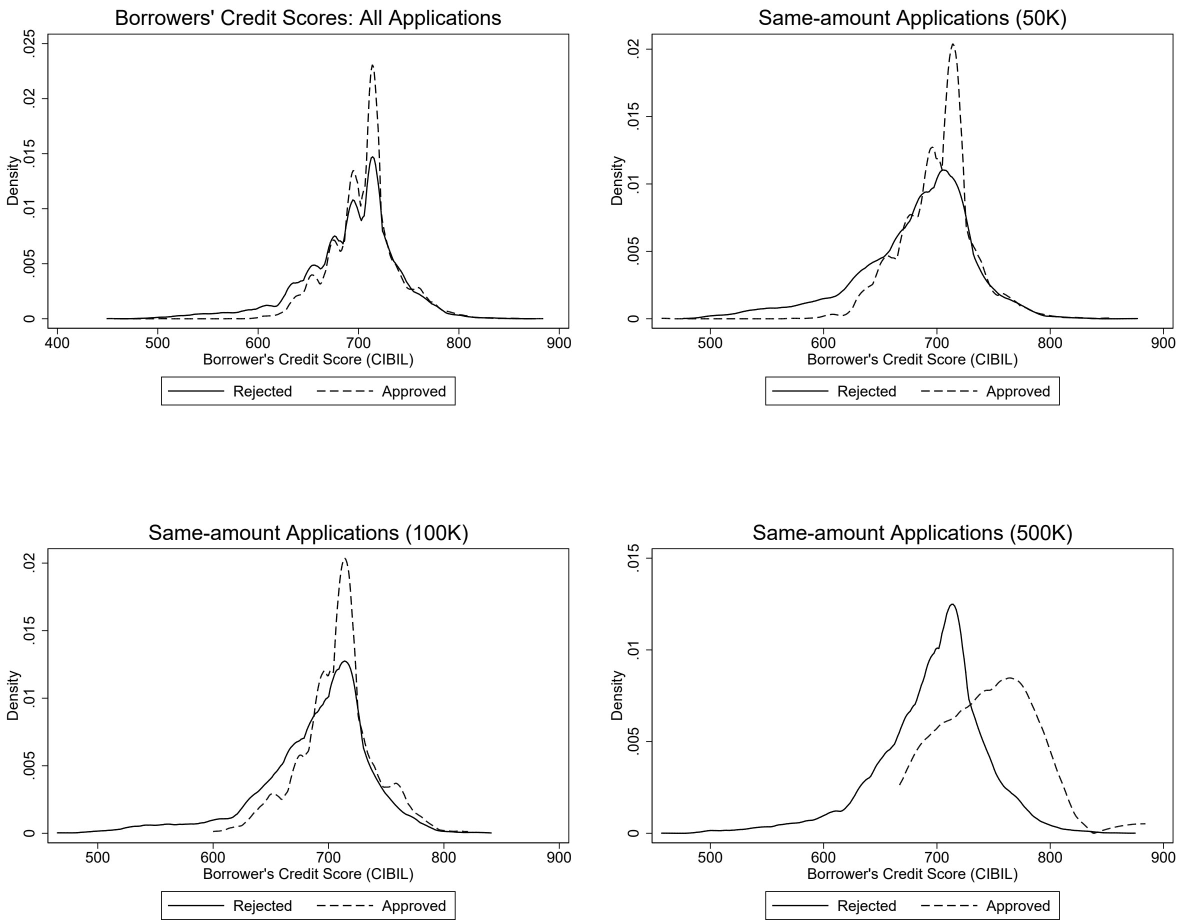

**Figure A.2: Platform's Screening of Borrowers 2:
Interest Rates, Maturities, and Loan Amounts by Credit Score**

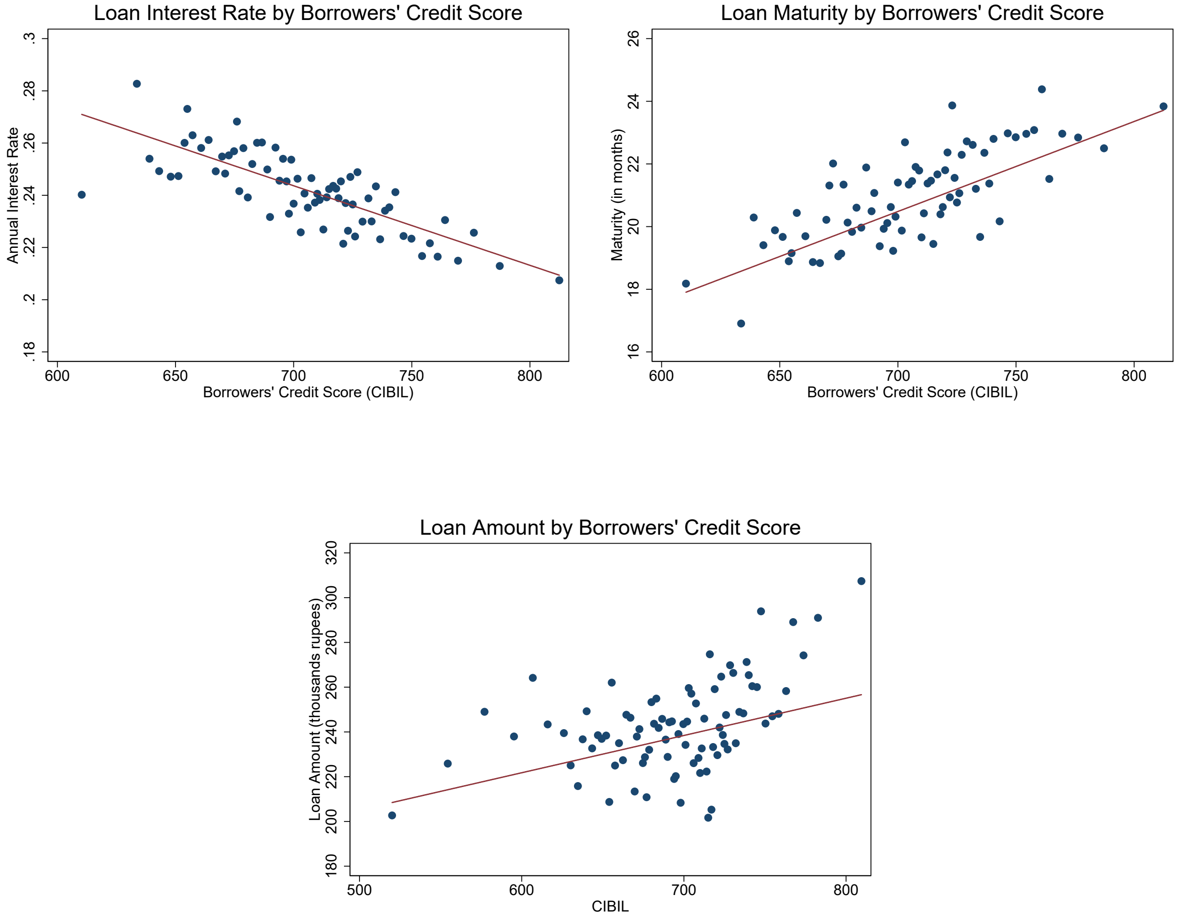

