Exploring the Role of Race and Gender in the Subprime Lending Crisis

Maya Sen[*]
msen@ur.rochester.edu

December 10, 2012

Abstract

The recent subprime mortgage crisis has brought to the forefront the possibility of discriminatory lending on the basis of race or gender. I explore these claims using approximately 10 million observations collected by the federal government in 2006 through the Home Mortgage Disclosure Act. I address two possible theories of discrimination: (1) structural discrimination, which is that any discriminatory lending patterns are picking up the fact that minority borrowers went to different lenders, and (2) individual discrimination, which is the possibility that individual lenders discriminated against identically situated borrowers. The results provide some evidence of both. However, a sensitivity test to examine the effect of missing controls (such as credit score) finds that these racial differences could be explained by a 50% difference in negative credit attributes between blacks and whites under structural discrimination, and 17% difference under individual discrimination.

This research was conducted with support from the John M. Olin Center for Law, Economics, and Business at Harvard Law School as well as from the Real Estate Academic Initiative at Harvard University. I am extremely grateful to Jennifer Hochschild, Gary King, Jim Greiner, Adam Glynn, Matt Blackwell, Richard Nielsen, Omar Wasow, Alisha Holland, Sparsha Saha, Burt Monroe and attendees of the Harvard Graduate Methods and Models Class for ongoing conversations about this project. In addition, I thank various participants at the 2010 Midwest Political Science Association conference panel on race and public

> \[A] lot of us who are older than 30 have some memory of disappointment or humiliation related to banks. The white guy in the suit with the same income gets a loan and you don't? So you turn to local brokers, even if they don't o er the best rates."
>
> { Colvin Grannum, President, Bedford-Stuyvesant Restoration Corporation[1]

On March 12, 2009, the NAACP led a lawsuit against several mortgage lending companies. Their claims were straightforward: they argued that African-American borrowers were up to thirty percent more likely than white borrowers to receive unfavorable lending terms in their mortgage loans. This, they claimed, amounted to nothing more than racial discrimination, and was unlawful under the Fair Housing Act, the Civil Rights Act, and the U.S. Constitution. More recently, the NAACP has been joined in other jurisdictions by various state and local governments, and the Justice Department is now deciding whether to take legal action against Wells Fargo for discriminating against black borrowers. Time

two interrelated explanations:

1. Individual discrimination. Minority borrowers and white borrowers might have gone to the same lenders. Any di erences in subprime lending would therefore suggest that lenders are treating identically situated borrowers di erently, perhaps because of their race or gender.

2. Structural discrimination. In addition, it could be that minority borrowers and white borrowers also sought out di erent lenders { perhaps because of predatory lending practices, or racial di erences in social and professional networks, or some combination of these and other factors. Because minority and white borrowers went to di erent lenders, and because di erent lenders will o er distinct products and services, it is possible that minority borrowers ended up with more subprime loans.

Untangling the e ects of each type of discrimination is di cult, but it is possible to gain traction on the question by leveraging the fact that the federal government routinely collects data at the lending agency level. As such, incorporating speci c agency into the modeling allows us to explore directly whether individual discrimination occurred, while disregarding the identity of the agency allows us to estimate a measure of structural and individual discrimination combined { a holistic measure of discrimination. To this extent, I attempt to applying matching techniques to HMDA in order to examine the racial di erences in pricing, both within and across lenders and lender types.

Although the results presented here suggest some discrepancies on the bias of race and gender in both contexts, a key complication with the real estate context is that many important personal data { including credit scores, employment histories, savings, and debt obligations { are highly proprietary and are not collected by the federal government via the HMDA. Thus, any conclusions based on HMDA data are thus subject to the critique that unobserved variables { not racial or gender discrimination { are driving the results.

To address these concerns, I subject the core ndings to sensitivity analyses. The sensitivity results suggest that evidence of individual-level discrimination are very sensitive missing confounders (e.g., credit scores), with evidence of structural discrimination being less so. Taken in tandem, the results strongly suggest that it is necessary to gather data on individual nancial characteristics like credit scores and debt-to-income ratios before de nitive assertions can be made by legal and policy actors.

This article proceeds as follows. Part 1 will provide an overview of recent research on these issues, noting how this paper reinforces and departs from previous work. Part 2 will discuss the data, which are 10,856,516 applications for home loans that were reported in 2006 to the federal government pursuant to the Home Mortgage Disclosure Act (HMDA). Part 3 and 4 will describe the methodology behind the analysis, noting in particular the identi cation strategies used. Part 5 will present the results, while Part 6 will discuss the sensitivity of conclusions made on the basis of the HMDA to unrecorded variables (such as credit scores or down payments). The article will conclude with a brief discussion of the substantive implications.

1 What We Know About Discrimination and Lending

Theories of discrimination have largely focused on individual attitudes. Does an employer discriminate between black and white job applicants (Fryer Jr and Levitt (2004), Bertrand and Mullainathan (2004))? Are white voters willing to vote for a black politician (Hopkins (2009))? To this extent, models of discrimination have largely focused on individuals, their attitudes, and their prejudices (e.g., Becker (1971), Myrdal and Bok (1996)). Less well measured has been discrimination based on broader, structural considerations. Consider a simple employment example. A woman job applicant might be discriminated against when applying to the same job as a similarly quali ed male applicant. But the same worker could also

height of the housing boom, data on nearly 11 million lending applications were recorded annually in this fashion.[2] For each lending application, the HMDA requires that lenders record data on each applicant's race, ethnicity (i.e., whether he or she is Hispanic or Latino), gender, and income to the nearest thousand.[3] Lenders are also asked to provide information on (1) whether the loan was originated (i.e., the loan was processed and the funds disbursed), (2) how much the loan was for, and (3) any applicable reasons for denial. Perhaps most importantly, lenders must report whether the loan was a high-cost loan { de ned as having an interest rate greater than 3 percent of that o ered by comparable U.S. Treasury instruments (usually 30-year Treasury Bonds). The federal government also requires that lenders report the interest rate on the loan, but *only* if the loan quali es as a high-cost loan.

