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 policy, the 2010 Penn State New Faces in Political Methodology Conference, and the 2009 Society for Political Methodology conference poster session for helpful comments and suggestions.

†Assistant Professor, Department of Political Science, University of Rochester, Harkness Hall 322, Rochester, NY 14627.

‡mayasen.org.

> "[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 offer the best rates."
>
> – Colvin Grannum, President, Bedford-Stuyvesant Restoration Corporation[1]

On March 12, 2009, the NAACP filed 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 will tell whether these lawsuits will be successful, but one thing is certain: these actors are hardly alone in thinking that minorities were disproportionately (and perhaps intentionally) affected by the subprime mortgage lending crisis. Allegations that companies targeted vulnerable populations – recent immigrants, non-English speakers, single mothers, African Americans, etc. – are abundant, and effectively borne out by the swaths of foreclosures currently gripping minority neighborhoods.

To explore these claims, I rely on some 11 million observations gathered in 2006 pursuant to the Home Mortgage Disclosure Act (HMDA). The paper departs from the existing literature on discrimination in lending in two important ways. First, while some observational studies have found correlations between subprime lending and race and gender, I take a causal approach by applying exact matching techniques within the context of the Neyman-Rubin potential outcomes framework. Second, this causal approach allows me to explore the kinds of discriminatory mechanisms in effect. When it comes to subprime lending, there are

[1] Quoted in Powell and Roberts (2009).

two interrelated explanations:

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

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

Untangling the effects of each type of discrimination is difficult, 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 specific 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 differences 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 findings 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 financial characteristics like credit scores and debt-to-income ratios before definitive 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 identification 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 qualified male applicant. But the same worker could also choose to apply only to companies that actively recruit women candidates like her. Thus, she could be discriminated against not just once, but twice. Both actions constitute discrimination colloquially understood, but both have different causes, consequences, and legal and policy solutions. Scholars have understood this conceptual distinction in areas as such as education, labor relations, and law, but measuring and distinguishing these discriminatory mechanisms quantitatively has proved elusive.

The personal finance industry is another area where individual-level and structural discrimination are widely thought to exist. Strong anecdotal evidence points to structural discrimination (particularly in the form of predatory lending – e.g., Kirchoff (2005), Blanton (2007), Powell and Roberts (2009)), while others believe that discrimination at the individual level might have played a significant role (e.g., the NAACP). There is no question that distinguishing between the two discriminatory mechanisms is important. Legal analysts and lawyers will care a great deal more about individual discrimination, as it forms a basis for potentially successful claims under anti-discrimination law (including under the U.S. Constitution). On the other hand, politicians, community advocates, and local officials might care just as much (if not more) about instances of predatory lending or firms targetting certain groups. Indeed, individual discrimination might have an easy redress in the courts, but structural discrimination might be more pernicious and require greater levels of regulatory oversight and governmental involvement.

Teasing apart the potential role that structural and individual discrimination play in any given issue area is, however, quite difficult. The empirical research into the home-buying industry has, for the most part, muddled the issue, pointing to a strong relationship between minority status and subprime lending without disaggregating potential discriminatory mechanisms. For example, a number of studies have relied on survey data, which allows researchers access to sensitive private information — for example, credit scores, employment histories, etc. Nonetheless, this kind of data do not allow researchers to distinguish between structural mechanisms versus discrimination happening at the individual level. Barr, Dokko and Keys (2007), for example, examine data from low- to moderate-income communities around the Detroit area, finding that, even within similarly situated low-income neighborhoods, African-American borrowers were more likely to report having loans with punitive or high-cost terms – such as ballooning interest rates or prepayment penalties. Other studies relying on survey data include Lax et al. (2004) and Courchane, Surette and Zorn (2004), both of which find that minority status is positively correlated with receiving a subprime loan. Other studies have used the federally collected HMDA data have reached similar conclusions (Wyly and Holloway, 2002).

Another subset of studies have relied on proprietary data (including credit scores and employment histories) that are not collected by the federal government. Perhaps the most far-reaching is Bocian, Ernst and Li (2008), which combined 2004 HMDA data with a proprietary database including credit scores, resulting in a database of approximately 120,000 observations. Using logit and least square regressions on this combined database, the authors find that being African American and Hispanic is closely correlated with a higher likelihood of receiving a subprime loan. The authors do not, however, explore whether this provides evidence for individual-level discrimination or if it is a "by-product" of different borrowers going to different lenders. Other studies have looked at similar proprietary data and have likewise have found correlations between concentrations of minorities and subprime lending activity (Williams, Nesiba and McConnell, 2005; Calem, Gillen and Wachter, 2004; Taylor, Silver and Berenbaum, 2004; Mayer and Pence, 2008). A strong exception to this line of research is, however, Haughwout, Mayer and Tracy (2009), which finds no evidence of any kind of discrimination against minority borrowers.

A related literature has taken a pair-testing approach to racial biases – an approach similar in many respects to the one used here. For example, Ross et al. (2008) examine Chicago- and LA-area testers who were assigned fictitious financial profiles. Minority testers were then paired with nearly identical non-minority (white) testers and sent to a representative sample of mortgage lenders. The key finding was that black and Hispanic testers were routinely treated differently than the non-minority testers. Other studies that have taken a similar paired-testing approach have been Ondrich, Ross and Yinger (2003), Yinger (1986), Smith and Cloud (1996). All have found that minority applicants are treated differently in the mortgage process, and all have made this conclusion on the basis of individual discrimination happening at the lender level. The similarities between these studies and the present one are discussed in greater depth below.

Also worth mentioning are the large number of unpublished studies and reports commissioned by non-profit and community advocacy organizations. California Reinvestment Coalition et al. (Unpublished 2009), for example, looked at subprime lending in minority communities and concluded that subprime lending is by far more prevalent in areas covered by the Community Reinvestment Act (CRA), a federal law encouraging lending to minorities. Other studies along these lines include Bocian, Ernst and Li (2006), Apgar and Calder (2005), and Bradford and Associates (2002), which find that African Americans (and in some instances Hispanics) are more likely than whites to receive subprime loans. By contrast, very few studies have examined the potential role of gender in the awarding of subprime or high-cost loans. Those that have, e.g., Fishbein and Woodall (2006), have found that women are overrepresented in the pool of subprime borrowers. In all of these studies, however, inferences about structural versus individual-level discrimination are not fully explored.

