http://clevelandfed.org/research/review
Economic Review 1996 Q3

Cultural Affinity and Mortgage Discrimination

by Stanley D. Longhofer

Stanley D. Longhofer is an economist at the Federal Reserve Bank of Cleveland. The author thanks Paul Bauer, Paul Calem, Charles Carlstrom, Ben Craig, Stephen Peters, João Santos, and Mark Schweitzer for helpful comments and discussions.

Introduction

In October 1992, researchers at the Federal Reserve Bank of Boston released their groundbreaking study on mortgage lending patterns in that area.[1] They found that black and Hispanic applicants were over 50 percent more likely to be denied mortgage loans than comparable whites, even after accounting for such factors as loan-to-value ratios, obligation ratios, and certain credit-history variables. In the end, they concluded that this disparity resulted from widespread, systematic discrimination in the Boston-area mortgage market. Although this study's validity has been hotly debated, since its publication a variety of theories have been developed to explain how such discrimination might persist in a market so many view as being highly competitive.[2]

This article reviews and expands one prominent theoretical source of discrimination in the residential mortgage market: the cultural affinity hypothesis proposed by Calomiris, Kahn, and Longhofer (1994; hereafter CKL). This theory argues that lenders find it easier (or less costly) to evaluate the creditworthiness of applicants with whom they have a common experiential background, or "cultural affinity." As a result, CKL contend that lenders make more mistakes when evaluating minority applicants, which gives them an incentive to discriminate against such applicants.

The CKL depiction of the mortgage market is based on the idea that lenders find it easier to assess the true creditworthiness of applicants with whom they have an affinity. This affinity may arise because the applicant and the loan officer share a common cultural background or because the lender has developed specialized expertise in evaluating the creditworthiness of members of a particular group. Thus, a lender's affinity may be considered inherent, learned, or both.[3] Regardless of its source, having an

1 Munnell, Browne, McEneaney, and Tootell (1992); hereafter Munnell et al. This paper was revised and published in the *American Economic Review* (Munnell, Tootell, Browne, and McEneaney [1996]).

2 For a sampling of criticisms of Munnell et al., see Day and Liebowitz (1993), Horne (1994), and Yezer, Phillips, and Trost (1994). Browne and Tootell (1995) provide a rebuttal.

3 Ferguson and Peters (forthcoming) emphasize the idea that cultural affinities may arise endogenously because of a lender's underwriting activity. Furthermore, lenders may choose to actively develop such affinities; an example of this specialized expertise would be a community development bank targeting low-income and minority neighborhoods. Conversely, some lenders may develop an affinity for white suburbanites simply because most of their applications come from such individuals.

affinity with a group enables a lender to gather additional information about the true creditworthiness of that group's members.

The traditional theory of taste-based discrimination was developed by Nobel laureate Gary Becker (1971). He argues that individuals discriminate against minorities for the same reasons they discriminate between products in the marketplace—personal preferences. Translated to the mortgage market, this means that rather than being "profit maximizers," bigoted lenders are "utility maximizers" who are willing to sacrifice profits in order to satisfy their "tastes for discrimination." They accomplish this by forgoing some marginally profitable loans to members of groups that they dislike.[4] In other words, bigoted lenders would require that members of disfavored groups meet a higher cutoff standard in order to be approved for loans.

Although the results of Munnell et al. may suggest that taste-based discrimination is a problem in the mortgage market, evidence on default rates contradicts this conclusion. In particular, if taste-based discrimination were prevalent in the home mortgage market, marginally qualified minority borrowers would default *less* frequently than their white counterparts.[5] Berkovec et al. (1994), however, analyze the performance of FHA mortgage loans and show that the opposite is true: Even after controlling for other factors associated with credit risk, black borrowers default significantly more often than their white counterparts.[6]

CKL's cultural affinity hypothesis is important because it helps reconcile the "Becker Paradox" posed by the seemingly inconsistent results of Munnell et al. and Berkovec et al. In particular, if lenders have an affinity with white applicants, minority applicants will be held to a higher cutoff standard and be denied loans more frequently.[7] In addition, minority applicants who are actually approved will tend to default more frequently on average than whites (although the default rate of the marginal applicants will be the same for both groups). If regulators insist that lenders treat all applicants the same, the average default rates for the two groups will diverge even further, with marginal minority applicants (that is, the least creditworthy applicants who are approved) defaulting more frequently than their white counterparts.

CKL's cultural affinity hypothesis is able to reconcile the seemingly inconsistent results of Munnell et al. and Berkovec et al., but their analysis has at least two limitations. First, they assume that lenders reject a majority of all applications they receive. Given that denial rates over the last several years have ranged between 15 and 40 percent, it is reasonable to question this assumption. One might interpret their analysis as beginning after some initial screen through which clearly qualified applicants are approved and obviously uncreditworthy applicants are weeded out.[8] Alternately, one might interpret CKL's model as focusing on the percentage of the *entire* population that receives loans, not just those who actually apply. Under either interpretation, CKL's assumption that most applicants are rejected might be more reasonable. Unfortunately, the deficiencies of existing data on denials make it nearly impossible to test any of the model's empirical predictions under either of these stories.

My analysis allows for either possibility, but concentrates on the assumption that lenders do, in fact, approve a majority of applications they receive. The primary by-product of this assumption is that lenders now have an incentive to discriminate against groups with whom they have an affinity, typically white applicants. This happens because the added information lenders receive about the quality of these applicants allows them not only to approve some who were previously deemed uncreditworthy, but also to detect and weed out uncreditworthy applicants who would otherwise have been accepted.