Although copious (n = 10,856,516), the data have signi cant shortcomings. For example, as many have observed, the de nition of a high-cost loan under the HMDA is not consistent with the colloquial understanding of a \subprime" loan. This is an accurate critique, as the HMDA does not require lenders to report idiosyncratic loan speci cs (length of the loan, down payments, repayment terms, etc.). More importantly, lenders do not need to report \private" information such as applicants' credit scores, employment histories, savings, or debt obligations. Unfortunately, even with identical incomes, property types, and geographical demographics, two applicants with di ering credit scores or debt obligations will nonetheless be o ered di erent loan products { including loans with di erent interest rates. The same is

[2] The HMDA was enacted by Congress for the speci c purpose of detecting discriminatory patterns in the home mortgage industry. Importantly, however, the HMDA covers only those lending agencies with large mortgage portfolios. Pursuant to criteria annually issued by the Federal Reserve Bank, only lenders with home purchasing loans exceeding 10 percent of all loans they have issued (or $25 million, whichever is greater) were obligated to ll out the HMDA paperwork in 2006. This requirement e ectively means that smaller, more specialized lenders can operate without having to disclose the kind of information mandated by the HMDA.

[3] For income, lenders were instructed to \enter the total gross annual income your institution relied on in making the credit decision. For example, if your institution relied on an applicant's salary to compute a debt-to-income ratio, and also relied on the applicant's annual bonus to evaluate creditworthiness, report the salary and the bonus. Report the amount in thousands, rounded to the nearest thousand ($500 should be rounded up to the next thousand)."

true for applicants with di erent levels of savings for down payments or closing costs. This is a key weakness of using the HDMA data, and will be discussed at length below.

3 Identi cation Strategies in the Lending Context

The volume of data allow the use of matching techniques. By dividing observations along a pre-speci ed set of control variables (discussed below), and pairing the observations that have identical values of these variables, matching in this context can e ectively mimic the

audit study, but, rather allowing individuals to go to di erent lenders, and then taking the average *ex post* e ects.

Here, we can again leverage the unique agency identi er in the HMDA data.[4] As discussed above, if we include the unique identi er, then we have controlled for across-agency uctuations (i.e., we have compared only people going to the same lender). The end result would be an estimate of the e ect of race that constitutes impermissible racial discrimination by an individual lender. On the other hand, *excluding* the agency identi er allows us to estimate the e ects of both individual-level discrimination and structural discrimination together. Under this speci cation, a minority borrower could also be compared to non-minority borrowers who went to di erent lenders. If we believe that minority borrowers were cajoled or prompted to seek out predatory lenders { ones more likely to o er toxic loans { then we would see this e ect re ected in these estimates. Important to note is the fact that failing to condition on the unique agency identi er does not neatly disaggregate the two discriminatory mechanisms. What we can do, however, is to compare and contrast the results under both modeling speci cations. Di erences in the two results can help us pinpoint the relative impact of the two discriminatory mechanisms.

4 Matching Methodology

The abundance of data (nearly 11 million observations) allowed for exact matching in all but a handful of states. Unfortunately, the staggering number of observations also meant that matching methods relying on functional form assumptions, such as propensity score matching, were impossible to implement (using even computing clusters and parallel processing). However, there are several signi cant advantages to matching observations exactly.[5] First,

[4] Note, however, that some literature has questioned whether the agency identi er in the HMDA captures meaningful distinctions between di erent lenders { e.g., Laderman and Reid (2008).

[5] The exactFirs(und0)-333(R)]TJc-39inpFiruteng).e[()u [G [(d)-33ccrimi-5429(gram9626 Tf24.47 -3.616267.492)-386(agR

and perhaps most obviously, exact matching is intuitive { we simply pair identical observations together and then see how the lending decisions di er.[6] For example, using exact matching, a borrower recorded as African-American, making $40,000 annually, looking to purchase a single-family home in Cape Coral, Florida, would be matched with a white borrower, also earning $40,000, and wanting to purchase a single-family home in the same neighborhood. Second, because the matching is done exactly, all of the control variables are by de nition balanced { an advantage over, for example, propensity score matching (which relies on a correct speci cation of the propensity score and can make balance worse in some instances (King et al., 2011)). On the other hand, matching exactly is data-exhausting; due to the curse of dimensionality, matching along control variables exactly (i.e., without any coarsening or distilling of any variables) means that many observations will be dropped, leading to concerns that the post-matched population is ill-suited to make generalizations. Here, the abundance of real estate activity at the height of the housing boom means that we have nearly eleven million initial observations (n = 10,856,516). Thus, even though matching exactly prunes large amounts of data (in this case, to be clear, a large fraction of the data), we still have enough data left over to make valuable inferences { some 20,000 observations under the most stringent of matching (see Table 2). To assuage concerns that the matching inadvertently truncates a portion of the population space, however, additional results from OLS models (that use all 10 million observations and do not discard any data) are included in the Appendix; they are substantively similar to the structural-discrimination matching results presented.