2 Data on the Lending Process

This paper uses data collected pursuant to the HMDA in 2006, a year chosen primarily because of availability and tractability, and also because it represents one of the years associated with the housing boom, which lasted from approximately 2004 to 2006. During the 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 – defined as having an interest rate greater than 3 percent of that offered 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 qualifies as a high-cost loan.

Although copious (n = 10,856,516), the data have significant shortcomings. For example, as many have observed, the definition 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 specifics (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 differing credit scores or debt obligations will nonetheless be offered different loan products – including loans with different interest rates. The same is

[2] The HMDA was enacted by Congress for the specific 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 fill out the HMDA paperwork in 2006. This requirement effectively 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 different 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 Identification Strategies in the Lending Context

The volume of data allow the use of matching techniques. By dividing observations along a pre-specified set of control variables (discussed below), and pairing the observations that have identical values of these variables, matching in this context can effectively mimic the ideal randomized experiment – one in which we would have an African American borrower compared with an identical white borrower. Once we have matched the data, we assume that the only difference between the treatment and control groups is that a "treatment" has been applied to one group and not the other. (Similar results are obtained form parametric regression, and are reported in the Appendix.)

This research design allows us to approximate audit studies on racial discrimination (Ross et al., 2008; Yinger, 1986). For example, in the context of employment discrimination Bertrand and Mullainathan (2004) randomly sent identical resumes with either African American names or white American names to employers, and the eventual employment decision (in that case, invitations for interviews) were compared. The "treatment" in these sorts of studies are the names on the resumes, while the experimental units are the potential employers. Here, the framework is the same, even though the observational context and the data are different. The "treatments" are the applicants – who, like the resumes, are identical on the available information. Unfortunately, using the HMDA data means we do not have access to key pieces of information such as creditor scores and down payments – information that we would need to make the applicants truly simile. But this research design, combined with sensitivity analyses on the potential role of these omitted variables (discussed below), takes us far in assessing the relationship of race and gender to subprime lending. Note also that this is an observational parallel to the experimental approaches of Ross et al. (2008) and by Yinger (1986).

The identity of the lender plays a key role in this analysis. If we believe that the predatory story is true, or if we simply believe that different people go to different lenders, then the lending outcome might be heavily influenced by a borrower's choice of bank or mortgage broker. Specifically, minority borrowers might be more inclined to go to certain lending agencies, and these lending agencies might be the ones more likely to offer punitive high-cost loans. To estimate the role that *individual* discrimination might play in lending outcomes, we must therefore condition on the identity of the lending agency itself, using it as a variable to match on (discussed below). Fortunately, the federal government records for each lending agency a unique identifier. This makes it possible to determine whether a borrower filed his or her application with lenders such as Countrywide, Wells Fargo, Washington Mutual, and a variety of regional agencies. Conditioning on this "agency ID" variable allows us to control for across-agency effects as well as to pinpoint discrimination operating at the lender level. The net result is that minority borrowers are compared only to non-minority borrowers who sought the services of the same lending agency – thereby allowing us to estimate the illegitimate reliance by individual lenders on the applicant's race or gender.

Structural Discrimination. The methodology discussion so far has focused on determining the effects of individual discrimination. Just as individual discrimination might have resulted in inequalities in lending between minority and non-minority borrowers, it is possible – and, indeed, interesting – that minority borrowers might have been the targets of structural discrimination as well (for example, in the form of predatory lending). Blacks, Hispanics, and Asian Americans might have simply gone to different lenders, and differences in lender business practices (good or bad) might be driving any observed differences. To this extent, the experimental analogy is different: we are no longer approximating as closely an audit study, but, rather allowing individuals to go to different lenders, and then taking the average *ex post* effects.

Here, we can again leverage the unique agency identifier in the HMDA data.[4] As discussed above, if we include the unique identifier, then we have controlled for across-agency fluctuations (i.e., we have compared only people going to the same lender). The end result would be an estimate of the effect of race that constitutes impermissible racial discrimination by an individual lender. On the other hand, *excluding* the agency identifier allows us to estimate the effects of both individual-level discrimination and structural discrimination together. Under this specification, a minority borrower could also be compared to non-minority borrowers who went to different lenders. If we believe that minority borrowers were cajoled or prompted to seek out predatory lenders – ones more likely to offer toxic loans – then we would see this effect reflected in these estimates. Important to note is the fact that failing to condition on the unique agency identifier does not neatly disaggregate the two discriminatory mechanisms. What we can do, however, is to compare and contrast the results under both modeling specifications. Differences 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 significant advantages to matching observations exactly.[5] First,

[4]Note, however, that some literature has questioned whether the agency identifier in the HMDA captures meaningful distinctions between different lenders – e.g., Laderman and Reid (2008).

[5]The exact matching was done with the statistical program `R` using code similar to that used by the Coarsened Exact Matching package (`CEM`), described in Iacus, King and Porro (2009). Exact matching was

and perhaps most obviously, exact matching is intuitive – we simply pair identical observations together and then see how the lending decisions differ.[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 $40,000, and wanting to purchase a single-family home in the same neighborhood. Second, because the matching is done exactly, all of the control variables are by definition balanced – an advantage over, for example, propensity score matching (which relies on a correct specification of the propensity score and can make balance worse in some instances (King et al., 2011)). On the other hand, matching exactly is data-exhausting; due to the curse of dimensionality, matching along control variables exactly (i.e., without any coarsening or distilling of any variables) means that many observations will be dropped, leading to concerns that the post-matched population is ill-suited to make generalizations. Here, the abundance of real estate activity at the height of the housing boom means that we have nearly eleven million initial observations ($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 filled out a mortgage application (be it a final 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 official 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 unaffected 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-treatment – after all, borrowers decide on a lender before they begin filling out the paperwork for a mortgage loan. On the other hand, if we believe that borrowers "sort" themselves and that minority borrowers are more likely to go to certain lenders, the identity of the lending agency is not necessarily pre-treatment. Thus, I at times do and do not match or control for the agency ID variable; this allows me to explore the differences between structural and individual discrimination outlined above.

