

SELF-SELECTION AND DISCRIMINATION IN CREDIT MARKETS^{*}

Stanley D. Longhofer
Wichita State University
Barton School of Business
1845 N. Fairmount
Wichita, KS 67260-0077
stan.longhofer@wichita.edu

Stephen R. Peters
Department of Finance
College of Business Administration
Kansas State University
177D Calvin Hall
Manhattan, KS 66506-0503

Preliminary Draft: November 1997
Current Draft: August 2002

* The initial work on this project was completed while Longhofer was at the Federal Reserve Bank of Cleveland. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Cleveland or of the Board of Governors of the Federal Reserve System. Peters acknowledges financial support from the Center for Real Estate and Urban Economic Studies at the University of Connecticut. The authors are grateful to Brent Ambrose, Ben Craig, Curt Hunter, Charlie Kahn, Mark Sniderman, Tony Yezer, seminar participants at the University of Tennessee and the 1999 American Real Estate and Urban Economics Annual Meeting, and an anonymous referee for helpful comments and discussions.

SELF-SELECTION AND DISCRIMINATION IN CREDIT MARKETS

Stanley D. Longhofer
Wichita State University

Stephen R. Peters
Kansas State University

JEL Subject Codes: G21, G28, J71

Abstract

In this paper we make two contributions toward a better understanding of the causes and consequences of discrimination in credit markets. First, we develop an explicit theoretical model of loan underwriting in which lenders use a simple Bayesian updating process to evaluate applicant creditworthiness. Using a signal correlated with an applicant's true creditworthiness and their prior beliefs about the distribution of credit risk in the applicant pool, lenders are able to evaluate an applicant's expected or "inferred" creditworthiness to determine which loans to approve and which to deny. Second, we explicitly model the self-selection behavior of individuals to show how market frictions like Beckerian "tastes for discrimination" can affect application decisions. Because these decisions shape banks' prior beliefs about the distribution of credit risk, they also affect the Bayesian posterior from which banks compute an applicant's inferred creditworthiness, implying that statistical discrimination can arise endogenously. In a market in which *only some* lenders have tastes for discrimination, we show that there are conditions under which lenders without racial animus will also discriminate. The model produces a number of empirical and policy implications.

1. Introduction

The issue of discrimination has remained as a fundamental challenge for policymakers throughout the latter half of the 20th century. As demonstrated by Gary Becker's seminal work on this topic (Becker, 1971), economics can play a vital role in helping to inform and guide policy. Most of the economic research in this area has focused on labor markets. Receiving less attention has been discrimination in product markets, although there are notable exceptions (e.g., Ayers and Seigleman, 1995). Recently, the question of whether financial service providers discriminate has come to the forefront as well (Munnell et al., 1992, 1996, Harrington and Neihaus, 1998, and Cavalluzzo and Cavalluzzo, 1998).

Ignited by empirical findings reported in the so-called "Boston Fed study" (Munnell et al., 1992, 1996) and the availability of new data through the Home Mortgage Disclosure Act (HMDA), mortgage markets have emerged over the last decade as the primary focus of most research on discrimination in credit markets. In the years following the initial Boston Fed study, academics, bankers, activists, and policy makers have struggled to agree on how best to rectify discriminatory practices in consumer credit markets. At the same time, no clear consensus has been reached on whether or not lenders actually do discriminate. At the heart of this conundrum is the difficulty in establishing what discrimination looks like and how it might be detected. This paper is an attempt to confront these issues by studying the loan underwriting process.

Virtually all of the research on lending discrimination has been in the form of empirical studies of mortgage discrimination. Much of it has focused on the validity of Munnell et al.'s results and on finding ways to analyze and detect discrimination in mortgage lending data.¹ In contrast, there has been little work on economic theories that might explain discriminatory behavior in credit markets or that would provide a framework for studying the loan underwriting process.² This lack of economic theory has not only forestalled policy debates, but also has hindered the design of appropriate empirical tests for lending discrimination.

¹ See, for example, Berkovec et al. (1994), Horne (1994, 1997), Yezer et al. (1994), Bostic (1996), Hunter and Walker (1996), Black et al. (1997), Bostic and Canner (1998), Day and Leibowitz (1998), and Ross (2000).

² Some theoretical contributions include Calomiris et al. (1994), Calem and Stutzer (1995), Ferguson and Peters (1995, 1997, 2000), and Han (1998).

Our paper makes two contributions toward a better understanding of the causes and consequences of discrimination in credit markets, yielding interesting insights into the behavior of both lenders and applicants. First, we develop an explicit theoretical model of the loan underwriting process that accounts for lenders' efforts to ascertain applicant creditworthiness in the presence of imperfect information. In our model, lenders use a simple Bayesian updating process to evaluate applicant creditworthiness, using both the information provided by the applicant and the lender's own prior underwriting experiences to determine which loans to approve and which to deny.

This structural model allows us to define discrimination with respect to observable variables, making it more useful to both empiricists and policy makers. The model can be used to design empirical tests to uncover any discrimination that may exist, as well as to reveal the underlying motivation that gave rise to this discrimination. For example, we show that the conditional default rate of minorities is lower than that of whites at bigoted banks (as suggested by Becker, 1993), while the opposite is true at banks that statistically discriminate against minorities (consistent with the findings of Berkovec, et al., 1994). Thus, while a Becker-style test on loan default rates would be capable of identifying bigoted lenders, it would allow statistical discriminators to discriminate with impunity. To avoid this problem, we describe a test involving applicant denial rates that not only can determine whether lenders discriminate, but also can ascertain whether that discrimination arises because of lender beliefs (statistical discrimination) or lender preferences (bigotry).

Our second contribution lies in our focus on the self-selection behavior of individual applicants. We examine how the market as a whole responds when a subset of lenders have "tastes for discrimination" (Becker, 1971). Our model illustrates how minority applicants' reactions to this discrimination alter the relative distributions of minority and white credit risk in bank applicant pools, giving *non-bigoted* lenders an incentive to statistically discriminate against minorities. This statistical discrimination arises endogenously as a result of applicant responses to bigotry.³

³ Throughout this paper, we use the term "bigotry" to refer to lenders with Beckerian tastes for discrimination.

The dearth of theoretical work addressing the problem of discrimination in credit markets may arise out of a belief that theoretical results developed for labor markets apply directly to credit markets as well. This, however, is often not the case. One example of this can be seen in analyses of “cultural affinities.” Cornell and Welch (1996) consider a labor market in which employers are better able to assess the true productivity of job applicants with whom they share a common cultural background and show that employers will naturally tend to hire applicants with whom they have a cultural affinity, even in the absence of any tastes for discrimination. In contrast, Longhofer (1996) analyzes virtually the same informational problem in the context of credit markets and shows that it gives lenders an incentive to discriminate *against* applicants with whom they have a cultural affinity.⁴ Thus, the same informational friction provides two diametrically opposing predictions for labor and credit markets, demonstrating the importance of modeling discrimination separately for each. Our paper provides a major step in filling this void in the literature.

Recent empirical work has highlighted the importance of the applicant self-selection behavior that we model in this paper. Bostic and Canner (1998) find that differences in mortgage applicant pools across black-owned, Asian-owned, and white-owned peer banks account for most of their denial-rate disparities. Rosenblatt (1997) finds strong evidence that an individual’s choice between conventional and FHA mortgage products (as well as whether to apply at all) is based on how well their personal characteristics match the requirements of each. Finally, Avery et al. (1994) show that cross-lender variation in minority and low-income originations primarily reflects differences in home mortgage applications rates, again supporting

Importantly, lenders who statistically discriminate are *not* bigots.

⁴ The intuition behind this anomalous result lies in understanding that affinities act to reduce the noise of an applicant’s signal. As a result, lenders and employers have a more difficult time assessing the true quality of applicants with whom they do not have a cultural affinity, who all tend to look like the average regardless of their true quality. This tends to work to the disadvantage of the best applicants (who are less able to demonstrate their merit) and to the advantage of low-quality applicants (who find it easier to pass themselves off as good). Because the employer’s problem is typically to select the most desirable applicant, it only seriously considers applicants sending signals well above the mean, benefiting the group with whom it has an affinity. In contrast, lenders generally approve all applicants that exceed a given threshold, often accepting a majority of the applications it receives. As a result, its cutoff threshold is below the mean of the applicant pool, in a range where applicants would just as soon have noisy signals.

Heckman (1998) discusses how this kind of problem arises in “audit pair” tests for labor market discrimination. He argues these tests are flawed because they typically do not control for the qualifications of the pair relative to the qualifications necessary for acceptance.

our notion that applicant decisions have a strong influence on lender behavior and other market outcomes.

In the next section, we review several theories of discrimination and discuss how these models relate to our own. In section 3, we introduce our model of loan underwriting and show how lenders use their past underwriting experiences and beliefs to evaluate applicant creditworthiness. In the following section, we formally define discrimination and show how different motives can give rise to this behavior. In section 5, we allow minorities to self-select among banks and demonstrate how this behavior can affect bank underwriting standards and other market outcomes. Because both individual application decisions and bank lending decisions are endogenous in our model, it provides a rich environment for analyzing a variety of empirical and policy questions; we discuss these in Section 6. Section 7 concludes, while proofs of all results and propositions are found in the Appendix.

2. Discrimination in Credit Markets

Becker (1971) pioneered the economic analysis of discrimination, developing a theory based on the preference a bigot has for one group over another. As Becker argues, a taste for discrimination makes a bigot willing to expend a cost (or forgo a benefit) to associate with a preferred group. As applied to credit markets, a bigoted lender will hold applicants from its preferred group to a lower credit standard than applicants from another group, causing the lender to make loans to high-risk applicants from the preferred group, while denying equally risky applicants from other groups.⁵ Alternatively, bigoted lenders may charge lower interest rates or fees to members of the preferred group.

Becker's theory is one of *preference-based* discrimination. In contrast, statistical or *belief-based* discrimination can arise when the characteristics of an individual's group are used to evaluate his or her personal characteristics.⁶ For example, if liquid assets or income are better

⁵ Ferguson and Peters (2000), however, show that if credit is rationed a bigoted lender can discriminate without forgoing profits and without creating a difference in marginal default rates between minorities and non-minorities.