The pre-treatment control variables (i.e., those variables matched on) were those taking place *before* the borrower lled out a mortgage application (be it a nal application or an application for pre-approval).[7] The variables matched on included any immutable

not feasible in two low-population states, Wyoming and Montana.

[6] Again, this provides an observational parallel to studies like Ross et al. (2008) and Yinger (1986).

[7] A full description of the variables is given by the HMDA o cial manual, at http://www.ffiec.gov/

attributes of the property, including where it was located (state, county, and census tract number based on the 2000 U.S. Census)[8] as well as the whether the property was owner-occupied as a principal dwelling versus not owner-occupied. Other information considered pre-treatment includes demographic information about the property location, including its location in an metropolitical statistical area (MSA), the percent of the MSA population that is minority (to the tenth of a percent), the MSA's median income (to the nearest thousand), and how many housing units (and family units) the MSA had.[9] The pre-treatment variables also include any borrower attributes una ected by the treatment, including the borrower's income. Here, the income { the only continuous variable { was recorded to the nearest thousand, and the wealth of data allowed for exact matching (to the nearest thousand) on this variable as well.[10]

It is worth mentioning that we may or may not include the unique lending agency ID with the pre-treatment variables. We have reason to think that the identity of the lending agency *is*

	Pre-matching	Structural Matching	Individual Matching
For Blacks			
N Blacks	977,847	211,241	21,238
N Whites	8,063,991	211,241	

subprime status, a binary \yes" or \no" variable, was regressed on the treatment variable.[12] The resulting estimate gives us the Average Treatment E ect (ATE), which in this case can be interpreted as the increased or decreased probability that the borrower in question received a high-cost loan. For purposes of this analysis, note that \high-cost" (or \subprime") refers to any loan that was (1) originated and (b) had an interest rate in excess of 3 percent of a comparable treasury instrument.[13]

Individual Discrimination

Matching on the pre-treatment control variables { including on the unique agency identi er { allows us delineate any possible individual-level discrimination that might be taking place. Starting from 10,856,516 individuals, 21,238 African Americans and whites, 99,249 women, 42,561 Hispanics, and 15,589 Asian Americans were matched exactly to corresponding borrowers from the control populations.[14] Note that these are borrowers who were matched *exactly* along all of the pre-treatment control variables and who, in addition, also led an application for a mortgage loan with the same lenders as their non-minority counterparts. Thus, any remaining e ect of the race or gender variable will, assuming no omitted variables, only be picking up impermissible racial considerations. Table 3 presents these gures for the entire United States, as well as by region.[15] Figure 2 presents the same data, but does so

[12] This was done using an ordinary least squares regression, even though a logit regression might have also been used. Because the treatment variable was the only independent variable, and because the treatment

state by state.

The results are the strongest and most interesting for those borrowers classi ed by the HMDA lenders as African American. Even though they constitute the largest minority in the United States, blacks comprise the second smallest sample in this study (behind Asian Americans); consequently, the con dence bounds for the state-by-state estimates are quite large in comparison to the other groups.[16] (This is especially the case for those African-American borrowers living in low population states such as Maine, Utah, and Connecticut.) Nonetheless, the e ects are quite sizable and, for about half the states, signi cant. For some states { Wisconsin, Mississippi, Indiana, and Nevada { being classi ed by a lender as African-American results in an approximate 10 percent higher probability of being awarded a high-cost loan. Given that we exactly matched on all available nancial and demographic variables and compared only borrowers going to the same lenders, this is quite a strong result.

The estimates are signi cantly more informative when looking at speci c regions of the country and also at the United States as a whole. Indeed, pooling over all of the states gives us a treatment e ect on the African-American variable of 0:575, and it is signi cant at the 1 percent level with 21,238 matched African-American borrowers. Thus, African Americans { *even when going to the same lender and even when displaying the same nancial pro les as white Americans* { are approximately 6 percent more likely than similarly situated white borrowers to receive high-cost loans. Note that African Americans living in the Midwest and Northeast appear more likely to be disadvantaged, while those living in the West and, surprisingly, the South are less so.

For women, the results are striking for precisely the opposite reason. The state-by-state distribution of the female treatment e ect looks approximately normal and centered

[16] It may very well be the case that African Americans as a whole are less likely than other minority groups to submit mortgage applications.