The **post-treatment variables** (i.e., those variables not matched on) are those data recorded *after* the lender has been exposed to the applicant's race or gender – in other

hmda/pdf/2006guide.pdf.

[8]The fine-grained nature of the data allowed matching to the census tract level; matching on the census tract presumably controls for regional and neighborhood variation, including differences across communities in terms of average credit scores, employment levels, and disposable income.

[9]Any MSA-specific attributes are matched on automatically when matching on the MSA.

[10]On this point, the data are "coarsened" automatically by the fact that income is recorded only to the nearest thousand. Matching exactly on this variable without further coarsening was possible thanks to the wealth of data initially available.

Pre-treatment Control Variables	Census tract (based on 2000 U.S. Census), owner occupancy of property (owner occupied, non-owner occupied), loan purpose (home purchase, home improvement, or refinancing), county of property, applicant income to the nearest thousand, MSA, MSA population, MSA minority population, MSA median income, MSA total number of units, and MSA number of family units
Post-treatment Variables	Reasons for denial (if any), the amount of loan, interest rate, "subprime" (high-cost) status, and whether the loan was subsequently purchased by another agency

Table 1: Pre- and post-treatment variables. The unique lending agency identifier is at times also treated as a pre-treatment control. Only pre-treatment variables were used in the matching.

words, those occurring after the lender has been "treated." The post-treatment variables therefore include anything having to do with the loan terms: whether the applicant was rejected or accepted, the terms of the loan, how much the loan was for, whether it was a high-cost loan, and whether the loan was eventually sold to another lender. (It is possible that the loan amount is decided in tandem with the rate; this possibility is explored via OLS regression, included in the Appendix. The substantive conclusions are unaffected.) Including these variables into the model could introduce post-treatment bias, as they have a good chance of affected by the treatment itself.

An important consideration is that the post-matched sample might not reflect accurately the population at large, or even the pre-matched borrowing population. For example, it could be the case that matching blacks to whites drops poorer blacks and impugns the generalizability of the results, in the process underestimating possible discrimination against poorer blacks. (In general, as Table 2 suggests, the population of borrowers is quite a bit wealthier than the general population; blacks have a median income of \$66,000, whereas census estimates put that figure for the general population closer to \$30,000.) OLS regressions provided in the Appendix assuage these concerns (and provide substantively similar

	Pre-matching	Structural Matching	Individual Matching
For Blacks			
N Blacks	977,847	211,241	21,238
N Whites	8,063,991	211,241	21,238
Mean Income Blacks	$81,489.98	$73,003.82	$71,473.77
Median Income Blacks	$66,000	$65,000	$62,000
Mean Income Whites	$100,268.33	$73,003.82	$71,473.77
Median Income Whites	$76,000	$65,000	$62,000
% Blacks buying owner occupied	0.89	0.97	0.97
% Whites buying owner occupied	0.88	0.97	0.97
For Women			
N Women	3,312,150	899,182	99,249
N Men	7,544,366	899,182	99,249
Mean Income Women	$84,840.68	$75,842.01	$76,892.08
Median Income Women	$66,000	$66,000	$66,00
Mean Income Men	$107,135.54	$75,842.01	$76,892.08
Median Income Men	$81,000	$66,00	$66,000
% Women buying owner occupied	0.88	0.97	0.96
% Men buying owner occupied	0.91	0.97	0.96
For Hispanics			
N Hispanics	1,334,738	380,184	42,561
N non-Hispanics	9,521,778	380,184	42,561
Mean Income Hispanics	$94,264.76	$80,850.19	$80,887.80
Median Income Hispanics	$78,000	$72,000	$72,000
Mean Income non-Hispanics	$101,184.46	$80,850.19	$80,887.80
Median Income non-Hispanics	$76,000	$72,000	$72,000
% Hispanics buying owner occupied	0.88	0.97	0.96
% non-Hispanics buying owner occupied	0.91	0.97	0.96
For Asian Americans			
N Asians	406,998	117,617	15,589
N Whites	8,063,991	117,617	15,589
Mean Income Asians	$131,123.33	$100,191.89	$100,328.80
Median Income Asians	$107,000	$89,000	$87,000
Mean Income Whites	$100,268.33	$100,191.89	$100,328.80
Median Income Whites	$76,000	$89,000	$87,000
% Asians buying owner occupied	0.89	0.95	0.94
% Whites buying owner occupied	0.86	0.95	0.943

Table 2: Pre- and post-matching population demographics. Matching was done exactly. Individual matching includes matching on the identity of the lending agent.

results to the matching). In addition, comparisons of the income levels of the pre- and post-matching populations show that the post-matched samples appears as a whole slightly less wealthy than the initial sample (Table 2); blacks are also less wealthy in the post-matched sample. (A likely reason for this is that minority populations tend to be less wealthy than non-minority populations; thus, matching minority populations to non-minority populations will result in the pruning of wealthy individuals, including possibly more non-minorities.) In addition, the mean income in the pre-matched population appears skewed by a handful of extremely wealthy individuals (for example, the mean income for Hispanics is $94,264.76, but the median is $78,000); this appears to be less of a concern in the matched sample (where the means and medians are equal or nearly equal). The matched samples also are more likely to be for homes that are intended to be owner occupied (as opposed to those intended to be investment properties). For purposes of this study, this means that we are focusing on buyers who are slightly less affluent than buyers as a whole (although wealthier than the population at large and with fewer outliers). These are also buyers interested in principal home ownership.