Figure A.3: Distribution of Probabilities that Borrowers are Shudra

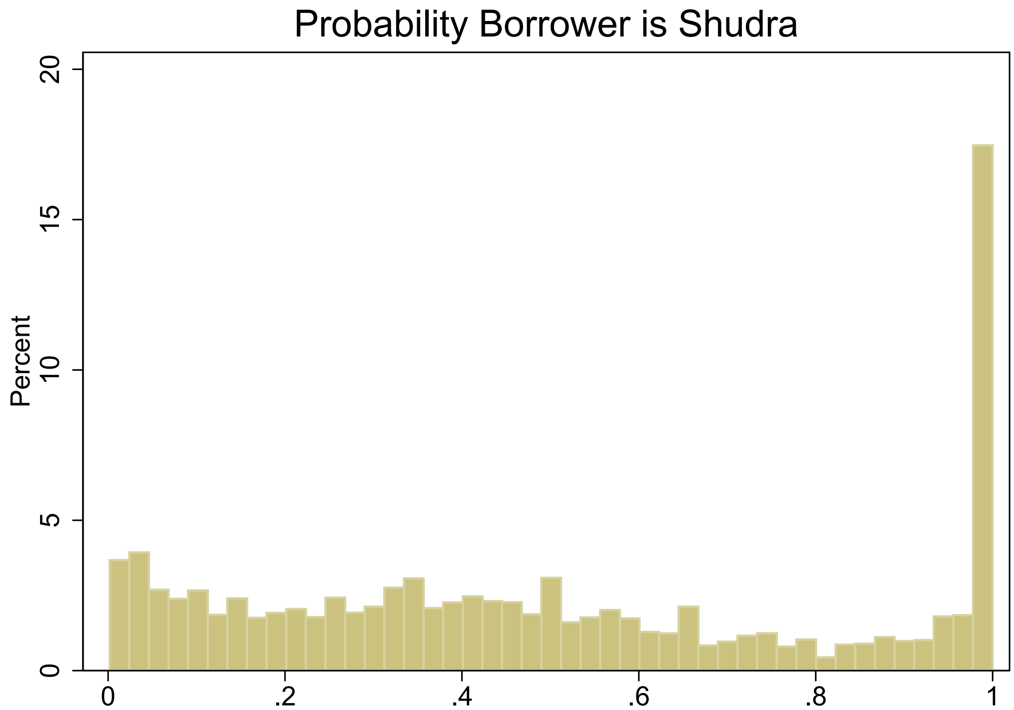

Figure A.4: Intensive Margin Performance Before and After Auto Invest: Full Sample