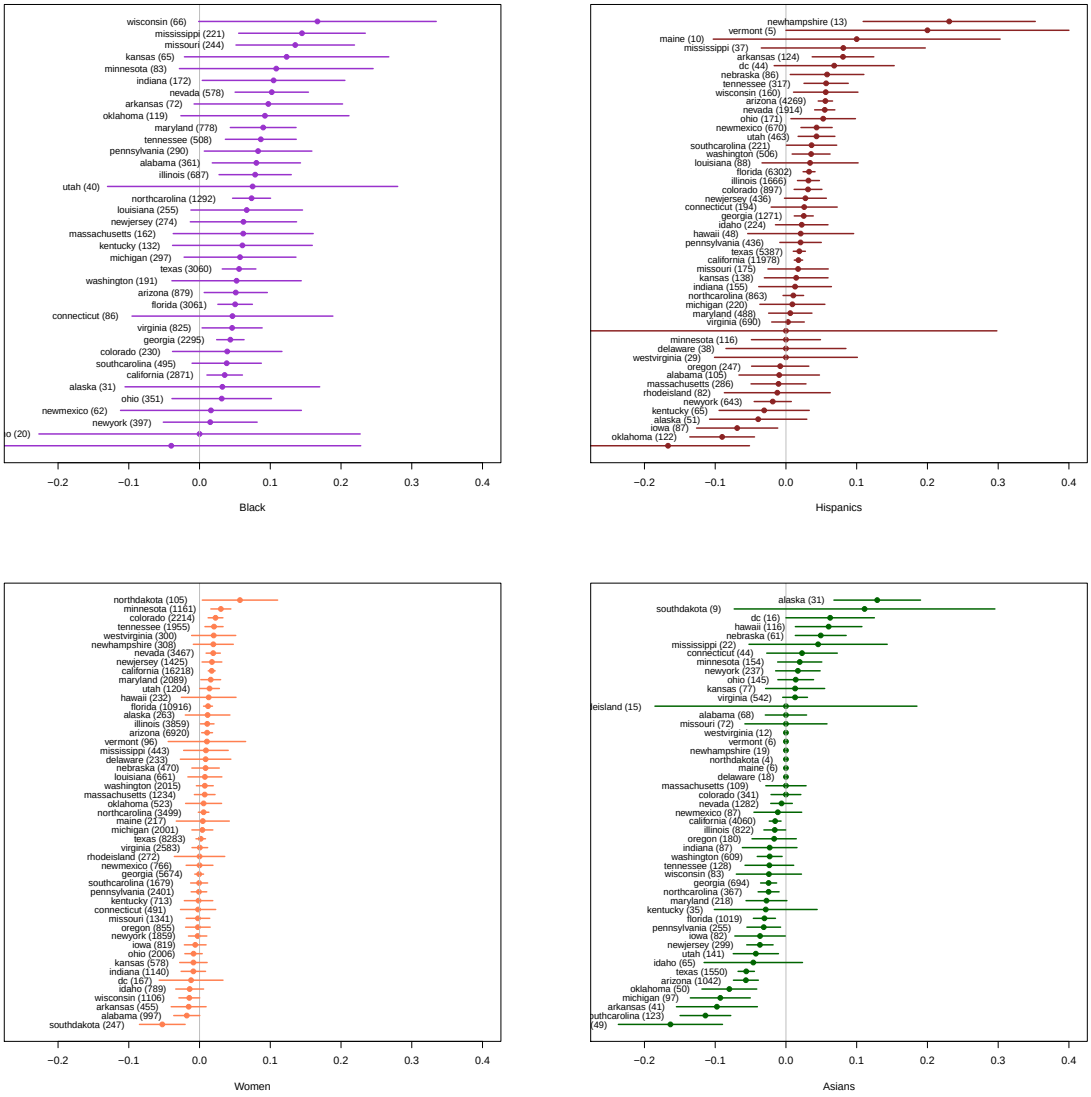

Figure 1: Likelihood of receiving a high-cost loan (individual-level measurements). The solid dots represent the point estimates for the average treatment e ect in each state, while the numbers in parentheses represent the number of treated units in the matching. In substantive terms, the point estimates are estimates of the increased or decreased probability that the treatment group in question will receive a high-cost loan. The line around each point represents the 95% con dence interval. If the spread includes the number zero, then we fail to reject the null hypothesis (at the = 0:05 level) that there is no e ect of a borrower's noted race or gender on receiving a high-cost loan. For the most part, the results suggest positive e ects associated with being African American, Hispanic, and female, but negative e ects for being Asian American.

	US Total		Midwest		South		West		Northeast	
	N	*ATE*	*N*	*ATE*	*N*	*ATE*	*N*	*ATE*	*N*	*ATE*
Blacks	21238	0.0575 (0.00443)	2062	0.0839 (0.0148)	12106	0.0565 (0.00567)	4929	0.0463 (0.00958)	2141	0.0635 (0.0142)

no e ect of being Hispanic on the ultimate lending decision. Similarly, pooling over all of the states gives a coe cient on the Hispanic variable of 0:0256. Although signi cant (with 42,561 Hispanic borrowers matched in this sample), this is not a particularly large e ect. In terms of regional variables, it appears that Hispanics living in the Midwest and Northeast are treated very similarly to other borrowers; the sole exception is the West coast, where being Hispanic is associated with an approximate 3 percent increase in the probability of being awarded a subprime loan.

Perhaps the most surprising results come in the Asian-American category. The hypothesis so far has been that being categorized as a minority would, if anything, lead a borrower to have a *greater* likelihood of being o ered a subprime loan. The data demonstrate that this is actually not the case for Asian Americans living in certain parts of the country. Indeed, with a few notable exceptions { namely Alaska { the results show that being categorized as Asian-American actually *lowers* the probability that a borrower will be awarded a subprime loans. The most dramatic results here come in the more rural areas of the country { Oklahoma, Arkansas, South Carolina { where being categorized as Asian-American results in an approximate 10 percent drop in the likelihood that a borrower will receive a subprime loan. Pooling over all fty states plus Washington D.C. results in a slightly more modest estimate

)-88(eing)-326(a)27(Bayesian)-347(A27ni(est)2(li)1(t)-372(H27ni6

question by the presence of omitted variables, an issue discussed at length below.