5 Results

The observations were matched four times, corresponding to four different "treatments" (Table 2). These were whether an applicant was (1) African American, (2) female, (3) Hispanic, or (4) Asian American. The baseline (or "control") groups were those categorized as white in the case of African Americans and Asian Americans, men in the case of women, and non-Hispanics in the case of Hispanics.[11] In addition to these four basic matching schemes, the matching either included lenders' unique identifiers or discarded that information, thus yielding the different substantive implications discussed above. After the data were matched,

[11] "Hispanic" is considered by the federal government to be an ethnicity, not a race. Applicants were therefore asked if they considered themselves Hispanic before they were asked to identify with one of the racial categories (white, black, Asian, Native American). It was therefore possible to be both Hispanic and black or Hispanic and Asian. For the sake of simplicity, Hispanics were therefore compared only to non-Hispanics, and blacks and Asian Americans to whites. There were not enough Native American borrowers to extract meaningful inferences.

subprime status, a binary "yes" or "no" variable, was regressed on the treatment variable.[12] The resulting estimate gives us the Average Treatment Effect (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 identifier – 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 filed an application for a mortgage loan with the same lenders as their non-minority counterparts. Thus, any remaining effect of the race or gender variable will, assuming no omitted variables, only be picking up impermissible racial considerations. Table 3 presents these figures 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 variable is binary (0 or 1), the inferences are the same.

[13]The data could, of course, be used to ask and answer a variety of related questions – for example, whether minorities and women were rejected for loans at disproportionate rates.

[14]In terms of generalizability, the matched borrowers were slightly less wealthy than the general sample, but not overwhelmingly so – see Table 2.

[15]"West" included Idaho, Nevada, Utah, Colorado, Arizona, New Mexico, Alaska, Washington State, Oregon, California, and Hawaii. "South" includes Washington D.C., Virginia, West Virginia, North Carolina, South Carolina, Georgia, Florida, Kentucky, Tennessee, Mississippi, Alabama, Oklahoma, Texas, Arkansas, and Louisiana. "Midwest" includes Wisconsin, Michigan, Illinois, Indiana, Ohio, North Dakota, South Dakota, Nebraska, Kansas, Minnesota, Iowa, and Missouri. Northeast includes Maine, New Hampshire, Vermont, Massachusetts, Rhode Island, Connecticut, New York, Pennsylvania, New Jersey, Delaware, and Maryland.

state by state.

The results are the strongest and most interesting for those borrowers classified 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 confidence 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 effects are quite sizable and, for about half the states, significant. For some states – Wisconsin, Mississippi, Indiana, and Nevada – being classified 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 financial and demographic variables and compared only borrowers going to the same lenders, this is quite a strong result.

The estimates are significantly more informative when looking at specific regions of the country and also at the United States as a whole. Indeed, pooling over all of the states gives us a treatment effect on the African-American variable of 0.575, and it is significant 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 financial profiles 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 effect 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.

