Working Paper
Center for Policy Research
Syracuse University

Pricing, Performance, and Discrimination in Home Insurance

John Yinger
November 2004

Introduction[1]

Pricing decisions by home insurance firms, like many other types of economic decisions, may involve discrimination based on race, ethnicity, sex or some other protected classification. This paper explores methods for uncovering this type of discrimination. We review the literature on pricing discrimination in general and examine the strengths and weaknesses of discrimination tests that appear in the literature. Several recent studies emphasize the advantage of including performance information in an analysis of discrimination. In the case of home insurance, performance is measured by a policyholder's claims. We present in detail a performance-based test for pricing discrimination devised by Ross and Yinger (2002) and show how it can be applied to the case of home insurance.

The Firm's Problem: Designing a Pricing Scheme

In home insurance, a firm's profit on a policy equals the price the customer pays, P, minus the loss, L, that is, the insurance payments made by the firm. Since the loss is a measure of the firm's investment, the profit rate is $(P-L)/L = (P/L) - 1$.

Now suppose a firm observes losses for a large sample of policies as a function of the characteristics of the customer and property observed at the time of application, say, X. Let M stand for a customer's minority status or the minority status of a neighborhood in which the insured house is located. Then the firm can build a pricing system by estimating the following regression, where α and β are parameters to be estimated and ε is a random error term:[2]

$$L = \alpha + \sum_i \beta_i X_i + \gamma M + \varepsilon. \tag{1}$$

The pricing problem applies to applications for which X and M are observed but L is not. Using the estimated values of the coefficients and the actual values of X and M, the form can obtain a predicted loss, $\hat{L}$, for any given application that is drawn from the same pool as the applications used to estimate equation (1). To make sure the profit rate is the same on every policy, the firm sets the price equal to a given proportion of the loss so that $P/\hat{L}$ is constant. In other words, it sets $P = \theta \hat{L}$, where θ is the same for all policies. The value of θ is determined by competition.

The most direct pricing policy is based on the predicted value from equation (1):

$$\hat{L} = \hat{L}_1 = \hat{\alpha} + \sum_i \hat{\beta}_i X_i + \hat{\gamma} M \tag{2}$$

and

$$P_1 = \theta \hat{L}_1 = \theta \left(\hat{\alpha} + \sum_i \hat{\beta}_i X_i + \hat{\gamma} M \right). \qquad (3)$$

Our civil rights laws do not allow a firm to use *M* in setting prices, however, even if it is statistically significant in equation (1). This type of practice is known as statistical discrimination. Statistical discrimination arises whenever a customer's membership in a protected class provides a firm with information about likely losses, and the firm uses that information to increase its profits. A statistically significant estimate for γ implies that membership in a protected class helps to predict the loss associated with a particular policy, because the membership is correlated with variables that the firm cannot observe and that lead to higher losses. A firm might not be aware, for example, of a customer's experience with household maintenance. In this case, accounting for a customer's minority status in the pricing scheme is profitable—but illegal. A firm is allowed to collect information on experience with maintenance or any other customer characteristics and use them in its price-setting procedure, but it is not allowed to use vary prices based solely on a customer's membership in a protected class.

In an attempt to avoid practicing discrimination, a firm might proceed in one of two ways. The first way is to set *M* at its average value in the sample of applications used to estimate equation (1), $\bar{M}$. In this case,

$$\hat{L} = \hat{L}_2 = \hat{\alpha} + \sum_i \hat{\beta}_i X_i + \hat{\gamma} \bar{M} \qquad (4)$$

and

$$P_2 = \theta \hat{L}_2 = \theta \left(\hat{\alpha} + \sum_i \hat{\beta}_i X_i + \hat{\gamma} \bar{M} \right). \qquad (5)$$

This approach provides the best possible pricing scheme that is consistent with existing civil rights laws. The weight placed on X_i, namely, $\hat{\beta}_i$, is based solely on the direct impact of that variable on loss, controlling for minority status. Another way to put it is that this weight considers only the within-group impact of this variable on loss, which is what the law requires.