Structural Discrimination

Matching Results. For borrowers categorized as African-American, the treatment effects are large, suggesting that African-American borrowers are worse off when they go to different lenders than white borrowers. In all states, being categorized as African American is linked with a higher likelihood of receiving a high-cost loan, even matching exactly for all available financial and demographic information. The effect is particularly striking in a small number of states where a African-American borrower is over 20 percent more likely than an applicant classified white to receive a high-cost loan. (These states include Minnesota, Utah, Mississippi, Wisconsin, South Carolina, Alabama, and Arkansas.) Note also that for all of the states – with the exception of Maine and Idaho[17] – the results are statistically

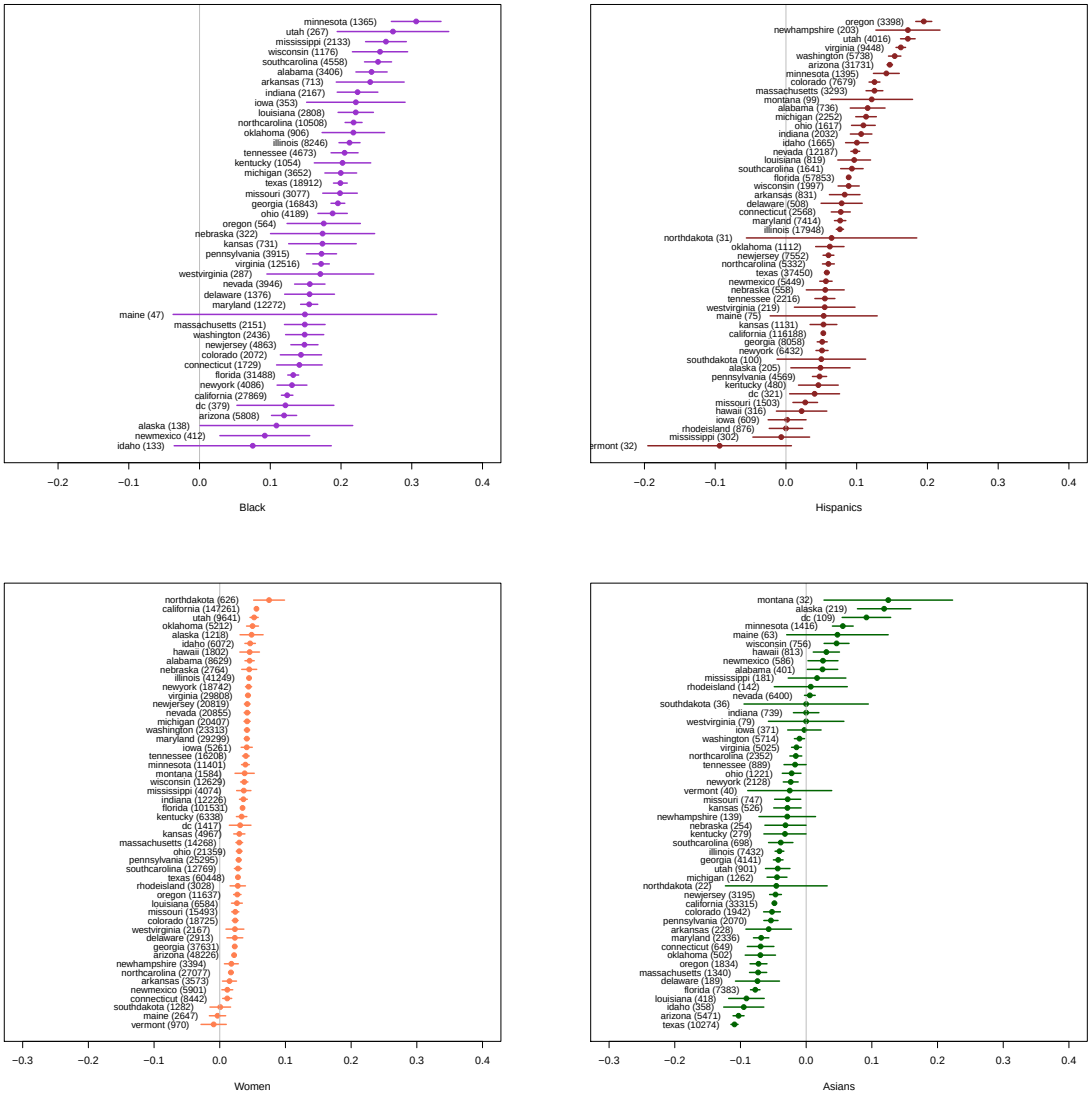

Figure 2: Likelihood of receiving a high-cost loan (structural-level measurements). The solid dots represent the average treatment e ects for each state. In substantive terms, these are estimates of the increased or decreased probability that the treatment group in question will receive a more expensive loan. The line around each point represents the 95 percent con dence interval. If the spread includes the number zero, then we fail to reject the null hypothesis (at the = 0:05 level) that there is no e ect of a borrower's noted race or gender on receiving a high-cost loan. For the most part, the results suggest positive e ects associated with being African American, Hispanic, and female, but negative e ects for being Asian.