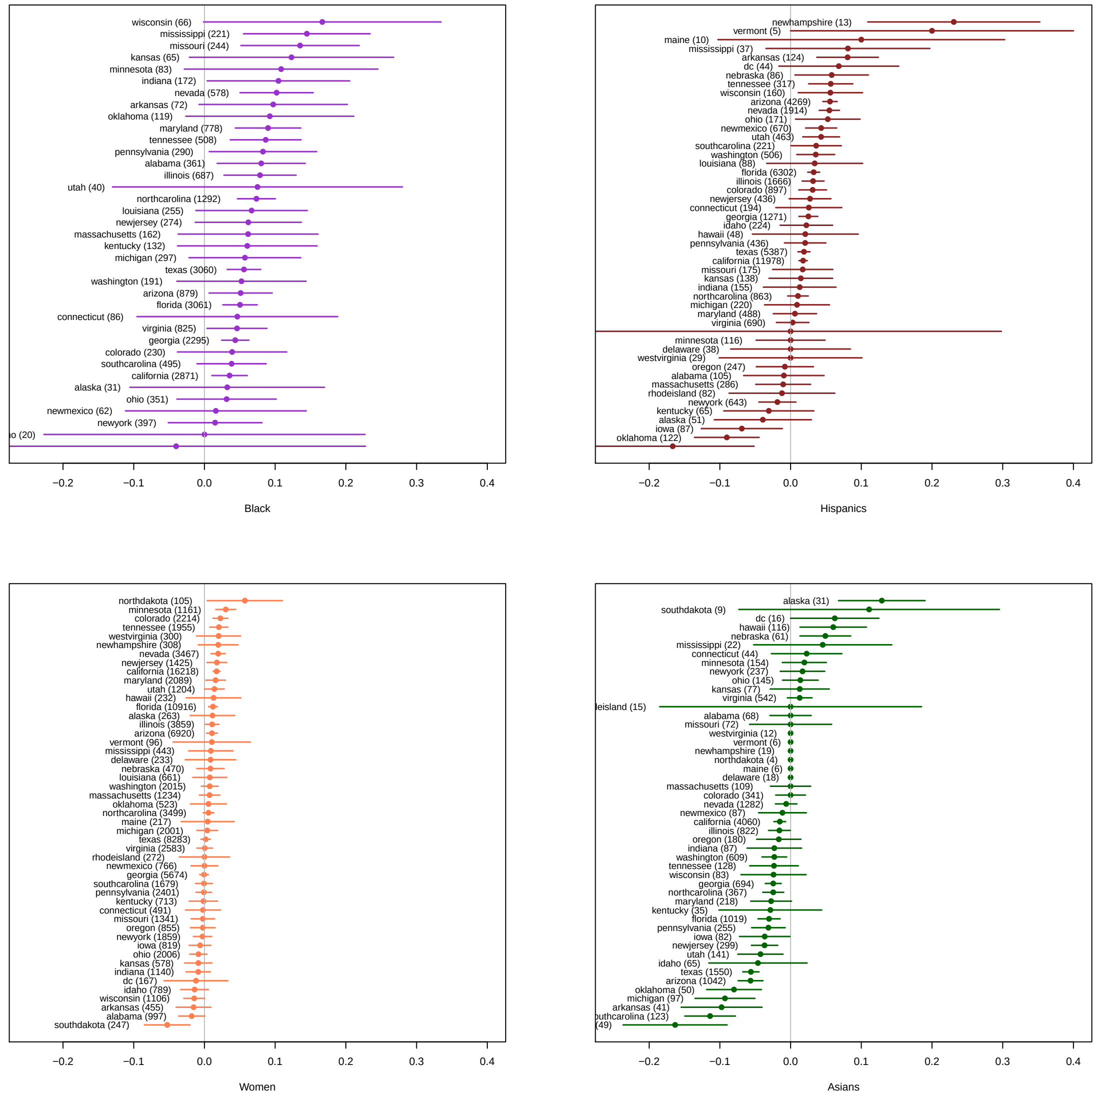

Figure 1: Likelihood of receiving a high-cost loan (individual-level measurements). The solid dots represent the point estimates for the average treatment effect in each state, while the numbers in parentheses represent the number of treated units in the matching. In substantive terms, the point estimates 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% confidence interval. If the spread includes the number zero, then we fail to reject the null hypothesis (at the $\alpha = 0.05$ level) that there is no effect of a borrower's noted race or gender on receiving a high-cost loan. For the most part, the results suggest positive effects associated with being African American, Hispanic, and female, but negative effects for being Asian American.

	US Total		Midwest		South		West		Northeast	
	N	ATE	N	ATE	N	ATE	N	ATE	N	ATE
Blacks	21238	0.0575	2062	0.0839	12106	0.0565	4929	0.0463	2141	0.0635
		(0.00443)		(0.0148)		(0.00567)		(0.00958)		(0.0142)
Women	99249	0.00788	14833	0.00175	38848	0.00494	34943	0.0142	10625	0.00649
		(0.00187)		(0.00470)		(0.00299)		(0.00323)		(0.00548)
Hispanics	42561	0.0256	2998	0.0254	15665	0.0244	21267	0.0307	2631	0.00684
		(0.00308)		(0.0113)		(0.00503)		(0.00439)		(0.0126)
Asian Am	15589	-0.0226	1693	-0.011223	4716	-0.0358	7954	-0.0182	1226	-0.0163
		(0.00381)		(0.00996)		(0.00617)		(0.00585)		(0.0114)

Table 3: Average Treatment Effect (individual-level measurements), subprime status as the outcome variable. The total number of treated individuals matched is on the left, while the coefficient estimate of the treatment variable is on the right. Standard errors are reported in parentheses.

around zero, with a handful of states above and below. The effects for some of the states are significant, but, for the most part, we are unable to reject the null hypothesis that being classified as a woman has no effect. Indeed, pooling over all of the states gives us a coefficient on the gender variable of 0.00778. Even though this is a statistically significant result (with 99,249 women in our matched sample, it is almost certainly guaranteed to be!), it is sufficiently close to zero to suggest remarkably little effect of the gender variable on the decision to award a subprime loan. The same is true for all regions of the country. This points to one conclusion: when women borrowers go to the same lenders as identically situated male borrowers, they are treated more or less the same.

The results for Hispanics are more provocative. Being categorized as Hispanic for the most part has a slightly positive effect on the probability of being awarded a subprime loan in many of the states. However, for many states, the 95 percent confidence bands around the estimates include zero, meaning that we cannot reject the null hypothesis that there is no effect of being Hispanic on the ultimate lending decision. Similarly, pooling over all of the states gives a coefficient on the Hispanic variable of 0.0256. Although significant (with 42,561 Hispanic borrowers matched in this sample), this is not a particularly large effect. 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 offered 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 fifty states plus Washington D.C. results in a slightly more modest estimate of -0.0226, but it is still significant at the 1 percent level. Note that Asian Americans living in the South appear to benefit the most, with an expected 3.58 percent *decrease* in the probability of being awarded a high-cost loan.

To summarize the results so far, it does appear that race factors in the decision to award a higher-cost loan. Even when matching exactly on all recorded financial information, and even when only looking at identical lenders, both African Americans and Hispanics still have a higher probability of being awarded high-cost loans in many states. Asian Americans, for their part, appear to have a lower probability of being awarded a high-cost loan. Women appear to be treated on par with men. Note that these results could still be called into 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 significant. The patterns are also strong when we pool across over all the United States and across various regions of the country. African-American borrowers living in the Midwest appear particularly disadvantaged, with a 21 percent greater probability of receiving a high-cost loan than similarly situated white borrowers. On the other hand, African-American borrowers living on the West coast fare better, with a 14 percent greater probability. These pooled results are all highly significant, with p-values less than 0.01.

Applicants classified as Hispanic fare better than African Americans, but there is still evidence that they are worse off than if they had gone to the same lenders as non-Hispanic borrowers. Compared to those categorized as non-Hispanics, these applicants experience an approximate 10-20 percent increased likelihood of receiving a high-cost loan in a large number of states, with the worst offenders being Oregon, Utah, Virginia, Washington, Arizona, and Minnesota. In fact, Hispanic borrowers living on the West coast are generally worse off, which stands in contrast with African-American borrowers, who are most disadvantaged

[17] There are relatively few African Americans living in both states, resulting in a smaller sample on which to draw inferences. The confidence intervals around the point estimate are consequently larger.