⁶ Statistical discrimination is sometimes distinguished from taste-based discrimination through labels such as "economic" or "rational" discrimination. Such labels can be misleading, however, since bigots are perfectly rational economic agents, maximizing their (albeit socially-condemned) preference function subject to constraints.

predictors of default risk for members of one group than another, even non-bigoted lenders may want to take group membership into consideration when making underwriting decisions. Arrow (1972a, 1972b, 1973) and, separately, Phelps (1972) were among the first to consider models of statistical discrimination, focusing on the problem of an employer with exogenously given beliefs that the average productivity of white labor is higher than that of minority labor. Callem and Stutzer (1995) model a similar problem in the context of credit markets using a Rothschild-Stiglitz (1976) framework.

More recent research has suggested that cultural differences between banks and borrowers (Calomiris et al., 1994, and Longhofer, 1996) or between employers and job applicants (Cornell and Welch, 1996) can lead to a form of statistical discrimination. The intuition behind this theory is that lenders may share a “cultural affinity” with members of one group due to a common ethnic background, race, religious belief, gender, education, or other social bond. Ferguson and Peters (1997) extend this idea, arguing that affinities may arise endogenously based on a bank’s experience in working with applicants from different groups. Whether cultural or experiential, however, an affinity allows lenders to more accurately evaluate the creditworthiness of applicants with whom they share this common bond.

Our underwriting model is general enough to incorporate any of these potential sources of discrimination. In this paper, we choose to focus on bigotry and the impact that it has not only on minority applicants at bigoted banks, but also on the behavior of non-bigoted banks and its applicants—both minority and white. We illustrate how the choices that individuals make regarding whether or not to apply for loans—and to which bank they apply—are based on the likelihood that they will be approved and, therefore, imply a correlation between the creditworthiness of a bank’s applicant pool and race. Profit- and utility-maximizing banks have an incentive to use the information this correlation reveals in order to more accurately assess credit risk. In other words, applicant self-selection behavior can lead to endogenous differences in the average creditworthiness of different groups. As a result, statistical discrimination arises *even if lenders believe that the distribution of creditworthiness is the same across groups in the general population*. In other words, bigotry on the part of some banks causes applicant self

selection behavior that makes non-bigoted banks statistically discriminate.

The empirical predictions that arise from our model are consistent with many existing empirical studies on discrimination in credit markets. In addition, our model clarifies the importance of the source of discrimination for testing and policy. It is important to keep in mind, however, that the primary contribution of our paper is not simply in the specific empirical predictions it may provide, but also in the structure it provides for thinking about how market frictions may affect the underwriting process. Our framework for analysis will provide a foundation for researchers and policymakers to better understand the true causes of observed racial differences in credit markets, and to analyze how various policy corrections may affect the market.

3. A Model of Loan Underwriting

Consider a world in which individuals want to make a purchase such as a house or other consumer good but lack sufficient funds to do so. As a result, they must obtain loans from a financial institution, which we will call a “bank.” Each individual in the population is assumed to have a true creditworthiness represented by $\theta \in [0,1]$. We interpret θ as an individual’s likelihood of repaying his loan (although this interpretation is not required for our analysis), and we assume that it captures all of the factors that might cause him to default, including disruptions to his income, changes in the value of the asset financed, and his personal compunction about defaulting on an obligation.

Creditworthiness is assumed to be distributed throughout the population according to the probability density function $f(\theta)$, with cumulative distribution function $F(\theta) = \int_0^\theta f(t)dt$; all banks and individuals share these prior beliefs about the distribution of true creditworthiness in the population.

Although each individual’s creditworthiness, θ , is given exogenously, an individual’s application decision is endogenous.⁷ Therefore, we must distinguish between the distribution of

⁷ Our model allows individuals to choose whether to apply for a loan given their creditworthiness, but it does not allow them to enhance or alter their true underlying θ , in contrast to Yezer et al. (1994). Such an extension to our model would be feasible, in principle.

creditworthiness in the population, $F(\theta)$, and the distribution of creditworthiness in the applicant pool. We will denote this latter distribution by $G(\theta)$ with corresponding density $g(\theta)$.

The mortgage market, like many other consumer loan markets, is a “posted-price” market. Lenders set a loan price ex ante, taking into account their cost of funds, the loan terms, market conditions (including their expectations of the applicant pool, and competition from other lenders), and then approve all applicants that are profitable at the posted price. Let r^* denote the competitive market interest rate and r_c the bank’s cost of funds. For an applicant of type θ , the bank’s expected profit is $\theta r^* - r_c$, implying that the bank would like to offer loans to all applicants for whom $\theta \geq \theta^* \equiv r_c / r^*$.⁸

Unfortunately, banks cannot perfectly observe θ . Instead, they observe a signal, $s = \theta + \varepsilon$, where ε represents the bank’s errors in assessing credit risk. We assume that this signal aggregates all the information banks collect about an applicant’s creditworthiness (i.e., liquid assets, credit history, collateral characteristics, debt ratios, and any other information that banks collect in the application process).

Let $p(s | \theta)$ be the likelihood that a bank observes signal s from an applicant of type θ . The usefulness of the signal will depend on the characteristics of this signal generation process. For our analysis, we assume that p is characterized as follows:

⁸ This begs the question of how the competitive market rate r^* is determined. To help preserve the tractability of the analysis we simply assume its existence. A more thorough treatment of this problem would require developing an equilibrium concept capable of supporting a pooling equilibrium in the context of our market structure. Unfortunately, existing equilibrium concepts such as a Riley Reactive Equilibrium or a Wilson Anticipatory Equilibrium are not directly applicable in the present case. Although our framework is directly translatable into a Rothschild-Stilitz insurance world, we have further complicated this model by adding a screening technology that is applied *after* lenders have achieved ex ante sorting of applicants through a menu of contracts. As a result, equilibrium concepts that have been used to support pooling equilibria in these models break down in our own.

Given the empirical relevance of the single-price assumption in credit markets, developing an equilibrium concept that can support a pooling equilibrium in our framework would be an extremely useful exercise. It would also entail a significant detour from our present course. Because the pooling outcome is a natural—and empirically observed—case to consider in our analysis, we proceed by simply assuming this single price exists and is adopted by all lenders in a pooling market. In section 6 we discuss how our results would change in a market that exhibits risk-based pricing.

ASSUMPTION A1: For every applicant type θ ,

- 1) $E[s | \theta] = \theta$;
- 2) $p(\theta - \delta | \theta) = p(\theta + \delta | \theta)$, $\forall \delta$;
- 3) $p'(s | \theta) > 0, \forall s < \theta$ and $p'(s | \theta) < 0, \forall s > \theta$;
- 4) $\frac{\partial p(s | \theta) / \partial \theta}{p(s | \theta)}$ is increasing in s ; and
- 5) $p(s | \theta)$ is strictly positive on its support, $(-\infty, +\infty)$, and has continuous partial and cross-partial derivatives.

Part 1 of this assumption simply states that s is an unbiased signal of an applicant's true underlying creditworthiness, θ . Part 2 is a symmetry assumption, and guarantees that applicants are equally likely to send erroneous signals in either direction. Part 3 ensures that p is unimodal, or "hump-shaped." Part 4 states that p satisfies the monotone likelihood ratio property. Finally, part 5 is a technical assumption to facilitate some of the proofs that follow. All of these assumptions are reasonably standard and are satisfied by a number of common distributions, including the normal, logistic, and non-central t distributions.

Based on this signal generation process and the distribution of its applicants, the unconditional density of signals observed by a bank is

$$\omega(s) = \int_T p(s | \theta) g(\theta) d\theta, \quad (1)$$

where T is the set of all types that apply in equilibrium. We then define the likelihood that an applicant sending signal s has true creditworthiness θ as

$$\pi(\theta | s) = \frac{p(s | \theta) g(\theta)}{\omega(s)}, \quad \forall \theta \in T, \quad (2)$$

and 0 otherwise. Thus, π is the bank's Bayesian posterior beliefs on θ .

Banks are not interested in an applicant's signal, per se. Rather, they grant credit based on an applicant's expected creditworthiness derived from this posterior π .⁹ We will often refer to this expected creditworthiness as an applicant's inferred "quality," which is denoted by

⁹ This Bayesian process formally models the intuitive notion that lenders use their past experiences to interpret the information on loan applications. This is true regardless of whether lenders formally subject their loan applications to scoring models.

$$q(s) = \int_T \theta \pi(\theta | s) d\theta . \quad (3)$$

Therefore, the bank's expected return from a loan to an applicant sending signal s is

$$q(s)r^* - r_c ; \quad (4)$$

banks will approve any applicants for whom this expression is positive.

RESULT 1: *Inferred creditworthiness is increasing in the observed signal: $q'(s) > 0$.*

RESULT 2: *There exists a unique cutoff signal $s^* \equiv q^{-1}(r_c)$, such that every applicant with signal $s \geq s^*$ is approved.*

Result 1 is simply a formalization of the intuition that applicants who send better signals tend to be more creditworthy. More importantly, it assures us that there exists a unique cutoff signal determining which applicants will be approved (Result 2).

Using these results, we can write the bank's problem as choosing the cutoff signal that maximizes its total expected profit:

$$\max_s \int_{s^*}^{\infty} [q(s)r^* - r_c] \omega(s) ds , \quad (5)$$

Total bank profit is therefore simply the sum of the profit the bank earns on all those loans it approves. Differentiating with respect to s^* allows us to derive the bank's zero marginal profit condition: $q(s^*) = \frac{r_c}{r^*} \equiv q^*$.

Given this structure of the underwriting process, we can now determine which individuals will apply for loans. Let $\alpha(\theta)$ denote the probability that a type- θ applicant is approved for a loan. Using Result 2,

$$\alpha(\theta) = \int_{s^*}^{+\infty} p(s | \theta) ds . \quad (6)$$

It is straightforward to verify that the probability of being accepted for a loan is increasing in the applicant's type: $\alpha'(\theta) > 0$.