	US Total		Midwest		South		West		Northeast	
	N	*ATE*	*N*	*ATE*	*N*	*ATE*	*N*	*ATE*	*N*	*ATE*
Blacks	211241	0.170	25331	0.210	111217	0.184	43716	0.129	30977	0.147
		(0.000)		(0.000)		(0.000)		(0.000)		(0.000)
Women	899182	0.0367	149664	0.0373	323466	0.0313	296235	0.0436	129817	0.0335
		(0.000)		(0.000)		(0.000)		(0.000)		(0.000)
Hispanics	380184	0.0806	31173	0.0812	126818	0.0803	188671	0.08316	33522	0.0672
		(0.000)		(0.000)		(0.000)		(0.000)		(0.000)
Asian Am	117617	-0.0465	14782	-0.0214	32959	-0.0635	57585	-0.0422	12291	-0.0511
		(0.000)		(0.000)		(0.000)		(0.000)		(0.000)

Table 4: Average Treatment Effect (structural-level measurements), subprime status as the outcome variable. The total number of treated individuals is on the left, while the coefficient estimate of the treatment variable is on the right. Note that the sample sizes are approximately ten times the size of the earlier analysis, which reflects the fact that many similarly situated borrowers are in fact going to different lenders. The standard errors are reported in parentheses, although all were extremely close to zero due to the large sample sizes.

in the Midwest and southern regions of the country. Moreover, for all but a handful of states with low Hispanic populations (Alaska, Mississippi, Washington DC, and Iowa) the confidence intervals do *not* include zero, meaning that, sensitivity aside, we can again be fairly certain that there exists some sort of positive treatment effect. Looking at the United States overall, Hispanic borrowers can expect to have 8 percent increased probability of receiving a high-cost loan, even when compared to exactly-situated non-Hispanic borrowers. This effect is, moreover, significant with a p-value of less than 0.01.

While the picture is brighter for women applicants, they are still more disadvantaged than if they had gone to the same lenders as their male counterparts. Unlike women who go to the same lenders as men, those who go to other lenders are more likely to be offered a high-cost loan, with the effect being somewhere around 3-5 percent. (The strongest effects appear to be in California, Utah, Alaska, Oklahoma, and Illinois.) This effect appears to be

relatively stable across regions of the country and stands in contrast to our earlier analysis,

6 Sensitivity Analysis

variable (for example, bad credit scores, insu cient down payments, poor working histories, or oppressive debt obligations) times more frequently than the control group. So, logically, if the treatment were truly randomized, and thus independent of all confounders, then = 1. But if the treatment were applied non-randomly, perhaps due to a missing confounder, then we would expect a value of > 1, such that the treatment group would be more likely to be a ected by the confounder. For example, when = 2, then the confounding variable is twice as likely to be present in the treatment group as opposed to the control group.

To determine the appropriate values, I use the speci cation developed by Rosenbaum (2002) and implemented by Keele (2009).[19] For each possible value of > 1, the sensitivity test recalculates the *p*-values associated with the treatment e ect. It continues increasing hypothetically the value until the *p*-values are no longer statistically signi cant. Thus, once the *p*-value for any given exceeds 0:05 (the maximum *p*-value that constitutes \statistical signi cance" in most social science literature), the sensitivity test \stops" because it has found the level at which the treatment e ect is no longer signi cant. This method therefore provides us with the amount of bias necessary in order for the results to be called into question.

Sensitivity analyses are relatively underused in policy research. As such, there appears to be little agreement as to what constitutes a suitable value for observational studies such as this one. Rosenbaum (2005) provides some guidance, particularly with regard to experimental studies, for which values greater than 3 or 4 appear to be the norm. But for observational studies, more attainable values might be found in Keele (2009), which reproduces the results of the Lalonde job training study. Replicating the treatment e ect estimated via experimental techniques yielded upper bound on possible values of 1.3. This benchmark represents a much lower value than those calculated by Rosenbaum (2005) for

[19] This paper uses a modi ed version of the Rbounds package (Keele (2009)), written for the statistical software R.

	black	women	Hispanic	Asian Am
North Carolina	1.40	-	-	-
Minnesota	-	1.00	0.00	0.00
South Carolina	-	0.00	-	1.70
Alabama	1.00	0.00	0.00	0.00
Oregon	0.00	0.00	0.00	-
Utah	-	-	-	-
Virginia	1.00	-	-	0.00
Texas	1.00	-	1.00	1.30
Colorado	-	1.00	-	0.00
Georgia	1.20	-	1.00	1.00

Table 6: Sensitivity results (individual-level measure of discrimination) for a handful of key states. The dash indicates that the original results were not statistically signi cant at the 5% level. Again, because represents the amount of bias we can tolerate before the treatment e ects lose their statistical signi cance, we would like high values. These values suggest that state-by-state results are actually quite sensitive to omitted variables, with a few exceptions { North Carolina (for African Americans) and South Carolina (for Asian Americans).

much" as the white borrowing population in order for the results to lose signi cance, and South Carolina, where Asian Americans would have to display a characteristic 1.70 as often as white borrowers for the results to be called into question.