While this result may initially seem counterintuitive, I argue that it actually reflects lender behavior in the mortgage market more accurately than does CKL's original analysis. First, lenders may not act symmetrically on new information they receive about white applicants, either because they can sell loans to the secondary market, or because such information is

4 It is sometimes argued that bigoted lenders might deny random applications, rather than selecting the least profitable members of the group they dislike. Although randomization is certainly possible, utility maximization would still require that they deny less profitable applicants more frequently than more profitable ones.

5 See Becker (1993) for a discussion of this idea.

6 Berkovec et al. (1994, p. 282) note, "For example, in the 1987 cohort, black borrowers are predicted to have cumulative default rates that are about two percentage points higher than white borrowers, all else equal." The competitiveness of the mortgage market also makes a conclusion of taste-based discrimination problematic. Longhofer (1995) provides a brief discussion.

7 CKL allow for the possibility that lenders can develop a screening technology for minorities that is equal to the one they use for white applicants. The results discussed here relate to the case where lenders choose not to invest in that technology because it is too costly.

8 Calem and Stutzer (1995) make this assumption in their model of statistical discrimination.

costly to obtain. Second, even if lenders do wish to discriminate against white applicants, these effects may be outweighed by the fact that the minority applicant pool is less creditworthy on average than the white one.

A second limitation of CKL's analysis is its assumption that the signals lenders receive from both applicant groups are directly comparable and that outside observers can objectively measure them. Much of the premise behind the cultural affinity hypothesis, however, depends on the fact that many indicators of creditworthiness are subjective and fully observable only by lenders. I therefore extend CKL's analysis to consider the consequences of this feature of mortgage underwriting, showing how it complicates efforts to detect discrimination.

In addition to correcting these two shortcomings of the model, the present analysis provides at least three important extensions to the cultural affinity hypothesis. First, it clarifies the process by which lenders update their prior beliefs about an applicant's creditworthiness, thus making the mortgage underwriting process more transparent.[9] Second, it allows for differences in average creditworthiness across races and analyzes the impact this might have on the theory's empirical predictions.[10] Finally, it examines how other features of the mortgage market, such as the secondary market and minority-owned lenders, might interact with cultural affinity to affect outcomes in the model.

In the next section, I develop a simple model of the mortgage underwriting process in which lenders have an affinity with members of one applicant group. I use this model to show how cultural affinity problems can affect relative denial rates across groups and create incentives to discriminate. I then extend the model, allowing groups to differ in their average creditworthiness and allowing lenders to sell their loans on the secondary market. In section II, I tie all of these results together, discussing some of the model's empirical and policy implications.

I. A Model of Mortgage Underwriting

Consider a world in which individuals want to buy a house but, lacking sufficient cash to do so, must obtain a mortgage loan. I assume that individuals are divided into two groups, W and M; when the variables or density functions introduced below differ between the two groups, I will use subscripts to denote this difference. In practice, we can think of these groups as identifying members of different races, genders, or other protected classes. Policymakers may wish to know how the distribution of inferred creditworthiness, the denial rate, or the likelihood of default differs between these two applicant groups. Most important, however, they would like to know if either group is discriminated against. To analyze these issues, I assume that lenders have an affinity with members of group W, the effects of which I will describe in a moment.

Suppose that each applicant has a true creditworthiness θ, which can be thought of as his or her probability of repaying a loan. In particular, this parameter is assumed to capture all factors that lead to default, including possible income disruptions, the value of the house being purchased, and the borrower's personal compunction about defaulting on an obligation. Based on its cost of funds and the competitive market interest rate it charges, each lender has a cutoff creditworthiness θ^* that defines which applications it will approve or reject.[11] That is, applicants with creditworthiness below θ^* will

9 For a similar approach, see Cornell and Welch (1996).

10 In this respect, the present analysis incorporates Ferguson and Peters' (forthcoming) extensions of the cultural affinity hypothesis. They show that cultural affinity problems, when combined with differences in average creditworthiness among races, lead to ambiguous implications about denial rates across racial groups.

11 I accept as a stylized fact that lenders offer a single interest rate to all applicants, rejecting those who are not profitable at that interest rate. Yet, it is worth asking why lenders do not accept all applicants and "price for risk." One explanation appeals to antidiscrimination laws: Since creditworthiness is correlated with race, it would rapidly become apparent that lenders charge minorities higher interest rates. Alternatively, CKL argue that if lenders have better information about an applicant's true creditworthiness than does the applicant, risk-averse borrowers prefer a single offered interest rate with a commitment to lend to all who qualify at that rate. Finally, Ferguson and Peters (1996) show how "portfolio effects" can make risk-based pricing suboptimal for lenders; in their model, the gains from making more loans offset the losses from loans to the least creditworthy borrowers who are approved. Under any of these justifications for offering a single rate to all applicants, lenders would make loans only to applicants who are sufficiently creditworthy.

Although individual lenders seem to announce a single rate to prospective applicants, the mortgage market as a whole does appear to apply risk-based pricing. By requiring deposit insurance for lenders with high loan-to-value ratios, lenders implicitly demand higher interest rates from riskier borrowers. Furthermore, although individual lenders do not appear to price for other factors associated with risk, Avery, Beeson, and Sniderman (1996) show that individual lenders sort themselves by selecting the rate–risk combination with which they are most comfortable. There is even a secondary market for so-called "B" and "C" loans, which are loans that fail to meet the underwriting guidelines established by Fannie Mae and Freddie Mac. Thus, although borrowers may not be able to obtain a full menu of prices from any one lender, they can nevertheless obtain a rate commensurate with their personal risk.

be rejected, while those who are more creditworthy will be accepted.[12]

Unfortunately, lenders cannot perfectly observe an individual's true creditworthiness. Instead, they observe a signal, s_1, that is correlated with θ. For example, lenders typically collect information about the applicant's property, loan-to-value ratio, obligation ratios, credit history, income, employment, and so forth. While this information can never predict default perfectly, it does allow a lender to infer the likelihood of this event. This signal is observed for members of both groups and can be verified by outsiders (such as regulators).