We assume that applying for a loan is costly, so that individuals will only do so if their chance of being approved is sufficiently high. Let B denote the gross benefit an individual

receives from a loan (the utility stream from owning the good financed by the loan) and C the individual's cost of applying for a loan (both shoe-leather and direct application costs). So long as $0 < C < B - r^*$, the fact that $\alpha'(\theta) > 0$ ensures that there exists a cutoff type θ^m such that

$$\alpha(\theta^m) = \int_{s^*}^{+\infty} p(s | \theta^m) ds = \frac{C}{B - r^*} \equiv \alpha^*, \quad (7)$$

implying that all individuals with $\theta \geq \theta^m$ will apply for loans, while those who are less creditworthy will not.¹⁰ We will often refer to θ^m as the “marginal” applicant type. Casual observation suggests that most individuals will only apply for a loan if their likelihood of being approved is relatively high. Thus, we make the following assumption.

ASSUMPTION A2: $\alpha^* > \frac{1}{2}$.

Given our assumption that p is symmetric, this implies that $s^* < \theta^m$ in any equilibrium.¹¹

We are now able to define an equilibrium in our model as the (s^*, θ^m) pair that simultaneously solves the following two equations:

$$\begin{aligned} q(s^*) &= q^* \\ \alpha(\theta^m) &= \alpha^*. \end{aligned} \quad (8)$$

4. Credit Market Discrimination

The model of loan underwriting developed in the last section is quite general, and can be applied to a wide variety of problems. In this section, we expand the model to consider the circumstances under which banks might discriminate in their underwriting decisions.

We divide the population of potential applicants into two groups, W and M , consistent

¹⁰ It is possible that $\theta^m \notin (0,1)$, in which case either all individuals apply for loans or remain out of the market. We will ignore this possibility throughout the rest of the paper.

¹¹ Note that this in turn implies that a majority of all loan applicants are approved. This is consistent, for example, with the findings from HMDA data, which indicate that over 80 percent of all mortgage applicants are approved. This does not suggest, however, that all applicants are acceptable to banks. In fact, it is easy to show that $\theta^m < q^*$ in any equilibrium; otherwise banks would accept all applicants without screening, which cannot occur in equilibrium.

with the common racial context of “white” and “minority.”¹² Similarly, we assume that individuals can apply at one of two different banks, B and N .¹³ Bank N has no preference for members of either group, while bank B has a Beckerian “taste for discrimination” against minority applicants (group M); we will commonly refer to bank B as the “bigoted” bank and bank N as the “non-bigoted” bank. Following Becker (1972), we employ a discrimination coefficient $\delta > 0$ as a measure of the utility loss a bigoted bank incurs when lending to a minority individual. Thus, the bigoted bank’s utility on a loan to a minority is

$$q_{MB}(s)(r^* - \delta) - r_c. \quad (9)$$

This implies that

$$q_{MB}^* \equiv \frac{r_c}{r^* - \delta} > q_{WB}^* = q_{MN}^* = q_{WN}^* \equiv \frac{r_c}{r^*}, \quad (10)$$

where q_{ij}^* is the lowest inferred quality a group i applicant may have and still be approved at bank j ; further uses of the subscripts W , M , B , and N should be interpreted similarly.¹⁴

To proceed, we must formally define what behavior constitutes “discrimination” in our model. Although this term is not explicitly defined in either the Equal Credit Opportunity Act or the Fair Housing Act, we employ a definition consistent with how these laws are enforced by regulators.¹⁵

DEFINITION: *A bank is said to **discriminate** against members of group i if it requires them to meet a higher cutoff signal than members of group j ($s_i^* > s_j^*$).*

Note that this definition of discrimination does *not* depend on the bank’s inferred creditworthiness standard, q^* . Rather, it is based solely on information that is directly verifiable by regulators and others outside the bank.

¹² This division may be along any racial, ethnic, cultural, gender, geographic, or other publicly observable characteristic that lenders can use to classify applicants; we merely use racial mnemonics for expositional convenience.

¹³ More accurately, we assume that there are two *types* of bank, all of which behave competitively. For simplicity, we will focus on the behavior of a representative bank of each type.

¹⁴ When a group’s treatment or behavior does not vary across banks we omit the bank subscript.

¹⁵ See Longhofer and Peters (1998) for a discussion of the confusion that has arisen because of the lack of a precise working definition of discrimination in the literature.

Alternatively, one might imagine that lenders could discriminate by applying the same cutoff signal to white and minority applicants, but instead charging minorities a higher interest rate (e.g., $r_M^* > r_W^*$). In a market characterized by a single price offered to all applicants who qualify—the mortgage market is a prime example—this will not occur. In particular, the fact that interest rates and other fees charged are observable to both regulators and potential applicants means that lenders will be unable to offer different prices to minority and white borrowers. First, minorities would avoid lenders posting such discriminatory prices, instead applying at non-discriminatory lenders. More practically, such lenders would be easy targets for both prosecution and negative publicity.¹⁶

It is also important to note that this definition of discrimination does not make a distinction among various motivations that might give rise to discriminatory behavior. To better understand this, consider Figure 1. Panel A depicts a market in which a bigoted lender's taste for discrimination causes it to set $q_M^* > q_W^*$. Assuming that the distribution of true credit risk (θ) is the same across both groups so that $q_W(s) = q_M(s)$ for all s , the bigoted lender will discriminate against minority applicants by setting $s_M^* > s_W^*$.

Of course, bigotry is not the only possible source of discriminatory behavior. Panel B depicts the actions of a lender that statistically discriminates when, for example, minorities are less creditworthy than whites on average. If this were true, $q_W(s)$ would lie everywhere above $q_M(s)$ and as a result, even though the lender has no taste for discrimination ($q_M^* = q_W^*$) its beliefs about the prior distribution of credit risk across the two groups would cause it to infer that a minority applicant sending a given signal s is less creditworthy than a white applicant sending that same signal. This gives the lender an incentive to statistically discriminate against minorities by setting $s_M^* > s_W^*$.

In this example, statistical discrimination arises because of the assumption that the

¹⁶ Recent empirical evidence supports this conclusion, showing that a lender's consumer loan terms vary surprisingly little with applicant characteristics (Sirmans and Benjamin, 1990, Duca and Rosenthal, 1994, Benjamin et al., 1995, and Avery et al., 1996). We should note, however, that charging different prices to white and minority applicants is more feasible in a market with risk-based pricing, because the schedule of prices offered makes it more difficult to determine whether an individual offer is discriminatory. We analyze this type of market in Section 6 below.

minority population is exogenously less creditworthy on average than the white population.¹⁷ More interesting, however, is the possibility that individual decisions about which bank to apply will affect the distribution of credit risk across banks, thereby giving rise to statistical discrimination even if the distribution of credit risk were the same across the two populations. We address this problem next.

5. The Impact of Self-Selection

Individuals' decisions over whether and at which bank to apply affect the loan underwriting process by altering the distribution of credit risk in applicant pools across groups and across banks. To analyze this phenomenon, we make

ASSUMPTION A3: $f_i(\theta) = 1$ for all $\theta \in [0,1]$, $i = M, W$.

Assuming a uniform distribution greatly simplifies our analysis of self-selection. For example, if there were no applicant or lender heterogeneity, $g(\theta) = 1/(1 - \theta^m)$, so that *applicant* creditworthiness would be uniformly distributed on the interval $[\theta^m, 1]$.

As discussed above, the bigoted bank has an incentive to discriminate against minority applicants, so that $s_{MB}^* > s_{MN}^*$. This in turn affects a minority applicant's chance of receiving a loan. Indeed, because $\alpha(\theta) = 1 - P(s^* | \theta)$ (where $P(s | \theta) \equiv \int_{-\infty}^s p(t | \theta) dt$), it must be the case that $\alpha_{MB}(\theta) < \alpha_{MN}(\theta)$ for every θ .

If minority applicants could costlessly and perfectly distinguish between the bigoted and non-bigoted bank, they would always choose to apply at the non-bigoted bank and minorities would never suffer from discrimination. Therefore, if discriminatory outcomes are to arise as a result of bigotry we must assume that the information that distinguishes between banks is costly and/or imperfect.

We assume applicants must pay a cost χ to acquire information about bank types, reflecting the time and effort required to inquire about a bank's reputation regarding its treatment

¹⁷ This would be the case, for instance, if $F_w(\theta)$ were first-order stochastic dominant over $F_M(\theta)$.

of minorities. To account for the noise inherent in such information, we assume that only a fraction $\gamma > \frac{1}{2}$ of individuals, uniformly distributed across the population, succeed in applying at their preferred bank. Thus, individuals know how they will be treated by each bank, but must expend a cost to learn which bank is which. Furthermore, because this information is imperfect some individuals who pay this cost nevertheless end up applying at the bigoted bank.

Given this setup, minority applicants will acquire information about banks' types if and only if¹⁸

$$(B - r^*)[\gamma\alpha_{MN}(\theta) + (1 - \gamma)\alpha_{MB}(\theta)] - C - \chi > (B - r^*)[\frac{1}{2}\alpha_{MN}(\theta) + \frac{1}{2}\alpha_{MB}(\theta)] - C, \quad (11)$$

which simplifies to

$$(B - r^*)(\gamma - \frac{1}{2})[\alpha_{MN}(\theta) - \alpha_{MB}(\theta)] > \chi. \quad (12)$$

The left-hand side of expression (12) is simply a minority's expected net benefit from acquiring information about bank types. It incorporates the gross benefit from obtaining a loan $(B - r^*)$, the increased probability of applying at the non-bigoted bank that results from this information $(\gamma - \frac{1}{2})$, and the higher probability of being approved at this bank $(\alpha_{MN}(\theta) - \alpha_{MB}(\theta))$. For any minority applicant with creditworthiness θ , if this net benefit exceeds the cost of acquiring information (χ) , he will attempt to self-select between banks.

Close examination of expression (12) verifies that if gathering information about banks' behavior is inexpensive (χ is small), all minority applicants will do so. On the other hand, if χ is large so that distinguishing between banks is quite costly, minorities will be better off avoiding these costs and taking their chances of ending up at the bigoted bank. In either case, the distribution of minority credit risk at the two banks is identical, and our basic underwriting model remains unchanged.

In contrast, for moderate costs of information acquisition minority self-selection behavior between the bigoted and non-bigoted banks can have a dramatic impact on the distribution of applicant credit risk at these two banks, and hence on lender discrimination incentives. This fact leads us to the following proposition.