The results measuring structural discrimination in addition to individual discrimination are more robust to omitted variables { but still open to the concern that unobserved confounders are playing a role. Consider African-American. For these applicants, certain states { among them Alabama, Arkansas, Georgia, Louisiana, Minnesota, North Carolina, Tennessee, and Utah { have values greater than 1.6. What this means is that an omitted confounder { perhaps poor credit scores, insu cient savings, or lackluster employment histories { would have had to be present in the treatment population at least 1.6 times more often in the African-American borrowing population than in the white borrowing population for the results to lose their signi cance. (With a value of 2.3, North Carolina is particularly striking { African Americans would have to have, for example, poor credit scores at twice

	US Total	Midwest	South	West	Northeast
Blacks	1.52	1.57	1.59	1.34	1.42
Women	1.11	1.11	1.09	1.10	1.10
Hispanics	1.23	1.21	1.22	1.23	1.17
Asian Am	1.22	1.07	1.37	1.16	1.30

Table 7: Sensitivity Results (structural-level measure of discrimination): The cells contain the largest values at which the treatment e ects observed are still statistically signi cant (e.g., have *p*-values less than 0.05). For example, borrowers classi ed by the HMDA as African American would have to have some trait approximately 1.57 as frequently as white borrowers in order for our earlier results to be rendered insigni cant. With the exception of the results for women (which are quite sensitive), the results are in the range of values typically seen in observational studies.

the rates of whites for the results to be a ected.) Unless some relatively large discrepancy exists between similarly earning blacks and whites in terms of terms of credit scores, savings, or employment histories, some structural discrimination probably exists in these regions.

The same cannot be said for the conclusions for the other groups. For female borrowers, all of the results have low values, suggesting that these treatment e ects are sensitive to potential missing confounders. For Hispanics and for Asian Americans, the sensitivity results are not particularly determinative either way and are, for the most part, squarely in the range of values that we would expect to see in observational studies. There are strong exceptions, however. Among these are Virginia, Utah, Oregon, for whom the results for Hispanic borrowers are actually fairly robust (= 1:50 and = 1:60), and Texas, for which the results for Asian Americans are also more robust (= 1:70).

7 Conclusions

	black	women	Hispanic	Asian Am
North Carolina	2.30	1.00	1.10	-
Minnesota	1.90	1.10	1.30	1.10
South Carolina	1.90	1.00	1.20	1.00
Alabama	1.80	1.10	1.20	-
Oregon	1.50	1.00	1.60	1.20
Utah	1.80	1.10	1.50	1.00
Virginia	1.50	1.10	1.50	1.00
Texas	1.50	1.00	1.10	1.70
Colorado	1.40	1.00	1.30	1.10
Georgia	1.80	1.00	1.10	1.20

Table 8: Sensitivity Results in ten interesting states (structural-level measure of discrimination): The cells contain the largest values at which the treatment effects observed are still statistically significant (e.g., have p-values less than 0.05). For example, applicants classified as African American in North Carolina would have to have some trait at least twice as frequently as applicants classified as white in order for the results presented earlier to disappear. The results for several Southern states are relatively robust (by observational study standards) for applicants classified as African American. Results for other groups are more sensitive with the exception of Texas (Asian Americans) and Oregon, Utah, and Virginia (Hispanics). The results displayed here are for the ten most "robust" states for the African-American group.

African-American borrowers were offered high-cost loans at a rate exceeding that of identically situated whites. These results are robust in some instances to omitted variables, suggesting that the distribution of unrecorded data like credit scores or employment histories would have to be somewhat skewed in order to call this conclusion into question. There is also more limited evidence of structural discrimination against borrowers categorized as Hispanic and, to a lesser extent, for women. These results are, however, more sensitive to potentially omitted variables, and, as a whole, are weaker than those for African Americans.

Second, the evidence regarding discrimination at the individual level is actually quite weak, owning mostly to the issue of sensitivity and unobserved financial characteristics. We do see positive effects associated with African American and Hispanic status, particularly in the Midwest and Northeast (for African Americans) and in the South and West (for

Hispanics). These effects are, however, quite sensitive to unobserved variables, and we simply cannot rule out the fact that such variables (including credit scores and accumulated savings) may be driving the results.

Third, Asian Americans represent a startling contrast with the other groups studied here. African Americans, Hispanics, and women are, if anything, disadvantaged by their minority status. Asian Americans, on the other hand, appear to benefit from their racial identity. We see this effect not just when we look at Asian Americans who have gone to different lenders, but also when we narrow the scope to borrowers going to the same agencies. Although these results are in some instances sensitive to potential omitted variables, they do suggest that Asian Americans are not disadvantaged in the lending process and that there is no evidence of individual-level or structural discrimination.

Beyond the substantive conclusions, this research has several broader implications for policy research. First, the paper shows how leveraging matching techniques and different identification strategies can be effectively used to measure different kinds of discrimination quantitatively. Here, it was the identity of the lending agency that provides the key bridge between structural and individual-level discrimination, but similar analogies exist in other areas { for example, in the education context, or in the labor setting, where, respectively, choice of school or employer identity would play a similar role. Second, the paper illustrates how sensitivity techniques can be used alongside the causal inference framework in order to quantify the role of key omitted variables. This is crucial in cases such as these, where necessary information is often withheld due to legal or political concerns.

Appendix

	Estimate	95% CI	Std Err	*p*-value
Not Controlling for Loan Amount				
African Americans	0.1829	0.1819, 0.1839	0.0005	0
Women	0.0411	0.0405, 0.0417	0.0003	0
Hispanics	0.0999	0.0991, 0.1008	0.0004	0
Asian Americans	-0.0558	-0.0572, -0.0544	0.0007	0
Controlling for Loan Amount				
African Americans	0.1792	0.1782, 0.1802	0.0005	0
Women	0.0399	0.0393, 0.0405	0.0003	0
Hispanics	0.1026	0.1017, 0.1034	0.0004	0
Asian Americans	-0.0500	-0.0514, -0.0487	0.0007	0

Table 9: Coe cients from separate OLS regression results testing structural discrimination. All models include controls for (1) loan type, (2) property type, (3) property occupancy, (4) borrower income, (5) MSA population, (4) MSA minority population, (5) median income, (6) MSA number of units, and (8) MSA number of family units. The second set of regressions also include controls for the amount of the loan.