Alternatively, the firm might estimate equation (1) without including *M*. This procedure appears to be group neutral in the sense that it does not account for minority status. As shown by Ross and Yinger (2000), however, this procedure actually builds disparate-impact discrimination into the pricing system.

To see why this is true, consider the following regression:

$$L = \alpha' + \sum_i \beta_i' X_i + \varepsilon'. \tag{6}$$

The predicted loss from this regression, and the associated pricing scheme, are

$$\hat{L} = \hat{L}_3 = \hat{\alpha}' + \sum_i \hat{\beta}_i' X_i \tag{7}$$

and

$$P_3 = \theta' \hat{L}_3 = \theta' \left(\hat{\alpha}' + \sum_i \hat{\beta}_i' X_i \right). \tag{8}$$

Because *M* is omitted from equation (6), the estimated coefficients are subject to omitted variable bias. Using the standard formula for this bias, we find that the expected value of $\hat{\beta}_i'$ is

$$E(\hat{\beta}_i') = \beta_i + \gamma r_i, \tag{9}$$

where β_i and γ are the parameters in equation (1) and r_i is the correlation between X_i and M. Under standard assumptions, the expected values of the estimated coefficients in equation (3) equal the underlying parameter values in equation (1), so we can compare the expected values of the pricing schemes defined by equations (3) and (8). Specifically, Ross and Yinger show that:

$$E(P_3) = E(P_2) + \gamma \sum_i r_i (X_i - \bar{X}_i) \ . \qquad (10)$$

Equation (10) shows that the weight on each variable in pricing scheme P_3 is adjusted by a factor that depends on the product of γ and r_i. Whenever the coefficient of the minority status variable is positive and significant (γ>0 and significant), this scheme increases the weight on any variable that is positively correlated with minority status ($r_i > 0$). In other words, it uses any such variable to help predict which applications come from a minority applicant. This is, by definition, disparate-impact discrimination. This scheme raises the weights on these variables because they help identify minority applicants, not because they help predict losses within a group. Any practice, including setting one of these weights, must be based on legitimate business considerations, or, to use the legal term, must be based on business necessity. Our civil rights laws make it clear that these considerations cannot include an individual's membership in a protected class.

Of course, a firm may decide not to use statistics to obtain a pricing scheme and it may not bother to avoid even disparate-treatment discrimination. A general way to write a pricing scheme that recognizes these possibilities is as follows:

$$P_3 = \theta\left(\tilde{\alpha} + \sum_i \tilde{\beta}_i X_i + \tilde{\gamma} M \right), \qquad (11)$$

where the tilde (~) over a parameter indicates that it is based on a rule of thumb, a tradition, a guess, or a desire to discriminate—not on a statistical procedure.

The Researcher's or Enforcement Official's Problem: Detecting Discrimination

This analysis leads to the key issue of this paper, namely, how to detect discrimination. Three types of tests are possible: without performance information, with a performance variable, and with predicted performance.

Regressions without Performance Information

The first type is to run a regression of a firm's prices on all the variables that are thought to influence loss.[3] This regression is

$$P = a + \sum_i b_i X_i + cM + \mu, \tag{12}$$

where *a*, *b*, and *c* are coefficients to be estimated and μ is a random error term. If all relevant *X* variables are included, the argument goes, then the coefficient of *M*, *c*, provides a test of discrimination. This type of approach has been used, for example, to study wage discrimination (studied reviewed in Altonji and Blank, 2002), and discrimination in overages charged by lenders (Courchane and Nickerson, 1997; Crawford and Rosenblatt, 1999).

As explained by Ross and Yinger (2002), however, this approach has two serious flaws. First, it can only capture disparate-treatment discrimination. Disparate-impact discrimination is hidden in the weights placed on the various explanatory variables, and this equation does not reveal whether an estimated weight is based on legitimate business considerations or not. This is equivalent to setting up a system that allows researchers or civil rights enforcement officials to observe bank robberies, but only if the robbers leave by the front door instead of the back door.