¹⁸ Note that white applicants have no incentive to distinguish between banks, since they face the same cutoff signal regardless of where they apply.

PROPOSITION 1: For moderate levels of the information acquisition cost χ ,

- 1) There exists a cutoff type $\theta^c \in (\theta_M^m, 1)$ such that all minority applicants with creditworthiness $\theta \in [\theta_M^m, \theta^c]$ will gather information on bank types and attempt to self-select to the non-bigoted bank; those with a higher θ will not gather this information and will apply at each bank with probability $1/2$;
- 2) The inferred quality of minority applicants at the bigoted bank will be higher than that of minorities at the non-bigoted bank and of whites at either bank: $q_{MB}(s) > q_{MN}(s), \forall s$ and $q_{MB}(s) > q_W(s), \forall s$; and
- 3) Bigoted banks will discriminate against minority applicants by holding them to a higher cutoff signal than they do white applicants, and to a higher cutoff signal than do non-bigoted banks: $s_{MB}^* > s_{MN}^*$ and $s_{MB}^* > s_W^*$.

The initial part of this proposition follows from the fact that the value of information about bank types is decreasing in an individual's true creditworthiness, θ . That is, high- θ minorities are likely to be approved at either bank, so distinguishing between banks is less important. In contrast, it is low- θ minority applicants that bear the brunt of the bigoted bank's preferences, because they are more likely to receive signals that the non-bigoted bank would approve while the bigoted bank would not. As a result, low-quality minorities attempt to self-select away from the bigoted bank.

This self-selection behavior alters the composition of applicant pools across banks, which in turn affects how banks assess the inferred creditworthiness of minority and white applicants. Recall that an applicant's inferred creditworthiness $q(s)$ is simply the bank's expectation of that individual's true θ given its prior beliefs about the distribution of creditworthiness in its applicant pool and the signal sent by the individual. The second part of Proposition 1 therefore follows from the first. Because low- θ minorities try to avoid the bigoted bank, their average creditworthiness at that bank is higher than that of whites or that of minorities at the non-bigoted bank.

The third part of this proposition simply states that this self-selection behavior on the part of minority applicants is not sufficient to offset the bigoted bank's incentive to discriminate against minorities. When low- θ minorities self-select, it raises the average creditworthiness of the minority applicant pool at the bigoted bank, and lowers it at the other. This tends to bring s_{MB}^* and s_{MN}^* closer together. Nevertheless, it cannot entirely eliminate the gap between the two,

because it is this gap that gives minority applicants an incentive to self-select in the first place. Similarly, s_{MB}^* is always larger than s_W^* since the self-selection effect that tends to lower s_{MB}^* is second-order to the bigotry effect that raises it.

Our most striking result is that bigotry on the part of some banks can affect the relative treatment of both white and minority applicants at other banks. The decision of relatively low-creditworthy minorities to seek out the non-bigoted bank reduces the average quality of its minority applicant pool, which in turn gives it an incentive to statistically discriminate against minorities. Thus, the self-selection behavior of applicants can lead non-bigoted banks to discriminate against minorities, even though they have no tastes for discrimination.

PROPOSITION 2: *Applicant self-selection behavior can cause the non-bigoted bank to discriminate against minorities: $s_{MN}^* > s_{WN}^*$.*

Self-selection has two counteracting effects. First, the fact that low-creditworthy minorities seek to avoid the bigoted bank lowers the average creditworthiness of minority applicants at non-bigoted banks. This tends to raise s_{MN}^* above s_{WN}^* . On the other hand, the fact that some lower- θ minorities stop applying for loans altogether (because bigotry makes it less likely that they will be approved) tends to raise the average creditworthiness of the minority applicant pool, making it possible that $s_{MN}^* < s_{WN}^*$. Which of these two effects dominates depends on the parameters of the problem but, in general, the higher θ_M^m , the more signals for which $q_{MN}(s)$ will be larger than $q_{WN}(s)$. Thus, as long as bigotry doesn't force too many lower-quality minorities out of the market, their self-selection behavior will give the non-bigoted bank an incentive to discriminate against minorities. On the other hand, if the bigoted bank's taste for discrimination is very strong then θ_M^m will rise sufficiently high to cause the non-bigoted bank to statistically discriminate against whites (see Figure 2).

Finally, we would like to analyze how bigotry and self-selection interact to affect relative *default* rates across groups at each of the banks. To do this, note that $q(s)$ is an unbiased estimator of an individual's true creditworthiness (i.e., his probability of repaying his loan). In the aggregate, therefore, (abusing the law of large numbers in the standard way) the function

$1 - q(s)$ gives us the ex post default rate for each group at each bank, *conditional on the observed signal s* . For example, of the minority borrowers ending up at the bigoted bank and sending signals s , a fraction $1 - q_{MB}(s)$ will default, on average.

Our previous results, then, lead immediately to the following proposition.

PROPOSITION 3: *Conditional on their signals,*

- 1) *Minority borrowers at the bigoted bank will default less frequently than minority borrowers at the non-bigoted bank and less frequently than white borrowers at either bank;*
- 2) *If the non-bigoted bank discriminates against minorities ($s_{MN}^* > s_{WN}^*$), minority borrowers at that bank will default more frequently than comparable whites; and*
- 3) *If the non-bigoted bank discriminates against whites ($s_{MN}^* < s_{WN}^*$), minority borrowers at that bank who send low signals will default less frequently than comparable whites, while those minorities who send relatively high signals will default more frequently than comparable whites.*

Figure 2 illustrates the effects outlined in Proposition 3. First, as shown in Proposition 1, $q_{MB}(s)$ lies everywhere above $q_{MN}(s)$ and $q_{WN}(s)$, showing that for any given signal s minority borrowers at the bigoted bank default less frequently than all other borrowers. This result is consistent with the prediction made by Becker (1993).

Second, if the non-bigoted bank discriminates against minority applicants, $q_{MN}(s) < q_{WN}(s)$ for every s that is approved by the bank (see Panel A). As a result, minorities at this bank default *more* frequently than their white counterparts, providing ex post confirmation for the beliefs that led to the statistical discrimination. Thus, if lenders statistically discriminate, the relative default rates between minorities and whites are the opposite of those that arise if lenders discriminate based on preferences.

Finally, Panel B shows what happens when the non-bigoted bank discriminates against white applicants. As noted above, this happens only when the bigotry effect is so strong that few minorities (only the most creditworthy) apply. In this case, because the minority applicant pool is of higher average creditworthiness, minorities sending low signals are less likely to default (i.e., are more likely to have higher θ s). On the other hand, the fact that the least creditworthy

minorities self-select toward the non-bigoted bank means that $q_{MN}(s) < q_{WN}(s)$ for high s . As a result, the non-bigoted bank observes minorities with very good signals defaulting *more* frequently than comparable whites.

It is important to recognize that the default rates we consider here are “conditional” default rates and are distinct from the “marginal” default rates often discussed in the mortgage discrimination literature. Unfortunately, the concept of a marginal default rate is often poorly defined, making it unclear as to whether it is intended to refer to the least creditworthy applicant approved for a loan, or to some larger group of high-risk applicants. In any event, statistical tests of default rates typically attempt to control for factors that might affect a borrower’s credit risk, and are thus inherently conditional default rates. As a result, the predictions outlined in Proposition 3 are more directly testable than any predictions predicated on some notion of a marginal borrower.

Finally, note that Proposition 3 makes predictions about the relative default rates of minorities and whites across banks, but does not tell us anything about their conditional default rates in the aggregate. Because minorities default more frequently than whites at some banks and less frequently at others, it is unclear whether their default rate in the market as a whole will be higher or lower than that of whites. As a result, our model does not provide firm predictions as to what one might expect to find in a study such as Berkovec et al. (1994), which focuses on default rates in an entire mortgage market. If one could control for the identity of banks, however, our model suggests a means of testing (using denial rates) which banks may be exhibiting tastes for discrimination. We explore this and other practical applications of our model and its relevance to recent policy issues in the next section.

6. Implications for Testing and Policy

Detecting Discrimination in Credit Markets

An ongoing debate in the mortgage discrimination literature has been on the proper method for detecting discrimination in credit markets. For example, Munnell et al. (1996) implicitly condition on applicant signals to determine whether applicant denial decisions are

racially correlated. In contrast, Becker (1993) argues that the proper way to test for discrimination is to focus on (marginal) default rates and other measures of the relative profitability of the two groups, because a discriminatory bank will forgo profitable lending opportunities to minorities.¹⁹ Berkovec et al. (1994) attempt to perform such a test using FHA default and loss data, and find that low-creditworthy minority borrowers are substantially more likely to default than comparable white borrowers.

Our model suggests an emphasis on denial rates. This is because both bigoted banks and banks that statistically discriminate will deny minority applicants at a higher rate than white applicants, conditional on their signal. However, as outlined in Proposition 3, minorities will default more often than whites at non-bigoted banks that statistically discriminate against minorities. Thus, Becker's test would misinterpret the behavior of non-bigoted banks, since statistical discrimination arises precisely because minority applicants are less creditworthy conditional on their signals ($q_{MN}(s) < q_{WN}(s)$).²⁰ On the other hand, if the bigoted bank's taste for discrimination is so strong that non-bigoted banks do not discriminate against minority applicants, then focusing on conditional default rates could lead to the inaccurate conclusion that those banks do discriminate against minorities, and do so because of racial animus.

Furthermore, and perhaps most importantly, since default-rate predictions differ based on whether a bank is a statistical discriminator or a bigot, studies focusing on a market that contains both types of discriminators are likely to be inconclusive; the lower minority default rate at bigoted banks may be offset by higher minority default rates at banks that statistically discriminate.²¹ In contrast, because denial decisions are based directly on the variable we use to measure discrimination (the applicant's signal), tests for discrimination using conditional denial rate data do not suffer from such errors.²²

¹⁹ Han (1998) analyzes models of bigotry and statistical discrimination, with an emphasis on the differential impact on bank profits across racial groups. He concludes that the ability to detect discrimination, whether statistical or taste-based—as well as what discrimination “looks like”—depends on the particular measure of profits employed. He analyzes several profit measures, including default rates, loan write-offs, expected profits, and realized rates of return, finding some evidence that loans to minorities are more profitable than loans to whites.