I model cultural affinity by assuming that lenders receive a second, private, signal (s_2) for group W applicants that they do not receive for group M applicants. As a result, lenders have more information with which to assess the creditworthiness of group W applicants than they have for group M applicants. Intuitively, we can think of the signal s_2 as encompassing any subjective information beyond the standard underwriting variables that lenders gather during the application process. Such information is often referred to as "compensating factors." For example, an applicant may provide information that explains a past default or job instability. Or a lender may be willing to make a "character loan" because he "knows" the borrower is a good credit risk. Alternatively, the loan interview may give the lender new information suggesting that an applicant really is a bad credit risk.[13] I am assuming that s_2 incorporates only these subjective compensating factors.[14]

Since the information that lenders receive about an applicant's creditworthiness depends on his group, this difference may give lenders an incentive to discriminate against members of one group or the other. I now analyze the lender's underwriting decision, starting with the more simple case of group M applicants.

Group *M* Underwriting

For simplicity, assume that applicant creditworthiness is distributed normally with mean $\bar{\theta}$ and variance σ_θ^2; denote the respective probability density function as $f(\theta)$. For now, I assume that this distribution is the same for both groups; this assumption will be relaxed later. Let $p(s_1|\theta)$ be the probability that the bank observes signal s_1 from an individual of type θ; assume that p is a normal density with mean θ and variance σ_s^2.[15]

Using this information, we see that

$$(1) \qquad \omega_M(s_1) = \int p(s_1|\theta) f(\theta) d\theta$$

is the density of signals received by a lender observing the entire population of group M applicants; DeGroot (1989, p. 304) shows that this distribution is normal with mean $\bar{\theta}$ and variance $\sigma_\theta^2 + \sigma_s^2$. Using this, we define

$$(2) \qquad \pi_M(\theta|s_1) \equiv \frac{p(s_1|\theta) f(\theta)}{\omega_M(s_1)}$$

as the likelihood that a group M applicant with signal s_1 has a true creditworthiness of θ.

Ultimately, lenders are interested not in the signal the applicant sends, but rather in the applicant's inferred "quality" given this signal. Under this setup, the inferred quality of any group M applicant is simply the expectation of $\pi_M(\theta|s)$, which Hogg and Craig (1978, p. 232) show to be[16]

$$(3) \qquad q_M(s_1) = \frac{s_1\sigma_\theta^2 + \bar{\theta}\sigma_S^2}{\sigma_\theta^2 + \sigma_S^2}.$$

Thus, an applicant's inferred quality is simply the weighted average of his signal and the average creditworthiness of the applicant pool; the weights are based on the precision of each of

■ **12** Clearly, lenders consider more than just the likelihood of default when evaluating a mortgage application. A more complete analysis would also incorporate such factors as the likelihood of delinquency (and the associated costs and fees) and any cross-selling profits (from credit card or consumer loans) the lender might earn. Such additions should not have any qualitative effects on my conclusions.

■ **13** Campbell and Dietrich (1983) provide evidence of adverse selection problems in the market for mortgage insurance around the point of an 80 percent loan-to-value ratio. Because lenders typically do not require applicants to purchase mortgage insurance when the loan-to-value ratio is below 80 percent, the fact that they require some such borrowers to obtain this insurance and that these borrowers default more often suggests that lenders do observe—and act on—negative information about applicants.

■ **14** It is possible to extend the model to allow these signals to be observed sequentially, with the lender deciding whether to invest in the second signal only after observing the value of the first. If the marginal cost of using the second signal is zero, the results are formally identical to what I derive below. If it is costly to obtain additional information about an applicant, then it is feasible that lenders will not wish to acquire this information for some applicants whose initial signal is particularly good or particularly bad. Nonetheless, I believe that my qualitative results would continue to hold in such an environment. If the cost of obtaining the second signal depends on the value of s_1, however, the results will differ.

■ **15** More generally, one might think of s_1 as the "sample mean" of n draws from this distribution, with each draw representing one piece of information collected by the loan officer.

■ **16** This $q(s)$ is analogous to the $p(s)$ used by CKL.

these factors. If mortgage applicants vary greatly in their underlying creditworthiness (a high σ_θ^2), lenders tend to discount the characteristics of the applicant population, instead relying more heavily on the individual's signal. On the other hand, if the signals are quite imprecise (a high σ_S^2), lenders are more likely to ignore this signal and treat all applicants alike by placing more weight on the population's average quality, $\bar{\theta}$. Based on this signal, a group M applicant will be accepted if his $q_M(s_1) \geq q^* = \theta^*$.

It will later be useful to know the distribution of inferred quality of the applicant pool. Since $q_M(s_1)$ is simply a linear function of s_1, "quality" in the group M applicant population is normally distributed with mean $\bar{\theta}$ and variance $\sigma_\theta^4/(\sigma_\theta^2 + \sigma_S^2)$.

Group *W* Underwriting

The decision to grant credit to group W applicants is essentially the same as for group M applicants, except that the lender receives the additional information, s_2, with which to make the decision. Like the first signal, s_2 is a random draw from a distribution with mean θ (the applicant's true creditworthiness) and variance σ_S^2.[17] I will denote the combined signal as $\bar{s} = (s_1 + s_2)/2$. For an individual applicant with underlying creditworthiness θ, $\bar{s}$ is distributed normally with mean θ and variance $\sigma_S^2/2$. Thus, distribution of $\bar{s}$ in the group W applicant population is normal, with mean $\bar{\theta}$ and variance $\sigma_\theta^2 + \sigma_S^2/2$.