Second, this approach provides a conservative estimate even of disparate-treatment discrimination, because it cannot determine whether the factors included in the regression are the factors actually used by the firm. If a regression includes a particular *X* that is not actually used by the firm but that is correlated with minority status, then disparate-treatment discrimination could, to some degree, be hidden in the coefficient of this *X* variable. The best way for a researcher to minimize this problem is to interview firms to determine which variables they actually use.

Regressions with a Performance Variable

A second type of test is possible when a scholar has data not only on prices and time-of-application characteristics, but also on performance. This approach has been extensively examined in the case of loan approval decisions. The basic point, made by Becker (1993), is that discrimination in loan approval constitutes a higher hurdle for minority applicants and therefore should result in higher performance for approved minority mortgages than for approved white mortgages. Berkovec at al. (1994, 1998) extend this approach by regressing loan performance on time-of-application characteristics and minority status. They argue that a positive coefficient for the minority variable is a sign of discrimination. This approach has been criticized by Galster (1996), Ross (1996a), and Yinger (1996).[4] These critiques are summarized in Ross and Yinger (2002), who also show how problems with this so-called default approach can be avoided by combining loan-approval and loan-performance information.

As we will see, a few studies have also introduced performance information into a study of pricing discrimination. In the case of home insurance, this approach is possible with data on the loss associated with each policy in the sample. In this case, the appropriate regression involves only loss and minority status as explanatory variables, or

$$P = a' + b'L + c'M + \mu'. \tag{13}$$

In this regression, the coefficient of L can be thought of as an estimate of θ in equation (3). Moreover, the coefficient of M picks up the extent to which people in a protected class face prices that are above the level justified by the losses associated with their policies. This coefficient has also been interpreted as a measure of discrimination.

An example of this approach can be found in the wage discrimination study by Hellerstein, Newmark, and Troske (1999).[5] Using a data set that combines firm and individual information for manufacturing workers, Hellerstein et al. find that, in most cases, inter-group differences in wages are accompanied by equivalent inter-group differences in productivity.[6] One key exception is for female workers, whose relative wages are lower than expected based on their relative productivity, which is a sign of discrimination.

This approach faces a much greater challenge when applied to a market in which prices are set based on predicted performance, namely, that it cannot measure statistical discrimination.[7] This condition clearly applies to the insurance market, where losses are only realized after a policy is issued. Now recall from the discussion of equation (1) that losses might be higher for customers in a protected class because of factors that are not observed by the firm, such as maintenance experience. In other words, losses might be higher due to factors that the firm is not allowed to consider in setting its prices. In this case, controlling for *L* is equivalent to slipping these factors back in.

To make this point in a formal way, consider pricing policy P_1, which is defined by equation (3). Now assume that the policies used to estimate the pricing scheme and the policies observed by the researcher are drawn from the same pool. Because the pricing scheme and the equation determining performance have exactly the same form, and because the expected values of the coefficient in the pricing equation (3) equal the parameter values in the performance equation (1), a regression of P_1 on *L* fully explains the variation in P_1. Moreover, the expected value of $\hat{b}'$ equals θ. Thus, the constant term and the minority status variable in equation (13) have no role to play and the expected values of the relevant coefficients equal zero; that is, $E(\hat{a}') = E(\hat{c}') = 0$. In other words, this regression will find no impact of minority status on pricing even though firms practice statistical discrimination. Of course, the coefficient of *M* will still pick up discrimination that is not based on profit maximization.[8]

One might argue that it is possible to avoid this problem by including the *X*'s in the regression. In fact, however, this step does not help at all. The coefficients of the *X* variables simply measure the difference between the weights placed on these variables in the firm's actual pricing scheme and the weights that would arise if the firm used pricing scheme P_1. Estimating this difference does not shed any light on the existence of statistical discrimination.