²⁰ Although some may argue that banks should not be punished for statistically discriminating, such behavior is illegal under the Equal Credit Opportunity Act.

²¹ This fact may explain the ambiguous nature of some of Berkovec et al.'s (1994) results.

²² Other researchers have also argued that default rates may provide a misleading picture of discrimination. See, for

Uncovering the Source of Discrimination

A related issue is the ability of testing procedures to identify whether evidence of discrimination suggests statistical discrimination or bigotry. As we have argued above, both types of discrimination are illegal under current fair-lending laws, making the distinction between the two less relevant from an enforcement perspective. Nevertheless, designing effective anti-discrimination policies depends crucially on identifying the underlying source of this discrimination.²³

Our Bayesian model of the underwriting process suggests a simple approach that might be used to distinguish empirically between statistical discrimination and bigotry. Because the difference between these two forms of discrimination is located in the relationship between q_W^* and q_M^* , the empiricist must be able to reconstruct the lender's Bayesian updating process (the underwriting process) to arrive at each applicant's inferred quality, $q(s)$. A finding that minority applicants are more likely to be rejected than whites, *conditional on* $q(s)$, would be consistent with preference-based discrimination. On the other hand, if the likelihood of being rejected varies across racial groups controlling for s , but does not vary after controlling for each applicant's inferred quality $q(s)$, statistical discrimination is the more plausible explanation.

Munnell et al. (1996) undertake the first step in this process by looking at applicant denial rates conditional on credit history and other personal information found in the loan file (the applicant's signal), and find that minority applicants are more likely to be denied loans than comparable whites. It would be premature to conclude that the source of this discrimination is bigotry, however, because they do not consider how these signals are transformed into an inferred credit risk for each applicant. That is, they *do not control for each applicant's* $q(s)$.²⁴ Essentially, this requires controlling for the distribution of true credit risk of each group in the lender's applicant pool, $g(\theta)$. In separate work, we intend to pursue this empirical question, using the structure provided by this model to exploit the information available in the denial-rate data used by Munnell et al.

example, Tootell (1993), Ferguson and Peters (2000), and Han (1998).

²³ See Longhofer and Peters (1998).

²⁴ Instead, they show that $s_M^* - s$ differs from $s_W^* - s$ for applicants sending an "average" s .

Additional Policy Implications

Our model is also capable of shedding light on other recent developments in credit markets. For example, advances in the automation of loan processing and credit scoring have reduced the overall costs of applying for loans, both in terms of up-front fees and the time involved in the process. Similarly, efforts to comply with CRA guidelines, expanded marketing efforts (like those observed in credit card and consumer lending), and the internet, may have lowered the shoe-leather costs associated with applying for loans, particularly for targeted groups. Our model suggests that such a reduction in the cost of applying for a loan should increase the number of applicants. In other words, as the cost of applying for a loan declines, previously sub-marginal individuals will find it worthwhile to apply (those with θ just below θ^m). This is consistent with trends in the mortgage market and other consumer loan markets through the 1990s.

Our model illustrates that reductions in these costs increase both unconditional and conditional denial rates. First, because more low-creditworthy individuals apply, unconditional denial rates are higher, which is clearly borne out in the HMDA data.²⁵ Although this is not surprising, our model shows that these efforts to expand credit availability impose an externality on all individuals. *Conditional* denial rates increase because these new applicants reduce the average creditworthiness of the applicant pool, which in turn requires lenders to establish a higher cutoff signal s^* in order to maintain the same creditworthiness threshold q^* . Thus, while application or participation rates increase with a reduction in the cost of applying for a loan, reducing these costs also makes it tougher for *all applicants* to qualify for a loan (see Figure 3). This result is particularly noteworthy because application rates are argued to be one of the most reliable measures of CRA compliance (see Avery et al., 1998).

Finally, the general nature of our underwriting model makes it an ideal framework for analyzing policy initiatives in credit markets. For example, it could be used to examine how improved access to information about bank lending practices through bank regulators or

²⁵ For example, in 1990 national HMDA data show that the unconditional denial rate was 33.9 percent for blacks and 14.4 percent for whites. By 1998 these denial rates has risen to 53.7 and 26.0 percent, respectively, although these denial rates have inched back down over the last two years (to 44.6 and 22.3 percent in 2000).

community groups (a reduction in χ or an increase in γ) might affect individual application decisions and the subsequent treatment of applicants by banks. Similarly, our underwriting model might also form the basis for analyzing the effects of self-selection in the presence of other market frictions such as cultural affinity (Longhofer, 1996).

Markets with Risk-based Pricing

In this paper we have focused on how applicant self-selection behavior affects discriminatory incentives in a pooling market (e.g., one characterized by lenders offering a single-price to all approved applicants regardless of their inferred creditworthiness). As discussed earlier, many consumer loan markets—the mortgage market is a prime example—exhibit this single-price feature. In recent years, however, technological advances have made risk-based pricing more common. It is reasonable to ask, therefore, how our results would differ in a market with risk-based pricing.

Our model can be readily altered to accommodate a market in which lenders price applicants based on their inferred quality. In this case, all applicants are approved and offered a price $r(s) = r_c/q(s)$ based on the signal they send to the bank.²⁶ Similarly, banks that exhibit tastes for discrimination will charge a higher interest rate $r(s) = r_c/q(s) + \delta$ to members of the disfavored group. With this market structure, the “denial rate” instead reflects the proportion of applicants that choose not to accept the loan offered them because the interest rate is too high ($r(s) > B$).

Our preliminary investigations into this market structure suggest that applicant self-selection incentives can be dramatically different when lenders set interest rates based on each applicant’s individual inferred quality. In particular, when a subset of lenders indulge their bigoted preferences against minorities, it is the *high-quality* minority applicants that have the most incentive to avoid the bigoted banks, the opposite of what we show for pooling markets. This is because minority applicants with high θ s are more likely to be offered an interest rate

²⁶ One significant advantage of modeling the market in this way is that it imposes additional constraints that allow us to solve for the function $r(s)$ endogeneously. The cost, however, is the enormous complexity associated with dealing with this endogenous schedule of prices, which in turn affects application rate, and hence the function $q(s)$. As a result, it is virtually impossible to solve for closed-form comparative static results.

they will find acceptable. In contrast, low- θ minority applicants are prone to receiving unacceptably high interest rates regardless of whether they end up at the bigoted bank. As a consequence, they have less incentive to invest in information that might help them determine which is which.

Many of the market outcomes discussed above may differ in a market characterized by risk-based pricing. For example, when lenders price applicants based on their inferred quality, non-bigoted banks will no longer have an incentive to statistically discriminate against minorities. Indeed, the fact that high-quality minority applicants are attempting to self-select toward these banks will cause them to want to statistically discriminate against their *white* applicants, giving their minority applicants more favorable pricing terms to reflect their better average quality.²⁷ In contrast, applicant self-selection behavior will exacerbate bigoted banks incentives to discriminate, as the distribution of credit risk in their minority applicant pools will be lower.

If bank's act on these incentives, market-wide studies may be unable to detect discriminatory behavior even if taste-based discrimination is rampant. The reason is that this behavior will be masked in the aggregate by the *reverse discrimination* practiced by the non-bigoted banks (those for which the minority applicant pool is of above-average quality). Thus, even if all banks use race as a factor in pricing loans, statistical tests may be unable to uncover this unless the identity of bigoted and non-bigoted banks is known in advance.

Default rate analyses will also be affected by the use of risk-based pricing. As in the single-price market, $1 - q(s)$ represents the expected conditional default rate of applicants sending signal s . In this case, the fact that high- θ minorities self-select away from the bigoted bank means that $q_{MN}(s) \geq q_{MB}(s)$ for every s , so that conditional default rates will be higher at the bigoted bank.²⁸ Notice, however, that this happens solely because of applicant self-selection behavior, not because of bank accept-reject decisions. In such a market, default-rate-based tests

²⁷ Of course, this will only be strictly true if the underlying distribution of credit risk is identical across the two populations before self-selection actions.

²⁸ Once again, this assumes that the underlying distribution of credit is the same across the two groups in the general population.

for discrimination will misinterpret the action of both types of banks, falsely suggesting that non-bigoted banks indulge in taste-based discrimination, while allowing bigoted banks to discriminate with impunity. As before, default rates should only be used to help inform about the underlying source of discrimination, not about its existence.

Given that recent trends in mortgage and other consumer credit markets are toward the broader use of credit scoring models, risk based pricing may become a more prevalent feature in these markets. Our analysis suggests that this change could have a profound effect on the way in which discriminatory behavior is manifested in the market, and in the way applicants respond to it. Bank examiners and others following credit market trends will need to be conscious of these differences and modify their procedures accordingly as these changes occur.

7. Conclusion

We have developed a model in which individuals choose to apply for loans based on the likelihood of being approved. Individuals self-select in such a way that the distribution of creditworthiness in the applicant pool differs across racial groups even though creditworthiness is identically distributed across groups in the general population.

This self-selection behavior on the part of applicants may give lenders an incentive to statistically discriminate against minority applicants, *even if they had no proclivity to discriminate against them in the first place*. Thus, in addition to introducing the general concept of applicant self-selection into a formal model of loan underwriting, our theory is capable of explaining how the behavior of lenders with tastes for discrimination can lead other lenders to discriminate as well.

We have also shown that if lenders statistically discriminate against minorities, conditional default rates for minorities are higher than those for whites, the opposite of the result obtained if lenders discriminate based on preferences. This fact suggests that researchers should use caution when attempting to draw conclusions about the presence of market-wide discrimination from default-rate studies, since both statistical discriminators and bigots may coexist in a single market.

Finally, because of the general nature of our model of loan underwriting, it provides a useful framework for analyzing policy initiatives in credit markets. As discussed in the previous section, our model is amenable to modifications that would allow for the investigation of the potential impact of a variety of different market frictions and policy actions. In addressing such issues, an overriding concern should be the impact on α , an applicant's likelihood of obtaining a loan. This, ultimately, should be the primary variable of concern.

8. Appendix

Proof of Results 1-2: These proofs are relatively standard and are omitted for brevity. Interested readers may obtain a copy from the authors upon request.