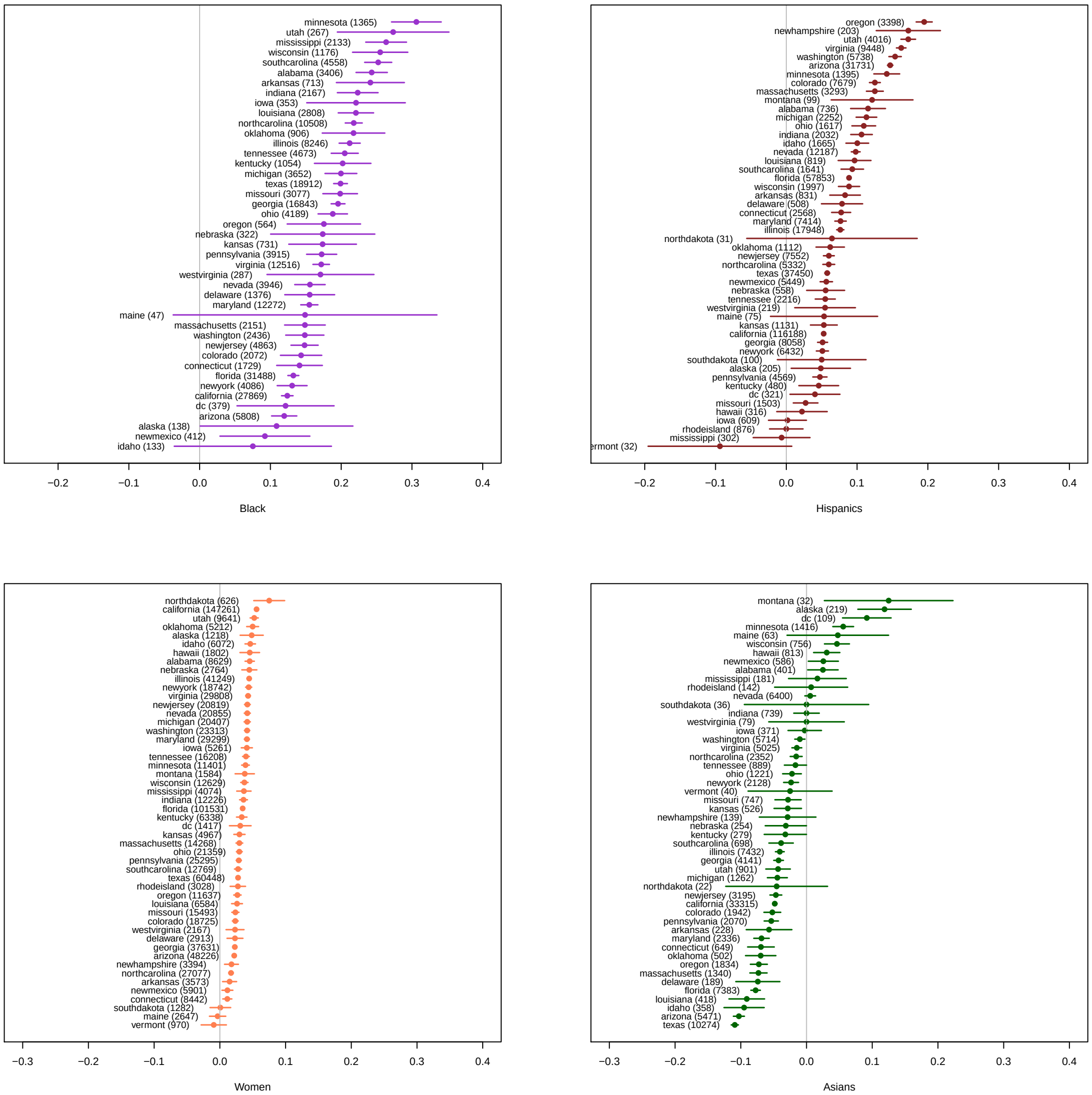

Figure 2: Likelihood of receiving a high-cost loan (structural-level measurements). The solid dots represent the average treatment effects for each state. In substantive terms, these are estimates of the increased or decreased probability that the treatment group in question will receive a more expensive loan. The line around each point represents the 95 percent confidence interval. If the spread includes the number zero, then we fail to reject the null hypothesis (at the $\alpha = 0.05$ level) that there is no effect of a borrower's noted race or gender on receiving a high-cost loan. For the most part, the results suggest positive effects associated with being African American, Hispanic, and female, but negative effects 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, in which we only looked at women borrowers who had gone to the same lenders as their male counterparts. In that case, the effect of being a woman was essentially nill. The positive effects seen here therefore lend some credence to the structural theory of discrimination.

Once again, borrowers categorized by the HMDA as Asian American present the most interesting, and perhaps unexpected, results. With the exception of Alaska, Maine, Minnesota, and Wisconsin – where Asian-American borrowers have a higher likelihood of being offered high-cost loans compared to their white counterparts – Asian-American borrowers are either statistically indistinguishable from applicants believed by lenders to be white (in which case the confidence interval contains zero) or, in fact, less likely to be offered high-cost loans. The trend is particularly striking in Texas, Arizona, and Florida, where applicants categorized by lenders as Asian American have up to a 10 percent lower likelihood than similarly situated white applicants to be offered a high-cost loan. In fact, Asian Americans living in the South have a 6 percent decreased probability of getting a high-cost loan (4.65 percent across the country overall) and this effect is statistically significant with a p-value less than 0.01.

What the results here demonstrate is that the general impact of being African-American, Hispanic, Asian-American, or a woman is amplified when we look at the entire lending industry, and not just at borrowers who went to the same lenders. To some extent, the results provide evidence in favor of the argument that structural discrimination might be at work. After all, if there were no structural differences among lenders – and minority and non-minority lenders went to the same sorts of lenders – then the results presented in this section would look similar to the ones presented earlier. That they do not suggests that minorities do in fact go to different lenders and these lenders are more (or less, in the case of Asian Americans) likely to offer loans with harsher terms.

6 Sensitivity Analysis

A potential problem with the results is that there might exist key variables that lenders do not report to the federal government. Unfortunately for researchers, the HMDA does not mandate that lenders report data on borrowers' credit scores, employment histories, savings available for down payments, and debt obligations (as they are considered both proprietary and confidential). These financial factors could have a substantial impact on the kind of loan terms a borrower is offered. After all, even if two people have the same income and are looking to buy a home in the same area, they will be offered very different loan products if one of the borrowers has a distressed credit score or a shaky employment history.[18] Likewise, individuals with different down payment amounts will also be offered different loan products.

Important to this discussion is that there exists strong evidence that borrowers from different racial and ethnic groups have different credit scores. For example, a 2007 report to Congress from the Federal Reserve found that normalized credit scores for whites was 54.0, while for Asian Americans it was 54.8, for Hispanics it was 38.2 and for blacks it was 25.6. However, the difference between these groups fell markedly once differences in income were accounted for. Indeed, the report noted that accounting for differences in income between the groups reduced the gap between blacks and whites by roughly half and the gap between whites and Hispanics by three-fourths. We can therefore expect that similarly earning blacks and Hispanics do have lower credit scores than whites, but but perhaps not by too much.