The fact that this combined signal is more precisely centered around the applicant's true creditworthiness allows lenders to make a better estimate of a group W applicant's quality:

$$(4) \qquad q_W(\bar{s}) = \frac{\bar{s}\,\sigma_\theta^2 + \bar{\theta}\sigma_S^2/2}{\sigma_\theta^2 + \sigma_S^2/2}.$$

Here we see that a group W applicant's inferred creditworthiness is again simply the weighted average of his signal and the average creditworthiness of the applicant pool. With the information added by the second signal, however, relatively less weight is placed on the applicant's group membership and relatively more on his individual signal.

As before, $q_W(\bar{s})$ is normally distributed with mean $\bar{\theta}$. The added information, however, increases the variance of this distribution to

$$(5) \qquad \frac{\sigma_\theta^4}{\sigma_\theta^2 + \sigma_S^2/2},$$

meaning that the distribution of inferred quality for group W applicants more closely resembles the true distribution of creditworthiness in the applicant pool than does the corresponding distribution for group M. Recall that the variance of $q_M(s_1)$ is $\sigma_\theta^4/(\sigma_\theta^2 + \sigma_S^2)$.

The final decision about whether to approve the loan is the same as it was before. Given the final value of the signal $\bar{s}$, the lender evaluates a group W applicant's inferred creditworthiness and approves the loan only if $q_W(\bar{s}) \geq q^*$.

Cultural Affinity and Discrimination

We are interested in analyzing how a lender's affinity with members of group W might affect its incentive to discriminate against members of one group or the other and, regardless of whether the lender discriminates, what impact affinity has on the relative denial rates of the two groups.

To do this, it is necessary to define what is meant by "discrimination." Simply stated, lenders discriminate if the cutoff signal, s^*, differs among groups, even if there are valid, profit-motivated reasons for their doing so. This definition corresponds to the legal definition of discrimination and is the one that is regularly used implicitly in policy debates. Of course, actual discrimination will often be difficult to detect, since I assume that outsiders (that is, regulators and econometricians) cannot observe the second signal. As a consequence, lenders will typically be accused of discrimination if the "apparent" cutoff using the initial signal, s_1^*, differs across groups.

I start by focusing on the difference in the denial rate—the proportion of the applicant pool that is rejected—across the two groups. In this simple framework, the denial rate is easy to calculate as the probability that a randomly selected applicant's inferred quality is less than q^*. The denial rate of group M applicants is then simply

$$(6) \qquad D_M = \Phi\left(\frac{q^* - \bar{\theta}}{\sigma_\theta^2/\sqrt{\sigma_\theta^2 + \sigma_S^2}}\right),$$

where Φ is the cumulative standard normal distribution function. Similarly, the group W denial rate is

17 As before, s_2 could be modeled as a sample mean of draws from the density p. By varying the number of observations used to make up s_1 and s_2, these two signals could be given different relative weights. All of my basic results would continue to hold in this more general model.

$$(7) \qquad D_W = \Phi\left(\frac{q^* - \bar{\theta}}{\sigma_\theta^2 / \sqrt{\sigma_\theta^2 + \sigma_S^2 / 2}}\right).$$

Expressions (6) and (7) make it clear that which group's denial rate is higher will depend on whether the minimum acceptable creditworthiness, q^*, is above or below $\bar{\theta}$, the average creditworthiness of the population.

If $q^* > \bar{\theta}$, it is straightforward to see that increases in the dispersion of creditworthiness in the applicant pool (increases in σ_θ^2) have the effect of decreasing the denial rate for both groups. In contrast, when applicant signals are imprecise (a high σ_S^2), lenders get relatively little useful information about an applicant's true quality, and therefore rely more heavily on the characteristics of the applicant pool. Thus, if the "average" applicant is uncreditworthy, a high σ_S^2 will tend to raise the denial rate for both groups.

Under this assumption about the creditworthiness of the "average" applicant, we get

RESULT 1: If lenders have an affinity with members of group W and if $q^* > \bar{\theta}$, the group M denial rate is higher than that for group W applicants.

This result mirrors that obtained by CKL (theorem 4). Essentially, the added information provided by s_2 helps to make the inferred quality of a group W applicant more accurate, making "good" signals from group W applicants more credible than "good" signals from group M applicants. Put another way, because their signals are not very precise, group M applicants tend to look more like the "average" applicant, and hence less creditworthy, than do group W applicants.

It is worth questioning, however, the reasonableness of the assumption that q^* is greater than $\bar{\theta}$. If this were true, the properties of the normal cumulative distribution function would imply that a majority of all mortgage applicants are rejected. Yet, raw denial rates for conventional home mortgage loans in the United States have ranged from 15 to 20 percent for whites and from 30 to 40 percent for blacks and Hispanics in the last several years, a fact seemingly at odds with this prediction.

One way to rationalize this assumption is to argue that the relevant population includes both mortgage applicants and those who do not apply because they believe they are not creditworthy. This is CKL's implicit assumption. In this case, it may be reasonable to assume that only a minority of all (potential) applicants obtain loans; some are denied loans, while the rest never even bother to apply, believing that they will be rejected. An alternate assumption would be that lenders use some initial screen to distinguish clearly qualified from unqualified applicants, obtaining the signals s_1 and s_2 only for "marginal" applicants.

The problem with both of these stories is that they make it impossible to test the model's empirical predictions. Home Mortgage Disclosure Act (HMDA) data are inadequate, since they include only households that actually apply for loans (not potential applicants) and do not identify marginal applicants. Similarly, general home-ownership rates cannot be used to test the model. Although it is likely that many households do not apply for a mortgage because they fear being rejected, there are other reasons for not becoming a homeowner. Indeed, many creditworthy households prefer the flexibility and smaller capital commitment of renting.