Proof of Propositions 1: In order to prove these propositions, we will need the following lemmas.

LEMMA 1: *Assume there exists some θ^c such that all applicants with $\theta \in [\theta^m, \theta^c]$ attempt to self-select toward bank j while those with larger θ randomize between bank j and bank k . Then $q_k(s) > q_j(s), \forall s$.*

Proof of Lemma 1: If applicants self-select in this way, the fact that creditworthiness in the general population is uniformly distributed on the interval $[0,1]$ allows us to write the density of creditworthiness in the applicant pools as

$$g_j(\theta) = \begin{cases} \frac{\gamma}{\gamma(\theta^c - \theta^m) + \frac{1}{2}(1 - \theta^c)}, & \theta \in [\theta^m, \theta^c] \\ \frac{\frac{1}{2}}{\gamma(\theta^c - \theta^m) + \frac{1}{2}(1 - \theta^c)}, & \theta \in (\theta^c, 1] \end{cases} \quad (13)$$

and

$$g_k(\theta) = \begin{cases} \frac{1 - \gamma}{(1 - \gamma)(\theta^c - \theta^m) + \frac{1}{2}(1 - \theta^c)}, & \theta \in [\theta^m, \theta^c] \\ \frac{\frac{1}{2}}{(1 - \gamma)(\theta^c - \theta^m) + \frac{1}{2}(1 - \theta^c)}, & \theta \in (\theta^c, 1]. \end{cases} \quad (14)$$

This implies that

$$\begin{aligned}
 q_j(s) &= \int_{\theta^m}^{\theta^c} \theta \frac{\frac{p(s|\theta)\gamma}{\gamma(\theta^c - \theta^m) + \frac{1}{2}(1 - \theta^c)}}{\int_{\theta^m}^{\theta^c} \frac{p(s|\theta)\gamma}{\gamma(\theta^c - \theta^m) + \frac{1}{2}(1 - \theta^c)} d\theta + \int_{\theta^c}^1 \frac{p(s|\theta)\frac{1}{2}}{\gamma(\theta^c - \theta^m) + \frac{1}{2}(1 - \theta^c)} d\theta} d\theta \\
 &\quad + \int_{\theta^c}^1 \theta \frac{\frac{p(s|\theta)\frac{1}{2}}{\gamma(\theta^c - \theta^m) + \frac{1}{2}(1 - \theta^c)}}{\int_{\theta^m}^{\theta^c} \frac{p(s|\theta)\gamma}{\gamma(\theta^c - \theta^m) + \frac{1}{2}(1 - \theta^c)} d\theta + \int_{\theta^c}^1 \frac{p(s|\theta)\frac{1}{2}}{\gamma(\theta^c - \theta^m) + \frac{1}{2}(1 - \theta^c)} d\theta} d\theta \quad (13) \\
 &= \frac{1}{\int_{\theta^m}^{\theta^c} p(s|\theta)\gamma d\theta + \int_{\theta^c}^1 p(s|\theta)\frac{1}{2} d\theta} \left[\int_{\theta^m}^{\theta^c} \theta p(s|\theta)\gamma d\theta + \int_{\theta^c}^1 \theta p(s|\theta)\frac{1}{2} d\theta \right]
 \end{aligned}$$

and

$$q_k(s) = \frac{1}{\int_{\theta^m}^{\theta^c} p(s|\theta)(1 - \gamma) d\theta + \int_{\theta^c}^1 p(s|\theta)\frac{1}{2} d\theta} \left[\int_{\theta^m}^{\theta^c} \theta p(s|\theta)(1 - \gamma) d\theta + \int_{\theta^c}^1 \theta p(s|\theta)\frac{1}{2} d\theta \right]. \quad (14)$$

Using this, we calculate

$$\begin{aligned}
 q_j(s) - q_k(s) &\propto \left[\int_{\theta^m}^{\theta^c} p(s|\theta)(1 - \gamma) d\theta + \int_{\theta^c}^1 p(s|\theta)\frac{1}{2} d\theta \right] \left[\int_{\theta^m}^{\theta^c} \theta p(s|\theta)\gamma d\theta + \int_{\theta^c}^1 \theta p(s|\theta)\frac{1}{2} d\theta \right] \\
 &\quad - \left[\int_{\theta^m}^{\theta^c} p(s|\theta)\gamma d\theta + \int_{\theta^c}^1 p(s|\theta)\frac{1}{2} d\theta \right] \left[\int_{\theta^m}^{\theta^c} \theta p(s|\theta)(1 - \gamma) d\theta + \int_{\theta^c}^1 \theta p(s|\theta)\frac{1}{2} d\theta \right] \quad (15) \\
 &= (\gamma - \frac{1}{2}) \left[\int_{\theta^m}^{\theta^c} \theta p(s|\theta) d\theta \int_{\theta^c}^1 p(s|\theta) d\theta - \int_{\theta^m}^{\theta^c} p(s|\theta) d\theta \int_{\theta^c}^1 \theta p(s|\theta) d\theta \right],
 \end{aligned}$$

where \propto indicates that this expression is proportional to the difference $q_j(s) - q_k(s)$. Now,

$$\int_{\theta^m}^{\theta^c} \theta p(s|\theta) d\theta \int_{\theta^c}^1 p(s|\theta) d\theta < \int_{\theta^m}^{\theta^c} p(s|\theta) d\theta \int_{\theta^c}^1 \theta p(s|\theta) d\theta < \int_{\theta^m}^{\theta^c} p(s|\theta) d\theta \int_{\theta^c}^1 \theta p(s|\theta) d\theta. \quad (16)$$

Recalling that $\gamma > \frac{1}{2}$, this implies that (15) is negative, thus proving $q_k(s) > q_j(s)$ for all s . \spadesuit

LEMMA 2: Define $\beta_{jk}(\theta) \equiv (\gamma - \frac{1}{2})[\alpha_j(\theta) - \alpha_k(\theta)] - \chi/(B - r^*)$. Suppose that $s_k^* > s_j^*$ and assume that $\chi > 0$ is sufficiently small such that $\beta_{jk}(\theta) > 0$ for some θ . Then $\beta_{jk}(\theta)$ has exactly two roots on the real line, one of which must lie below θ^m .

Note that $\beta_{jk}(\theta)$ is derived from (12) in the text, the condition determining which applicants choose to self-select toward the bank with the lower s^* ; any applicants with θ such that

$\beta_{jk}(\theta) \geq 0$ will choose to self-select toward bank j , while the rest will not.

Proof of Lemma 2: First note that because $\gamma > 1/2$ and $s_k^* > s_j^*$, $(\gamma - 1/2)[\alpha_j(\theta) - \alpha_k(\theta)] > 0$ for all θ . Thus, \exists some χ such that $\beta_{jk}(\theta) > 0$ for some θ . Furthermore, by the symmetry of p , $\alpha(\theta) = P(\theta | s^*)$. As a result, $\lim_{\theta \rightarrow \infty} [\alpha_j(\theta) - \alpha_k(\theta)] = 0$ and similarly as $\theta \rightarrow -\infty$. This implies that, given any fixed $\chi > 0$, $\beta_{jk}(\theta) < 0$ for sufficiently large and sufficiently small θ . Consequently, β must have at least two roots on the real line. Next, note that

$$\beta'_{jk}(\theta) = p(\theta | s_j^*) - p(\theta | s_k^*). \quad (17)$$

By the symmetry and unimodality of p (parts 2 and 3 of A1), this is positive for any $\theta < (s_j^* + s_k^*)/2$ and negative for all $\theta > (s_j^* + s_k^*)/2$. Hence, $\beta_{jk}(\theta)$ is quasi-concave and bounded, reaching its maximum at $\theta = (s_j^* + s_k^*)/2$. Thus, $\beta_{jk}(\theta)$ must have exactly two roots.

It remains only to be shown that the lower root of $\beta_{jk}(\theta)$ lies below θ^m . Let $\underline{\theta}$ denote this lower root, and note that $\underline{\theta} < (s_j^* + s_k^*)/2$. By the symmetry of p (part 2 of A1),

$$1/2 \alpha_j\left(\frac{s_j^* + s_k^*}{2}\right) + 1/2 \alpha_k\left(\frac{s_j^* + s_k^*}{2}\right) = \frac{1}{2}. \quad (18)$$

Define $\tilde{\theta}$ such that

$$\gamma \alpha_j(\tilde{\theta}) + (1 - \gamma) \alpha_k(\tilde{\theta}) - \chi / (B - r^*) = \frac{1}{2} \quad (19)$$

(see Figure 4). By the continuity of α ,

$$\gamma \alpha_j(\tilde{\theta}) + (1 - \gamma) \alpha_k(\tilde{\theta}) - \chi / (B - r^*) > 1/2 \alpha_j(\tilde{\theta}) + 1/2 \alpha_k(\tilde{\theta}). \quad (20)$$

Thus, $\tilde{\theta} > \underline{\theta}$. Recall now that θ^m is defined as the θ such that

$$\max\{\gamma \alpha_j(\theta^m) + (1 - \gamma) \alpha_k(\theta^m) - \chi / (B - r^*), 1/2 \alpha_j(\theta^m) + 1/2 \alpha_k(\theta^m)\} = \alpha^*. \quad (21)$$

Since $\alpha^* > 1/2$ by assumption, the monotonicity of the left-hand side of (21) implies that $\theta^m > \tilde{\theta} > \underline{\theta}$. ♠

First we show that $s_{MB}^* > s_{MN}^*$. Assume to the contrary that $s_{MN}^* \geq s_{MB}^*$ and define

$$\beta_{NB}(\theta) \equiv (\gamma - 1/2)[\alpha_{MN}(\theta) - \alpha_{MB}(\theta)] - \chi / (B - r^*) \quad (22)$$

as the payoff to a type- θ minority applicant from attempting to self-select toward the non-bigoted bank (as in Lemma 2 above). This expression must be nonpositive, since $\gamma > 1/2$ and $\alpha_{MN}(\theta) \leq \alpha_{MB}(\theta)$ for all θ (because $s_{MN}^* \geq s_{MB}^*$). Thus, no minority applicant will want to self-

select toward the non-bigoted bank when $s_{MN}^* \geq s_{MB}^*$.²⁹

Now define

$$\beta_{BN}(\theta) \equiv (\gamma - \frac{1}{2})[\alpha_{MB}(\theta) - \alpha_{MN}(\theta)] - \chi / (B - r^*) \quad (23)$$

as the payoff to a type- θ minority applicant from attempting to self-select toward the *bigoted* bank. If χ is sufficiently large, β_{BN} will be negative for all θ and no minority applicant will attempt to self-select toward either bank. In this case, $q_{MN}(s) = q_{MB}(s)$ for all s and the monotonicity of q implies that $q_{MN}^* > q_{MB}^*$, a contradiction.