To gain leverage on these omitted variables, I proceed using sensitivity analysis of the kind described by Rosenbaum (2002), which posits that the treatment group might have some

[18]An advantage to matching is that it allows us to account to some extent for information not reported specifically in the data. For example, we might believe that employment correlates with income and education levels, and we know that these vary from community to community. If we match on borrower neighborhoods (via the census tract variable – which *is* recorded by the HMDA and *was* matched on) then we essentially control for this regional fluctuation and thus, to some extent, control for varying work histories. Nonetheless, employment histories (and likewise credit reports and debt obligations) could still influence our results, leading us to biased estimates.

variable (for example, bad credit scores, insufficient down payments, poor working histories, or oppressive debt obligations) Γ times more frequently than the control group. So, logically, if the treatment were truly randomized, and thus independent of all confounders, then $\Gamma = 1$. But if the treatment were applied non-randomly, perhaps due to a missing confounder, then we would expect a value of $\Gamma > 1$, such that the treatment group would be more likely to be affected by the confounder. For example, when $\Gamma = 2$, 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 Γ values, I use the specification developed by Rosenbaum (2002) and implemented by Keele (2009).[19] For each possible value of $\Gamma > 1$, the sensitivity test recalculates the p-values associated with the treatment effect. It continues increasing hypothetically the Γ value until the p-values are no longer statistically significant. Thus, once the p-value for any given Γ exceeds 0.05 (the maximum p-value that constitutes "statistical significance" in most social science literature), the sensitivity test "stops" because it has found the level at which the treatment effect is no longer significant. 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 Γ value for observational studies such as this one. Rosenbaum (2005) provides some guidance, particularly with regard to experimental studies, for which Γ values greater than 3 or 4 appear to be the norm. But for observational studies, more attainable Γ values might be found in Keele (2009), which reproduces the results of the Lalonde job training study. Replicating the treatment effect estimated via experimental techniques yielded upper bound on possible Γ values of 1.3. This benchmark represents a much lower Γ value than those calculated by Rosenbaum (2005) for

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

	US Total	Midwest	South	West	Northeast
Blacks	1.17	1.16	1.18	1.07	1.11
Women	1.01	-	-	1.07	-
Hispanics	1.07	1.01	1.05	1.08	-
Asian Am	1.13	-	1.29	1.04	-

Table 5: Sensitivity Results (individual-level discrimination): The cells contain the largest Γ values at which the treatment effects observed are still statistically significant (e.g., have p-values less than 0.05). The dashed line are instances where the original results were not statistically significant. For example, African-American borrowers living in the South would have to possess a trait 1.18 times as much as whites in order for the results to lose their statistical significance. This could beNote that because Γ represents the amount of bias we can tolerate before the treatment effects lose their statistical significance, we would like high Γ values. While true that Γ values of 3 or 4 would be ideal, but perhaps values between 1 and 1.5 are more realistic for observational studies (Keele, 2009).

the experimental studies, but, as this is the norm for observational studies, it is the general rule of thumb followed here.

Tables 4 and 6 display the regional and country-wide results of sensitivity analyses conducted on results of the matching analyses, while Tables 5 and 7 display the results from ten key states. (The entirety of the state-by-state results are reported separately in the Appendix.) When we look at the results that measure out individual-level discrimination, they are quite sensitive – even for African-American borrowers, which displayed the strongest results. For example, African-American borrowers living in the South would roughly have to have a characteristic (such as a poor credit score) 1.17 times as often as white borrowers in order for the results in that region of the country to become statistically insignificant. The strongest (more robust) results are actually those of Asian Americans, whose 3.5 percent decreased likelihood of receiving a subprime loan would become insignificant so long as they possessed some characteristic 1.29 times as often as the white borrower population. This could come in the form of credit scores above 620, or having a 20% down payment. The state-by-state analyses are likewise fairly sensitive. The few exceptions are North Carolina, where the African-American borrowing population must possess some trait 1.4 times "as

	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 significant at the 5% level. Again, because Γ represents the amount of bias we can tolerate before the treatment effects lose their statistical significance, we would like high Γ 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 significance, 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 Γ values greater than 1.6. What this means is that an omitted confounder – perhaps poor credit scores, insufficient 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 significance. (With a Γ 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 Γ values at which the treatment effects observed are still statistically significant (e.g., have p-values less than 0.05). For example, borrowers classified 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 insignificant. With the exception of the results for women (which are quite sensitive), the results are in the range of Γ values typically seen in observational studies.

the rates of whites for the results to be affected.) 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 Γ values, suggesting that these treatment effects 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 Γ 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 ($\Gamma = 1.50$ and $\Gamma = 1.60$), and Texas, for which the results for Asian Americans are also more robust ($\Gamma = 1.70$).

7 Conclusions

The substantive results of this research are roughly threefold. The first of these is that there appears to be somewhat limited (and perhaps sensitive) evidence of structural discrimination. In particular, the data provide some support for the widely held belief that