	black	women	Hispanic	Asian
North Carolina	2.30	1.00	1.10	-
Hawaii	1.90	1.00	-	0.00
Minnesota	1.90	1.10	1.30	1.10
South Carolina	1.90	1.00	1.20	1.00
Alabama	1.80	1.10	1.20	0.00
Georgia	1.80	1.00	1.10	1.20
Utah	1.80	1.10	1.50	1.00
Arkansas	1.70	-	1.10	-
Mississippi	1.70	1.00	0.00	0.00
Louisiana	1.60	1.00	1.10	1.20
Tennessee	1.60	1.00	1.00	-
Wisconsin	1.60	1.10	1.10	1.00
Illinois	1.50	1.10	1.10	1.10
Indiana	1.50	1.00	1.20	0.00
Kentucky	1.50	1.00	-	-
Oklahoma	1.50	1.10	1.00	1.10
Oregon	1.50	1.00	1.60	1.20
Tex.1	1.501.0lpo68.1(1)6COregon			

	black	women	Hispanic	Asian
North Carolina	1.40	-	-	-
Georgia	1.20	0.00	1.00	1.00
Tennessee	1.20	-	-	-
Maryland	1.10	-	-	-
Missouri	1.10	0.00	-	0.00
Nevada	1.10	1.00	1.10	-
Alabama	1.00	0.00	0.00	0.00
Arizona	1.00	-	1.10	1.10
California	1.00	1.00	1.00	-
Florida	1.00	-	1.00	1.00
Illinois	1.00	-	1.00	-
Mississippi	1.00	-	-	0.00
Pennsylvania	1.00	0.00	-	-
Texas	1.00	-	1.00	1.30
Virginia	1.00	-	-	0.00
Idaho	0.00	0.00	-	-
Northeast	0.00	-	0.00	*
Oregon	0.00	0.00	0.00	-
Rhineland	0.00	0.00	0.00	0.00
vestigial	0.00	-	0.00	*
Alaska	-	-	0.00	*
Arkansas	-	0.00	-	-
colorado	-	1.00	-	0.00
Connecticut	-	0.00	-	0.00
DC	-	0.00	-	*
Delaware	-	-	0.00	*
Hawaii	-	-	-	0.00
Iowa	-	0.00	0.00	-
Indiana	-	0.00	-	-
Kansas	-	0.00	-	0.00
Kentucky	-	0.00	0.00	-
Louisiana	-	-	-	1.00
Massachusetts	-	-	0.00	0.00
Maine	-	-	-	*
Michigan	-	-	-	1.10
Minnesota	-	1.00	0.00	0.00
Nebraska	-	-	-	1.20
New Jersey	-	-	-	1.00
Mexico	-	0.00	1.00	-
New York	-	0.00	0.00	0.00
Ohio	-	0.00	-	0.00
Oklahoma	-	-	1.10	-
South Carolina	-	0.00	-	1.70
Utah	-	-	-	-
Washington	-	-	-	-
Wisconsin	-	0.00	-	-
New Hampshire	*	-	-	*
South Dakota	*	0.00	1.20	0.00
Vermont	*	-	-	*

Sensitivity results (individual-level measure). Note that the results are ordered by the sensitivity of the African-American treatment variable, from highest to lowest. The asterisk indicates that the sensitivity analysis failed to yield a interpretable value because only members of the treated group received subprime loans. The dash indicates that the original results were not statistically signi cant at the 5% level.

References

Fishbein, A. and P. Woodall. 2006. \Women are Prime Targets for Subprime Lending: Women are Disproportionately Represented in High-Cost Mortgage Market." *Consumer Federation of America, December* . 7

Fryer Jr, R.G. and S.D. Levitt. 2004. \The Causes and Consequences of Distinctively Black Names." *Quarterly Journal of Economics* 119(3):767{805. 4

Haughwout, A., C. Mayer and J. Tracy. 2009. \The Impacy of Race, Ethnicity, and Gender on the Cost of Borrowing." *Federal Reserve Bank of New York Sta Report* (368). 6

Hopkins, D.J. 2009. \No more wilder e ect, never a whitman e ect: when and why polls mislead about black and female candidates." *The Journal of Politics* 71(03):769{781. 4

Ross, S.L., M.A. Turner, E. Godfrey and R.R. Smith. 2008. \Mortgage lending in Chicago and Los Angeles: A paired testing study of the pre-application process." *Journal of Urban Economics* 63(3):902{919. 6, 9, 10, 12

Smith, S.L. and C. Cloud. 1996. \The role of private, nonpro t fair housing enforcement organizations in lending testing." *Mortgage Lending, Racial Discrimination, and Federal Policy* pp. 589{610. 7

Taylor, J., J. Silver and D. Berenbaum. 2004. \The Targets of Predatory and Discriminatory Lending: Who Are They and Where Do They Live?" *Why The Poor Pay More: How To Stop Predatory Lending* p. 25. 6

Williams, R., R. Nesiba and E.D. McConnell. 2005. \The changing face of inequality in home mortgage lending." *Social Problems* 52(2):181{208. 6

Wyly, E.K. and S.R. Holloway. 2002. \The Disappearance of Race in Mortgage Lending*." *Economic Geography* 78(2):129{169. 6

Yinger, J. 1986. \Measuring racial discrimination with fair housing audits: Caught in the act."