Because of these limitations, I extend CKL's analysis to consider what happens when the cutoff for creditworthiness is below the average in the applicant pool. I maintain this assumption throughout the rest of this article. Yet, readers should remember that it is also possible to assume the opposite, and some of the results that follow are reversed when this is done.

On the surface, it might seem strange that lenders accept below-average applications. Indeed, if evaluating applications is costly, the implication is that lenders would be better off accepting all applicants without screening. This, of course, ignores the fact that the makeup of the applicant pool depends on the use of a screen; lenders that accept all comers will soon find that everyone will apply, and the distribution of their applicant pool will be much worse than average. Although I do not model or discuss this feature of the mortgage market further, it is an important one to keep in mind.

Once we assume that q^* is less than $\bar{\theta}$, changes in the model's parameters have the opposite effects from those they had before: The higher σ_θ^2 and the lower σ_S^2 are, the *higher* the denial rate will be for both groups. In addition, the presence of the signal s_2 now has the opposite effect on relative denial rates.

RESULT 2: If lenders have an affinity with members of group W and if $q^* < \bar{\theta}$, the denial rate is higher for group W than it is for group M.

Note that, although lenders in this world require members of both groups to be equally creditworthy to be approved for a loan, they will require that group W applicants meet a higher cutoff signal. That is, the cutoff value of the final signal $\bar{s}^*$ is, in fact, higher than s_1^*. Why? Because their signal does less to distinguish creditworthiness, members of group M tend to look like the average applicant regardless of their signal; in contrast, the inferred quality of group W applicants is more disperse. Since by assumption the cutoff creditworthiness level is below the average in the applicant pool (that is, $q^* < \bar{\theta}$), lenders infer that "marginal" group W applicants (that is, those with below-average signals) are *worse* credit risks than essentially similar group M applicants: $q_W(s) < q_M(s)$ for all $s < \bar{\theta}$. Thus, rational lenders would require group W applicants to have a better signal than group M applicants.

Of course, the "better" signal required of group W applicants does not necessarily mean that this discrimination will be detected, since outsiders cannot observe its value. Nonetheless, lenders' use of it increases the average s_1 of all group W applicants, both approved and denied, giving the appearance that lenders require members of this group to clear a higher hurdle.

RESULT 3: If lenders have an affinity with group W applicants, they will want to discriminate against members of group W. This discrimination will be apparent to outsiders.

Although lenders will seem to require members of group W to meet a higher cutoff signal, this standard will appear more flexible than it is for members of group M.

RESULT 4: If lenders have an affinity with members of group W, it will appear that members of group M are held to a rigid cutoff standard, while regular exceptions are made for members of group W.

Thus, the initial underwriting decision for group W applicants will often be overridden. As it turns out, positive overrides will be outnumbered by negative ones.

Why do negative overrides dominate? With any signal, lenders make mistakes—they approve some loans that should be rejected and reject some that should be approved. Since the cutoff quality q^* is below the average for the applicant pool, lenders know that truly "bad" applicants are uncommon. Hence, they are willing to accept applicants with relatively low signals, knowing that these signals are quite noisy. As they gather more information about applicants, however, they place relatively less weight on the population's average creditworthiness and more on the individual applicant's signal. Consequently, many group W applicants who really are uncreditworthy will send a second, marginal signal, and lenders will correctly ascertain that this information is more reliable and deny the loan. Of course, there will be positive overrides as well, but these will be relatively less numerous precisely because lenders were giving applicants "the benefit of the doubt."

Hence, cultural affinity, combined with the assumption that $q^* < \bar{\theta}$, leads to three strong empirical predictions: Groups that share an affinity with lenders will 1) be denied loans more frequently, 2) appear to be required to meet a higher cutoff signal, and 3) be held less rigidly to that cutoff signal.

One might argue that, on the surface, these results suggest that the cultural affinity hypothesis is incorrect, since the prediction that group W applicants are denied loans more frequently is counterfactual. After all, we are inclined to think of group W as representing white applicants, who are rejected much less frequently than blacks and Hispanics. Likewise, it seems odd to conclude that lenders want to discriminate against white applicants. It becomes clear, then, that when $q^* < \bar{\theta}$, the cultural affinity hypothesis in isolation cannot fully explain any apparent discrimination that may exist in the residential mortgage market.

There are, however, at least two important characteristics of the mortgage market that I have thus far left out of the model. First, it is well known that black and Hispanic mortgage applicants are less creditworthy on average than their white counterparts, which can affect lenders' inferences about an individual's likelihood of repaying a loan.[18] Second, the presence of a large, active secondary market for mortgage loans gives lenders an incentive to "game" any subjective information they obtain about an applicant. Once these facets of the market are incorporated into the model, we see that the cultural affinity hypothesis can explain many features of the residential mortgage market.

18 For evidence of this fact, see Munnell et al. (1996), table 1.

Differential Group Creditworthiness

The analysis above assumes that the two applicant groups are identical; lenders treat them differently only because they receive more precise signals from members of group *W*. Yet, black and Hispanic mortgage applicants are less creditworthy on average than their white counterparts. If the factors that lenders consider in deciding whether to accept an application cannot fully account for this, lenders will have a profit-based incentive to discriminate against such applicants.[19] As a result, the effects of cultural affinity may simply be partially offsetting the stronger effects of average group creditworthiness, whether or not any actual discrimination occurs.