	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 Γ 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: Coefficients 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
Texas	1.50	1.00	1.10	1.70
Virginia	1.50	1.10	1.50	1.00
Colorado	1.40	1.00	1.30	1.10
Delaware	1.40	-	1.00	1.00
Iowa	1.40	1.10	-	-
Kansas	1.40	1.00	1.00	-
Maryland	1.40	1.00	1.10	1.20
Michigan	1.40	1.00	1.20	1.00
Missouri	1.40	1.00	-	-
Ohio	1.40	1.00	1.20	-
Pennsylvania	1.40	1.00	1.00	1.20
California	1.30	1.10	1.10	1.10
Florida	1.30	1.00	1.20	1.20
Massachusetts	1.30	1.00	1.20	1.40
New Jersey	1.30	1.10	1.10	1.20
Nevada	1.30	1.10	1.20	0.00
Washington	1.30	1.10	1.40	-
Arizona	1.20	1.00	1.30	1.30
Connecticut	1.20	-	1.10	1.20
DC	1.20	-	-	1.60
Nebraska	1.20	1.10	1.00	-
New York	1.20	1.10	1.00	1.00
vestigial	1.20	-	-	0.00
Mexico	1.10	-	1.10	0.00
Alaska	1.00	1.00	-	1.20
Rhineland	1.00	1.00	0.00	0.00
Idaho	-	1.10	1.20	1.10
Maine	-	0.00	-	0.00
Montana	-	1.00	1.00	0.00
Northeast	-	1.10	-	-
New Hampshire	-	-	1.20	-
South Dakota	-	-	-	0.00
Vermont	-	0.00	0.00	-

Table 10: Sensitivity results (structural-level measure). Note that the results are ordered by the sensitivity of the African-American treatment variable, from highest to lowest.

	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 Γ value because only members of the treated group received subprime loans. The dash indicates that the original results were not statistically significant at the 5% level.

References

Apgar, W.C. and A. Calder. 2005. "The dual mortgage market: the persistence of discrimination in mortgage lending." *The Geography of Opportunity: Race and Housing Choice in Metropolitan America* pp. 101–23. 7

Barr, M.S., J.K. Dokko and B.J. Keys. 2007. "Who Gets Lost in the Subprime Mortgage Fallout? Homeowners in Low-and Moderate-Income Neighborhoods.". 6

Becker, G.S. 1971. *The Economics of Discrimination*. University of Chicago Press. 4

Bertrand, M. and S. Mullainathan. 2004. "Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination." *American economic review* 94(4):991–1013. 4, 9

Blanton, Kimberly. 2007. "A 'smoking gun' on race, subprime loans: Fresh evidence shows minorities in Boston more likely than whites to get them when buying home." *Boston Globe* . 5

Board of Governors of the Federal Reserve System. 2007. "Report to Congress on Credit Scoring and Its Effect on the Availability and Affordability of Credit.".

Bocian, D.G., K.S. Ernst and W. Li. 2006. "Unfair lending: The effect of race and ethnicity on the price of subprime mortgages." *Center for Responsible Lending* . http://www.responsiblelending.org/mortgage-lending/tools-resources/rr011-Unfair_Lending-0506.pdf. 7

Bocian, D.G., K.S. Ernst and W. Li. 2008. "Race, ethnicity and subprime home loan pricing." *Journal of Economics and Business* 60(1-2):110–124. 6

Bradford, C. and L. Associates. 2002. "Risk or Race? Racial Disparities and the Subprime Refinance Market." *Center for Community Change* . http://www.cccfiles.org/shared/publications/downloads/Risk_or_Race_5-02.pdf. 7

Calem, P.S., K. Gillen and S. Wachter. 2004. "The neighborhood distribution of subprime mortgage lending." *The Journal of Real Estate Finance and Economics* 29(4):393–410. 6

California Reinvestment Coalition, Community Reinvestment Association of North Carlina, Empire Justice Center, Massachusetts Affordable Housing Alliance, Neighborhood Economic Development Advocacy Project, Ohio Fair Lending Coalition and Woodstock Institute. Unpublished 2009. "Paying More for the American Dream III: Promoting Responsible Lending to Lower-Income Communities and Communities of Color." http://www.empirejustice.org/assets/pdf/publications/reports/american-dream-III.pdf. 7

Courchane, M.J., B.J. Surette and P.M. Zorn. 2004. "Subprime borrowers: Mortgage transitions and outcomes." *The Journal of Real Estate Finance and Economics* 29(4):365–392. 6

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 Impact of Race, Ethnicity, and Gender on the Cost of Borrowing.” *Federal Reserve Bank of New York Staff Report* (368). 6

Hopkins, D.J. 2009. “No more wilder effect, never a whitman effect: when and why polls mislead about black and female candidates.” *The Journal of Politics* 71(03):769–781. 4

Iacus, S.M., G. King and G. Porro. 2009. “Causal Inference Without Balance Checking: Coarsened Exact Matching.”. 11

Keele, L. 2009. “An overview of rbounds: An R package for Rosenbaum bounds sensitivity analysis with matched data.”. 27, 28

King, G., R. Nielsen, C. Coberley, J.E. Pope and A. Wells. 2011. “Comparative effectiveness of matching methods for causal inference.” *Unpublished manuscript* . 12

Kirchoff, Sue. 2005. “Minorities depend on subprime loans.” *USA Today* . 5

Laderman, E. and C. Reid. 2008. Lending in low-and moderate-income neighborhoods in California: the performance of CRA lending during the subprime meltdown. In *Federal Reserve Bank of San Francisco*. pp. 2008–05. 11

Lax, H., M. Manti, P. Raca and P. Zorn. 2004. “Subprime lending: An investigation of economic efficiency.” *HOUSING POLICY DEBATE-WASHINGTON-* 15:533–572. 6

Mayer, C.J. and K. Pence. 2008. *Subprime mortgages: what, where, and to whom?* National Bureau of Economic Research Cambridge, Mass., USA. 6

Myrdal, G. and S. Bok. 1996. *An American dilemma: The Negro problem and modern democracy*. Transaction Pub. 4

Ondrich, J., S. Ross and J. Yinger. 2003. “Now You See It, Now You Don’t: Why Do Real Estate Agents Withhold Available Houses from Black Customers?” *Review of Economics and Statistics* 85(4):854–873. 7

Powell, Michael and Janet Roberts. 2009. “Minorities Affected Most as New York Foreclosures Rise.” *New York Times* . 2, 5

Rosenbaum, P.R. 2002. *Observational studies*. Springer Verlag. 26, 27

Rosenbaum, P.R. 2005. “Heterogeneity and causality: Unit heterogeneity and design sensitivity in observational studies.” *The American Statistician* 59(2):147–153. 27

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, nonprofit 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.” *The American Economic Review* 76(5):881–893. 7, 9, 10, 12