Now if χ is relatively small, Lemma 2 implies that β_{BN} will have two roots on the real line. Call the upper of these two roots θ^c and note that if $\theta^c \geq 1$, *all* minority applicants attempt to self-select toward the bigoted bank; once again $q_{MN}(s) = q_{MB}(s)$ for all s , implying the contradiction $q_{MN}^* > q_{MB}^*$.

By Lemma 2, the only remaining case to consider is $\theta^c \in (\theta_M^m, 1)$. In this case, Lemma 1 implies that $q_{MN}(s) > q_{MB}(s)$ for all s . This, together with our assumption that $s_{MN}^* > s_{MB}^*$ would again imply the contradiction $q_{MN}^* > q_{MB}^*$.

Next we prove that $s_{MB}^* > s_{WB}^* = s_{WN}^*$. Suppose to the contrary that $s_{MB}^* \leq s_{WB}^*$. In this case, the fact that $s_{MN}^* < s_{MB}^*$ implies that

$$\alpha_{WB}(\theta) < \frac{1}{2}\alpha_{MN}(\theta) + \frac{1}{2}\alpha_{MB}(\theta) < \gamma\alpha_{MN}(\theta) + (1 - \gamma)\alpha_{MB}(\theta) \quad (24)$$

for all θ (i.e., for any given θ , a minority applicant is more likely to be approved than a white applicant, regardless of whether he attempts to self-select toward the non-bigoted bank). This, however, implies that $\theta_M^m < \theta_W^m$, a contradiction. Finally, note that $s_{WB}^* = s_{WN}^*$ trivially because $q_{WB}^* = q_{WN}^*$ and $q_{WB}(s) = q_{WN}(s)$ for all s . Note that this proves the third part of the proposition.

Recall that $\beta_{NB}(\theta)$ is defined in expression (22) above as the payoff to a type- θ minority applicant from attempting to self-select toward the non-bigoted bank. The first term in this expression must be positive, since $\gamma > \frac{1}{2}$ and $\alpha_{MN}(\theta) < \alpha_{MB}(\theta)$ for all θ (because $s_{MB}^* > s_{MN}^*$ as shown above). Thus, for sufficiently large χ , β_{NB} will be negative for all $\theta \in [\theta_M^m, 1]$. For smaller values of χ , Lemma 2 implies that β_{NB} will have two roots on the real line, one of which

²⁹ It is worth noting that this fact rules out the possibility of a perverse signaling equilibrium, in which high- θ minority applicants attempt to self-select toward the bigoted bank in order to signal their quality.

is lower than θ_M^m . Call the upper of these two roots θ^c and note that if $\theta^c \geq 1$, *all* minority applicants attempt to self-select toward the non-bigoted bank. For moderate levels of χ , however, $\theta^c \in (\theta_M^m, 1)$ and Lemma 2 implies $\beta_{NB}(\theta) > 0$ only for $\theta \in [\theta_M^m, \theta^c]$. This proves the first part of the proposition.

Finally, note that $\theta_W^m < \theta_M^m$. If this were not the case, the same logic used in the proof Lemma 1 would imply that $q_{WN}(s) > q_{MN}(s)$ for all s . Because $q_{WN}^* = q_{MN}^*$ (the non-bigoted bank applies the same creditworthiness threshold to both minority and white applicants), this would require that $s_{WN}^* < s_{MN}^*$. Combining this with the fact that $s_{MB}^* > s_{MN}^*$ (proven above) would mean that

$$\alpha_{WN}(\theta) > \gamma\alpha_{MN}(\theta) + (1-\gamma)\alpha_{MB}(\theta) - \chi/(B-r^*), \quad \forall \theta. \quad (25)$$

This, however, would imply that $\theta_W^m < \theta_M^m$, a contradiction.

Using this fact ($\theta_W^m < \theta_M^m$), Lemma 1 implies that $q_{MB}(s) > q_{WB}(s) = q_{WN}(s)$, proving the second part of the proposition. ♠

Proof of Proposition 2: We prove existence by means of a numerical example. Equilibrium in our model with self-selection is given by the $(s_{MN}^*, s_{MB}^*, \theta_M^m, \theta^c)$ combination that solves the following system of equations:

$$\begin{aligned} q_{MN}(s_{MN}^*) &\equiv \frac{\int_{\theta_M^m}^{\theta^c} \theta p(s_{MN}^* | \theta) \gamma d\theta + \int_{\theta^c}^1 \theta p(s_{MN}^* | \theta) \frac{1}{2} d\theta}{\int_{\theta_M^m}^{\theta^c} p(s_{MN}^* | \theta) \gamma d\theta + \int_{\theta^c}^1 p(s_{MN}^* | \theta) \frac{1}{2} d\theta} = q_{MN}^* \\ q_{MB}(s_{MB}^*) &\equiv \frac{\int_{\theta_M^m}^{\theta^c} \theta p(s_{MB}^* | \theta) (1-\gamma) d\theta + \int_{\theta^c}^1 \theta p(s_{MB}^* | \theta) \frac{1}{2} d\theta}{\int_{\theta_M^m}^{\theta^c} p(s_{MB}^* | \theta) (1-\gamma) d\theta + \int_{\theta^c}^1 p(s_{MB}^* | \theta) \frac{1}{2} d\theta} = q_{MB}^* \\ \hat{\alpha}(\theta_M^m) &\equiv \gamma\alpha_{MN}(\theta_M^m) + (1-\gamma)\alpha_{MB}(\theta_M^m) - \frac{\chi}{B-r^*} = \alpha^* \\ \beta(\theta^c) &\equiv (\gamma - \frac{1}{2})[\alpha_X(\theta^c) - \alpha_Y(\theta^c)] - \frac{\chi}{B-r^*} = 0; \end{aligned} \quad (26)$$

as before, s_W^* and θ_W^m are determined by (8) in the text.

Let $p(s|\theta)$ be a normal density with mean θ and variance .1, and assume that $q_{WN}^* = q_{MN}^* = .575$, $q_{MB}^* = .6$, $\alpha^* = .6$, $\chi = .01$, $B-r^* = 1$, and $\gamma = .85$. Numerical calculation of


(26) and (8) using these parameter values gives us the following solution: $s_{MB}^* = .493$, $s_{MN}^* = .485$, $s_{WN}^* = .478$, $\theta_W^m = .504$, $\theta_M^m = .514$, and $\theta^c = .544$. ♠

Proof of Proposition 3: Follows immediately from Proposition 1 and the discussion in the text. ♠

9. Figures


Figure 1

Panel A
Discrimination motivated by bigotry.



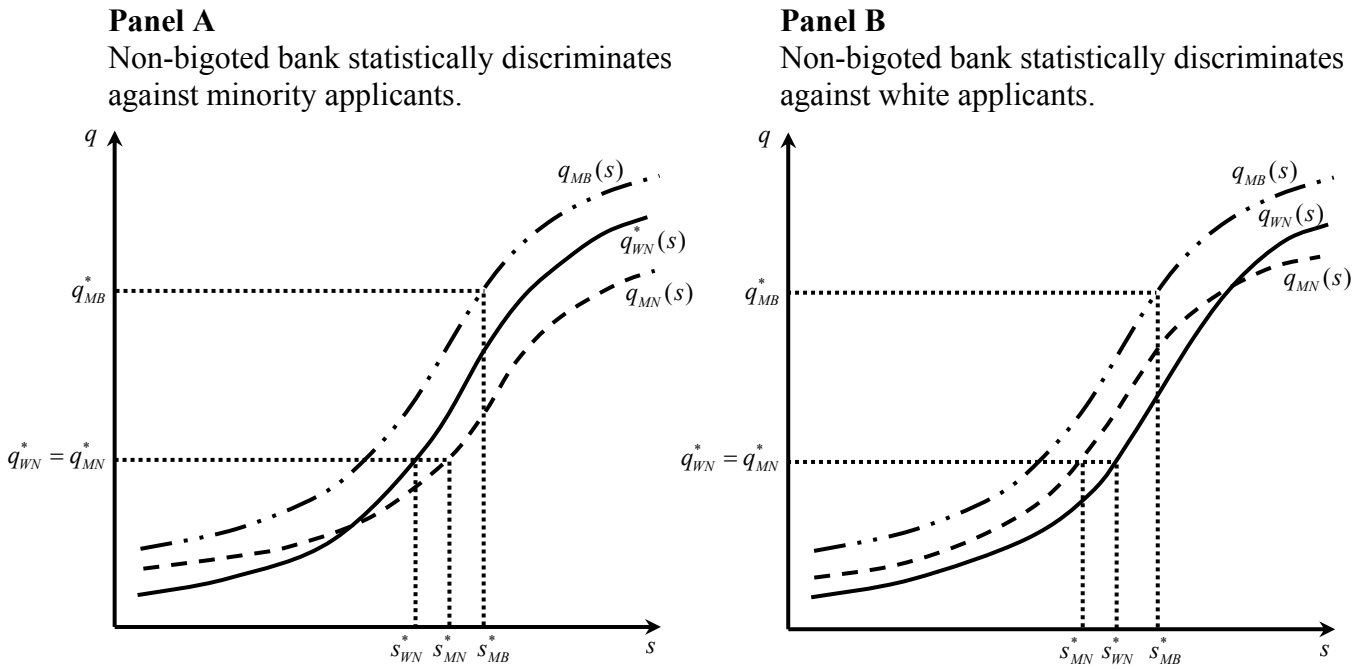
The higher q^* required by a bigoted lender leads it to discriminate against minority applicants by holding them to a higher cutoff signal.

Panel B
Statistical discrimination.