To understand this idea more fully, assume that lenders have no affinity with either group (they receive only signal s_1 from members of both groups). Suppose that group *W* applicants are, on average, better credit risks than group *M* applicants (that is, $\bar{\theta}_W > \bar{\theta}_M$) and that lenders know this. Now recall that an applicant's inferred quality, $q(s_1)$, is the weighted average of his signal and the average creditworthiness of his group. It follows that an applicant from group *W* is more creditworthy than a group *M* applicant who sends the same signal:

$$(8)\quad q_W(s_1) = \frac{s_1\sigma_\theta^2 + \bar{\theta}_W\sigma_s^2}{\sigma_\theta^2 + \sigma_s^2} > \frac{s_1\sigma_\theta^2 + \bar{\theta}_M\sigma_s^2}{\sigma_\theta^2 + \sigma_s^2} = q_M(s_1).$$

Since lenders are willing to make loans only to applicants whose inferred creditworthiness is at least q^*, they will have an incentive to use a different cutoff signal, s_1^*, for group *M* than they do for group *W*. This idea is summarized in

RESULT 5: If the average creditworthiness of group M is below that of group W, lenders will rationally want to discriminate against members of group M by requiring them to meet a higher cutoff signal than is required for members of group W.

Notice that lenders here do not ask members of group *M* to be *more* creditworthy; the minimum acceptable quality is the same for both groups. However, because lenders cannot perfectly observe an applicant's true creditworthiness, and since the applicant's group is correlated with creditworthiness, lenders will want to account for this when deciding whether to make a loan. Since q is increasing in s_1, lenders in this world will want to set a higher cutoff signal, s_1^*, for group *M* applicants than for group *W* applicants. This behavior would constitute classic statistical discrimination.

The denial rate of each group in such a world is easily calculated as

$$(9)\quad D_W = \Phi\left(\frac{s_W^* - \bar{\theta}_W}{\sqrt{\sigma_\theta^2 + \sigma_s^2}}\right), \text{ and}$$

$$D_M = \Phi\left(\frac{s_M^* - \bar{\theta}_M}{\sqrt{\sigma_\theta^2 + \sigma_s^2}}\right).$$

From this expression, it is clear that even if lenders did not discriminate and used the same cutoff signal for both groups, the denial rate would be higher for group *M*; the fact that lenders would like to use a more stringent cutoff signal for members of group *M* just exacerbates this effect, making the difference in denial rates between the two groups even more dramatic.[20]

RESULT 6: When members of group M are less creditworthy on average than members of group W, the denial rate for group M will be higher than that for group W, even in the absence of discrimination; when lenders discriminate, this disparity becomes still greater.

Cultural Affinity with Differential Group Creditworthiness

When we return to the assumption that lenders have an affinity with members of group *W*, we see two countervailing effects. On the one hand, the denial rate for group *M* is higher because of its members' lower average creditworthiness and any actual statistical discrimination that occurs. On the other hand, the added information provided by the second signal causes lenders to deny loans to group *W* applicants more frequently and to discriminate

■ **19** Calem and Stutzer (1995) develop an alternative model of statistical discrimination based on adverse selection in the mortgage market.

■ **20** It is worth noting that this result does not depend on the assumption that $q^* < \bar{\theta}$.

against them. The final outcome in the mortgage market depends on the balance between these two effects.[21]

By comparing the denial rates of the two groups, it is easy to see that group *M* applicants will be denied loans more frequently if their average creditworthiness is sufficiently low:

$$(10) \qquad \bar{\theta}_M < q^* + (\bar{\theta}_W - q^*) \frac{\sqrt{\sigma_\theta^2 + \sigma_s^2/2}}{\sqrt{\sigma_\theta^2 + \sigma_s^2}}.$$

RESULT 7: Assume lenders have an affinity with members of group W. If the average creditworthiness of group M is sufficiently low, the denial rate of group M will exceed that of group W. When lenders discriminate against members of group M, such discrimination may or may not be detected by outsiders.

This last conclusion follows because cultural affinity creates the appearance that lenders are discriminating against members of group *W*, offsetting the effects of their discrimination against members of group *M*.

Secondary Market Distortions

Another factor that can alter the effects of cultural affinity in the home mortgage market is the presence of a secondary market for mortgage loans. Recall that I originally assumed that lenders used the signal s_2 to update their evaluation of group *W* applicants' creditworthiness, approving some applications that were initially rejected and rejecting some that were initially approved. Ultimately, I argued, negative overrides outweighed positive ones, leading to higher denial rates for group *W* applicants.

Yet, anecdotal evidence suggests that many lenders seem only rarely to reject applicants who have passed the initial screen, raising the question of whether negative overrides really do outnumber positive ones. Once a group *W* applicant has been approved using the first (objective) signal, lenders may choose to ignore any additional "bad" information they receive about that applicant or, perhaps more likely, may never bother to observe the second signal at all. If we treat the first signal as a proxy for the information that secondary market institutions—Fannie Mae and Freddie Mac—require to purchase or guarantee a loan, lenders that sell their loans to the secondary market may have no incentive to consider negative information about applicants whom secondary market agencies are willing to approve.[22]

Positive information, however, will always be used by a lender. First of all, Freddie Mac and Fannie Mae guidelines allow originators to consider compensating factors when evaluating an application. If the originator can document an applicant's creditworthiness, his or her loan will still be salable on the secondary market if it fails to pass muster based on the initial signal. Furthermore, even if the lender cannot document a loan's quality to the secondary market's satisfaction, lenders can choose to hold obviously creditworthy applications in their own portfolios. Thus, there are at least two strong reasons why an originator might use the second signal to make positive overrides, while never rejecting applicants who pass the initial screen.

In such a world, we get

RESULT 8: Assume that lenders have an affinity with members of group W. When lenders ignore "bad" information contained in the signal s_2, group M applicants will be denied loans more frequently than group W applicants. Furthermore, even in the absence of discriminatory behavior, lenders will appear to discriminate against group M applicants (that is, appear to require them to pass a more stringent signal). Finally, it will be apparent to outsiders that group M applicants are held more rigidly to traditional underwriting standards than are group W applicants.

Effectively, lenders collect only the second signal from group *W* applicants with $s_1 < s_1^*$; the rest are immediately approved and sold to the secondary market. Now, it is group *W* applicants who are given the "benefit of the doubt," since they have the opportunity to overcome a poor initial signal with new information. In contrast to before, cultural affinity problems can exacerbate differences in denial rates that arise if group *M* applicants are less creditworthy on average.[23]

21 Note that if $q^* > \bar{\theta}$, the effects of cultural affinity and differential group creditworthiness compound one another to the detriment of group *M*.

22 Alternatively, one could imagine that an individual loan officer might have little incentive to relay negative information about an applicant to the loan committee, since doing so would reduce the chance of earning a commission.

23 If $q^* > \bar{\theta}$, both the disparity in denial rates and the apparent discrimination are even more pronounced.

Finally, note that I would reach the same conclusions even without a secondary market, as long as lenders found it more costly to verify "good" initial signals than "bad" ones. This might occur, for instance, if regulations make it difficult to justify denials based on the information contained in the second (subjective) signal. If this were the case, lenders would make only positive overrides, and result 8 would hold even in the absence of a secondary market.

A Note on Default Rates

In the introduction, I argued that Munnell et al.'s conclusion of taste-based discrimination was unsatisfying partly because of evidence that marginal black and Hispanic borrowers default more frequently than their white counterparts. Up to this point, we have not discussed the impact of cultural affinity on default rates.

In its purest form, discrimination that arises either from cultural affinity or from differential group creditworthiness will have the effect of *equalizing* the probability that a marginal applicant in each group defaults. Lenders choose their cutoff signals to ensure that the inferred quality of the last applicant approved will earn them a non-negative expected return. But this just means that their expected probability of default is the minimum acceptable to the firm:

$$q_M(s_1^*) = q_W(\bar{s}^*) = \theta^*.$$

As a result, discrimination of this sort will not lead to cross-racial differences in the default rates of marginal applicants. Indeed, it will tend to equalize these default rates, even if other factors tend to pull them apart. In the real world, however, setting different cutoff signals for different groups is patently illegal. As a result, fair lending laws can create disparities in the default rates of marginal applicants across groups, even where none would have existed otherwise.

To see this, consider the case of a portfolio lender—that is, a lender that does not sell its loans on the secondary market—with a cultural affinity for group W applicants in a market where group M applicants are less creditworthy on average. Suppose also that $\bar{\theta}_M$ is low enough to make lenders want to discriminate against group M applicants by using a higher cutoff signal. Finally, assume that the more a lender's cutoff signal varies across groups, the more likely regulators are to detect and punish that lender for discrimination. In such a world, lenders may well choose to discriminate against group M applicants. Nonetheless, the possibility of detection and punishment will cause them to discriminate less than they would in the absence of fair lending laws. As a result, they will select a lower cutoff signal for group M applicants and a higher one for group W applicants, making the implied creditworthiness of marginal group W applicants higher than that of group M applicants. In other words, enforcement of fair lending laws can lead to default rate disparities like those found by Berkovec et al. (1994).

II. Concluding Thoughts

Of course, in the real world there are many different lenders. Most high-volume lenders sell a large proportion of their loans to the secondary market, while other, typically smaller, institutions hold most of their loans in their own portfolios. As a consequence, we would expect to see different market behaviors by different types of lenders. Furthermore, there is strong evidence to suggest that minority applicants are, in fact, less creditworthy on average than their white counterparts. Putting all of these ideas together allows us to make some strong empirical predictions that seem consistent with what we see in the mortgage market.

Assuming that most lenders have an affinity with white applicants and that the cutoff "quality" is below the average creditworthiness of both groups, we would expect to see large lenders that sell most of their loans to the secondary market rejecting minority applicants more frequently than white ones, whether or not they discriminate. Furthermore, because they regularly make exceptions for marginally qualified white applicants (positive overrides), these lenders will appear to discriminate against minority applicants by holding them to a more stringent credit standard. If such lenders actually do discriminate against minority applicants, these effects will merely be exacerbated. Finally, the few loans that such lenders hold in portfolio will tend to be from white applicants and will perform better than ostensibly similar loans sold to the secondary market.

In contrast, small portfolio lenders will show fewer outward signs of discrimination. Although they too will likely reject relatively more minority applications, the difference in denial rates will be less stark than at larger institutions. Indeed, if the distribution of minority creditworthiness in the applicant pool is sufficiently good,

there may be no difference in denial rates at these institutions. This moderated denial rate differential, however, will not arise because of more frequent lending to minorities. Rather, it will result from the fact that such lenders do not pass off their "bad" white loans to the secondary market, but reject them instead. Finally, if these small lenders do discriminate against minority applicants, it may be difficult to detect this behavior. Because they refuse more white applicants than larger lenders do, the disparity in their denial rates may still appear less severe, even when they hold minority applicants to a more stringent signal.

Existing empirical evidence seems consistent with these implications of the cultural affinity hypothesis. Black and Hispanic mortgage applicants are rejected more frequently than whites, as evidenced by HMDA data over the last six years. Furthermore, Hunter and Walker (1996) show that higher denial rates for blacks and Hispanics seem to result from lenders giving more "breaks" to white applicants, meaning that blacks and Hispanics seem to be held much more rigidly to standard underwriting criteria.[24]

Finally, it is worth asking what happens when lenders have a cultural affinity with black and Hispanic applicants rather than with whites. This might be true, for instance, for minority-owned institutions and other lenders that make a particular point of marketing to minority communities. Typically, such institutions are smaller portfolio lenders. In contrast to white-owned portfolio lenders, however, these institutions tend to deny minority loans at an even higher rate than other lenders. This happens because their affinity makes them better at weeding out unprofitable minority applications and keeping only those that are sufficiently creditworthy. Just as white-owned institutions appear to discriminate against white applicants, black-owned institutions seem to hold minority applicants to harsher, if more flexible, underwriting criteria. However, in contrast to its effect on the white applicants discussed earlier, this discrimination exacerbates the higher denial rate that would arise from the deficiencies of minority applicants' average credit prospects.