The lower average creditworthiness of the minority applicant pool lowers the bank's estimate of a minority applicant's inferred quality for any given signal s . This leads the bank to statistically discriminate against minority applicants by holding them to a higher cutoff signal, even though it has no taste for discrimination.

Figure 2




Minority self-selection behavior lowers the average creditworthiness in the non-bigoted bank's minority applicant pool (i.e., shifts $q_{MN}(s)$ down), giving this bank an incentive to statistically discriminate against minorities. Because $1 - q(s)$ is also a measure of the ex post conditional default rate, the non-bigoted bank observes higher conditional default rates from minority borrowers than it does from its white borrowers.

If enough minority individuals drop out of the applicant pool, it raises their average creditworthiness at the non-bigoted bank, giving it an incentive to statistically discriminate against *white* applicants. In this case, minority borrowers that send relatively low signals default more frequently than their white counterparts, while minority borrowers that send very good signals default *less* frequently than comparable whites.

Figure 3


Effect of a reduction in mortgage application costs.



Reducing the cost of applying for a mortgage has the effect of reducing α^* to α_2^* . This, in turn, lowers the creditworthiness of the marginal applicant to θ_2^m . The resulting lower average creditworthiness in the mortgage applicant pool causes lenders to *raise* their required cutoff signal to $s_2^* > s_1^*$, reducing all applicants' probabilities of being approved for a loan.

Figure 4

Proof of Lemma 2.



$\underline{\theta}$ must lie below θ^m because, by the monotonicity of α , it must lie below $\tilde{\theta}$, which is smaller than θ^m by the continuity of α .

10. References

- Arrow, Kenneth J. (1972a). "Models of Job Discrimination," in *Racial Discrimination in Economic Life*, Anthony H. Pascal, ed. Lexington, Mass.: Lexington Books, 1972, pp. 83-102.
- Arrow, Kenneth J. (1972b). "Some Mathematical Models of Race Discrimination in the Labor Market," in *Racial Discrimination in Economic Life*, Anthony H. Pascal, ed. Lexington, Mass.: Lexington Books, 1972, pp. 187-203.
- Arrow, Kenneth J. "The Theory of Discrimination," in *Discrimination in Labor Markets*, Orley Ashenfelter and Albert Rees, eds. Princeton, NJ: Princeton University Press, 1973, pp. 3-42.
- Avery, Robert B., Patricia E. Beeson, and Mark S. Sinderman. "Cross-Lender Variation in Home Mortgage Lending." Federal Reserve Bank of Cleveland, *Economic Review*, vol. 30, no. 4, 4th Quarter 1994, pp. 15-29.
- Avery, Robert B., Patricia E. Beeson, and Mark S. Sinderman. "Posted Rates and Mortgage Lending Activity." *Journal of Real Estate Finance and Economics*, vol. 13, no. 1, July 1996, pp. 11-26.
- Avery, Robert B., Patricia E. Beeson, and Mark S. Sinderman. "Neighborhood, Race and the Mortgage Market: Evidence from HMDA Data." Federal Reserve Bank of Cleveland, unpublished manuscript, 1998.
- Ayres, Ian and Peter Siegelman. "Race and Gender discrimination in Bargaining for a New Car." *American Economic Review*, vol. 85, no. 3, June 1995, pp. 304-321.
- Becker, Gary S. *The Economics of Discrimination*, 2nd ed. Chicago: University of Chicago Press, 1971.
- Becker, Gary S. "Nobel Lecture: The Economic Way of Looking at Behavior." *Journal of Political Economy*, 1993, 385-409.
- Benjamin, J.D., Andrea J. Heuson, and C.F. Sirmans. "The Effect of Origination Strategies on the Pricing of Fixed-Rate Mortgage Loans." *Journal of Housing Research*, vol. 6, 1995, pp. 137-148.
- Berkovec, James A., Glenn B. Canner, Stuart A. Gabriel, and Timothy H. Hannan. "Race, Redlining, and Residential Mortgage Loan Performance." *Journal of Real Estate Finance and Economics*, vol. 9, no. 3, 1994, pp. 263-94.
- Black, Harold A., M. Cary Collins, and Ken B. Cyree. "Do Black-Owned Banks Discriminate Against Black Borrowers?" *Journal of Financial Services Research*, vol. 11, nos. 1/2, February/April 1997, pp. 189-204.
- Bostic, Raphael W. "The Role of Race in Mortgage Lending: Revisiting the Boston Fed Study." Board of Governors of the Federal Reserve System, unpublished manuscript, December 1996.

- Bostic, Raphael W. and Glenn B. Canner. "Cultural Affinity Mechanisms and Variation in Mortgage Lending." Board of Governors of the Federal Reserve System, unpublished manuscript, 1998.
- Calem, Paul and Michael Stutzer. "The Simple Analytics of Observed Discrimination in Credit Markets." *Journal of Financial Intermediation*, vol. 4, no. 3, July 1995, pp. 189-212.
- Calomiris, Charles W., Charles M. Kahn, and Stanley D. Longhofer. "Housing-Finance Intervention and Private Incentives: Helping Minorities and the Poor." *Journal of Money, Credit, and Banking*, vol. 26, no. 3, August 1994, part 2, pp. 634-674.
- Cavalluzzo, Ken S. and Linda C. Cavalluzzo. "Market Structure and Discrimination: The Case of Small Business Lending." *Journal of Money Credit and Banking*, vol. 30, no. 4, November 1998, pp. 771-792.
- Cornell, Bradford and Ivo Welch. "Culture, Information, and Screening Discrimination." *Journal of Political Economy*, vol. 104, no. 3, June 1996, pp. 542-571.
- Day, Ted and Stan J. Leibowitz. "Mortgage Lending to Minorities: Where's the Bias." *Economic Inquiry*, vol. 34, January 1998, pp. 3-28.
- Duca, John V. and Stuart S. Rosenthal. "Do Mortgage Rates Vary Based on Household Default Characteristics? Evidence on Rate Sorting and Credit Rationing." *Journal of Real Estate Finance and Economics*, vol. 8, 1994, pp. 99-113.
- Ferguson, Michael F. and Stephen R. Peters. "What Constitutes Evidence of Discrimination in Lending." *Journal of Finance*, vol. 50, no. 2, June 1995, pp. 739-748.
- Ferguson, Michael F. and Stephen R. Peters. "Cultural Affinity and Lending Discrimination: The Impact of Underwriting Errors and Credit Risk Distribution on Applicant Denial Rates." *Journal of Financial Services Research*, vol. 11, nos. 1/2, February/April 1997, pp. 153-168.
- Ferguson, Michael F. and Stephen R. Peters. "Is Lending Discrimination Always Costly?" *Journal of Real Estate Finance and Economics*, vol. 21, no. 1, July 2000, pp. 23-44.
- Han, Song. "Credit Risk, Discrimination, and Regulation in Credit Markets." Ph.D. Dissertation, University of Chicago, November 1998.
- Harrington, Scott E. and Greg Niehaus. "Race, Redlining, and Automobile Insurance Prices." *Journal of Business*, vol. 71, no. 3, July 1998, pp. 439-470.
- Heckman, James J. "Detecting Discrimination." *Journal of Economic Perspectives*, vol. 12, no. 2, Spring 1998, pp. 101-116.
- Horne, David K. "Evaluating the Role of Race in Mortgage Lending." *FDIC Banking Review*, vol. 7, no. 1, Spring/Summer 1994, pp. 1-15.
- Horne, David K. "Mortgage Lending, Race, and Model Specification." *Journal of Financial Services Research*, vol. 11, nos. 1/2, February/April 1997, pp. 43-68.
- Hunter, William C. and Mary Beth Walker. "The Cultural Gap Hypothesis and Mortgage Lending Decisions." *Journal of Real Estate Finance and Economics*, vol. 13, no. 1, July 1996, pp. 57-70.

- Longhofer, Stanley D. "Cultural Affinity and Mortgage Discrimination." Federal Reserve Bank of Cleveland, *Economic Review*, vol. 32, no. 3, 3rd Quarter 1996, pp. 12-24.
- Longhofer, Stanley D. and Stephen R. Peters. "Beneath the Rhetoric: Clarifying the Debate on Mortgage Lending Discrimination." Federal Reserve Bank of Cleveland, *Economic Review*, 4th Quarter, 1998, pp. 2-19.
- Munnell, Alicia H., Lynn E. Browne, James McEneaney, and Geoffrey M.B. Tootell. "Mortgage Lending in Boston: Interpreting HMDA Data." Federal Reserve Bank of Boston Working Paper No. 92-7, October 1992.
- Munnell, Alicia H., Geoffrey M.B. Tootell, Lynn E. Browne, and James McEneaney. "Mortgage Lending in Boston: Interpreting HMDA Data." *American Economic Review*, vol. 86, no. 1, March 1996, pp. 25-53.
- Phelps, Edmund. "The Statistical Theory of Racism and Sexism." *American Economic Review*, vol. 62, no. 4, September 1972, pp. 659-661.
- Rosenblatt, Eric. "A Reconsideration of Discrimination in Mortgage Underwriting with Data from a National Mortgage Bank." *Journal of Financial Services Research*, vol. 11, nos. 1/2, February/April 1997, pp. 109-131.
- Ross, Stephen L. "Mortgage Lending, Sample Selection, and Default." *Real Estate Economics*, vol. 28, no. 4, Winter 2000, pp. 581-621.
- Rothschild, Michael and Joseph E. Stiglitz. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics*, vol. 90, no. 4, November 1976, pp. 630-49.
- Sirmans, C.F. and J.D. Benjamin. "Pricing Fixed-Rate Mortgages: Some Empirical Evidence." *Journal of Financial Services Research*, vol. 4, 1990, pp. 191-202.
- Tootell, Geoffrey M.B. "Defaults, Denials, and Discrimination in Mortgage Lending." *New England Economic Review*, September/October 1993, pp. 45-51.
- Yezer, Anthony M.J., Robert F. Phillips, and Robert P. Trost. "Bias in Estimates of Discrimination and Default in Mortgage Lending: The Effects of Simultaneity and Self-Selection." *Journal of Real Estate Finance and Economics*, vol. 9, no. 3, 1994, pp. 196-215.