It may seem strange to suggest that black-owned banks are more likely to discriminate against black applicants, but recent evidence suggests that this may well be the case. Black, Collins, and Cyree (forthcoming) use logistic regression techniques (similar to those of Munnell et al.) to show that black-owned banks also appear to require that black applicants meet a higher cutoff credit standard. The theory of cultural affinity can, therefore, explain this seemingly counterintuitive finding.

The above empirical predictions make it clear that detecting and eradicating discrimination in a market with cultural affinities can be quite difficult. In particular, the observable behavior of a discriminating lender will vary according to its position in the market (that is, whether it is a portfolio lender or one that sells its loans to the secondary market, and whether it specializes in making minority loans). Focusing solely on a lender's denial rates provides little, if any, information about its true actions. Instead, examiners must concentrate their efforts on understanding the makeup of the lender's applicant pool and the screen it is using to evaluate those applicants. By comparing a lender's denial rate with that implied by its credit standards and applicant population, regulators can more accurately determine what is causing the denial rate disparities: the effects of cultural affinity, differential group creditworthiness, or illegal discrimination.

24 See also Bostic (1995). Interestingly, he finds that minorities receive favorable treatment regarding loan-to-value ratios, but face negative biases with respect to obligation ratios.

References

Avery, R.B., P.E. Beeson, and M.S. Sniderman. "Posted Rates and Mortgage Lending Activity," *Journal of Real Estate Finance and Economics*, vol. 13, no. 1 (July 1996), pp. 11–26.

Becker, G.S. *The Economics of Discrimination*, 2d ed. Chicago: University of Chicago Press, 1971.

_____. "Nobel Lecture: The Economic Way of Looking at Behavior," *Journal of Political Economy*, vol. 101, no. 3 (June 1993), pp. 385–409.

Berkovec, J.A., G.B. Canner, S.A. Gabriel, and T.H. Hannan. "Race, Redlining, and Residential Mortgage Loan Performance," *Journal of Real Estate Finance and Economics*, vol. 9, no. 3 (November 1994), pp. 263–94.

Black, H.A., M.C. Collins, and K.B. Cyree. "Do Black-Owned Banks Discriminate against Black Borrowers?" *Journal of Financial Services Research*, forthcoming.

Bostic, R.W. "Patterns of Racial Bias in Mortgage Lending," Board of Governors of the Federal Reserve System, unpublished manuscript, August 1995.

Browne, L.E., and G.M.B. Tootell. "Mortgage Lending in Boston—A Response to the Critics," *New England Economic Review*, September/October 1995, pp. 53–78.

Calem, P., and M. Stutzer. "The Simple Analytics of Observed Discrimination in Credit Markets," *Journal of Financial Intermediation*, vol. 4, no. 3 (July 1995), pp. 189–212.

Calomiris, C.W., C.M. Kahn, and S.D. Longhofer. "Housing-Finance Intervention and Private Incentives: Helping Minorities and the Poor," *Journal of Money, Credit, and Banking*, vol. 26, no. 3, part 2 (August 1994), pp. 634–74.

Campbell, T.S., and J.K. Dietrich. "The Determinants of Default on Insured Conventional Residential Mortgage Loans," *Journal of Finance*, vol. 38, no. 5 (December 1983), pp. 1569–81.

Cornell, B., and I. Welch. "Culture, Information, and Screening Discrimination," *Journal of Political Economy*, vol. 104, no. 3 (June 1996), pp. 542–71.

Day, T., and S.J. Liebowitz. "Mortgages, Minorities, and Discrimination," University of Texas at Dallas, unpublished manuscript, 1993.

DeGroot, M.H. *Probability and Statistics*, 2d ed. Reading, Mass.: Addison-Wesley, 1989.

Ferguson, M.F., and S.R. Peters. "A Symmetric-Information Model of Credit Rationing," University of Connecticut, unpublished manuscript, August 1996.

_____, and _____. "Cultural Affinity and Lending Discrimination: The Impact of Underwriting Errors and Credit Risk Distribution on Applicant Denial Rates," *Journal of Financial Services Research*, forthcoming.

Hogg, R.V., and A.T. Craig. *Introduction to Mathematical Statistics*, 4th ed. New York: Macmillan, 1978.

Horne, D.K. "Evaluating the Role of Race in Mortgage Lending," *FDIC Banking Review*, vol. 7, no. 1 (Spring/Summer 1994), pp. 1–15.

Hunter, W.C., and M.B. 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, S.D. "Rooting Out Discrimination in Home Mortgage Lending," Federal Reserve Bank of Cleveland, *Economic Commentary*, November 1995.

Munnell, A.H., L.E. Browne, J. McEneaney, and G.M.B. Tootell. "Mortgage Lending in Boston: Interpreting HMDA Data," Federal Reserve Bank of Boston, Working Paper No. 92–7, October 1992.

Munnell, A.H., G.M.B. Tootell, L.E. Browne, and J. McEneaney. "Mortgage Lending in Boston: Interpreting HMDA Data," *American Economic Review*, vol. 86, no. 1 (March 1996), pp. 25–53.

Yezer, A.M.J., R.F. Phillips, and R.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 (November 1994), pp. 196–215.