Consider the general pricing scheme defined by equation (11). A comparison of this pricing scheme and equation (1), which explains *L*, reveals that

$$E(a') = \theta(\tilde{\alpha} - \alpha)$$
$$E(b_i') = \theta(\tilde{\beta}_i - \beta_i) \qquad (14)$$
$$E(c') = \theta(\tilde{\gamma} - \gamma)$$

These coefficients are related to discrimination, but do not accurately measure it. Deviations between the weights actually used, marked with a tilde, and the weights that actually influence loss raise the possibility of disparate-impact discrimination. These weights differ, for example, if a firm builds statistical discrimination into its weights by using pricing scheme P_2. However, this equation does not reveal whether existing deviations actually have a disparate impact. Moreover, the expected value of the minority status coefficient still equals zero if the actual pricing scheme is based on standard statistical discrimination $(\tilde{\gamma} = \gamma)$.

A version of this second approach that has been suggested for studying insurance (Harrington and Niehaus, 1998) is to use the loss ratio, L/P as the dependent variable. This is equivalent to estimating equation (13) with a couple of restrictions. To be specific, this approach assumes that[9]

$$\frac{L}{P} = a'' + b''M \qquad (15)$$

or

$$P = \left(\frac{1}{a'' + b''M}\right)L \qquad (16)$$

Equation (15) is similar to equation (13) except that it assumes the impact of M operates exclusively through the ratio L/P, which, in the terms defined earlier, is $1/\theta$. In this case, discrimination involves setting a higher value of θ, that is, a higher price for any given expected performance, for people in the protected class. The test for discrimination, therefore is whether b'' is negative, which implies that θ is higher when $M = 1$. This approach cannot detect a higher price for minority applicants that is not associated with θ. It also cannot detect statistical discrimination.

Recall that the value of θ is determined by competition. This approach is equivalent, therefore, to assuming that firms cannot discriminate unless they have market power.[10] There is no need to impose this assumption, however. The following regression makes it possible to determine whether or not discrimination operates through θ, that is, whether or not discrimination is associated with market power:

$$P = a''' + b'''L + c'''M + d'''LM + \mu'''. \tag{17}$$

In this equation, a positive, significant estimate for c''' indicates pricing discrimination not associated with θ and a positive, significant estimate of d''' indicates pricing discrimination that takes the form of a higher θ for customers in the protected class. This form still cannot detect statistical discrimination, however.

Harrington and Niehaus (1998) do not have data for individual insurance policies, but instead apply equation (15) to automobile insurance using zip-code level data with the percentage of the population that is black as the measure of minority status. They argue that in the case of automobile insurance, θ depends not only on competition but also on a “loading factor” for non-claim costs, such as claim-administration costs, and they include many zip-code demographic and coverage related factors to account for this possibility. In the notation of equation (15), they make a'' a function of these factors. This is a legitimate concern. Even in a competitive environment, the P/L might vary over time and might vary across locations. Including variables to account for these possibilities can minimize the possibility that the minority status coefficient is biased because minority customers buy insurance a different time or in a different place from other customers.[11]

Unfortunately, however, Harrington and Niehaus cannot test the assumption that the loading factor depends on their control variables; if it does not, the inclusion of these variables may, for the reasons given earlier, simply hide the role of minority composition. Moreover, even if this assumption is correct, their approach cannot uncover statistical discrimination.

Regressions with a Predicted Performance Variable

To uncover statistical discrimination when pricing depends on predicted performance, we must turn to the third approach. This approach, which was developed by Ross and Yinger (2002) for discrimination in mortgage pricing, is to regress *P* on $\hat{L}_2$, as defined by equation (4), which is a prediction of loss that does not involve statistical discrimination. In symbols, this regression is

$$P = a^* + b^* \hat{L}_2 + c^* M + \mu^* . \qquad (18)$$

Now if a firm practices statistical discrimination using pricing scheme P_1, as defined by equation (3) (or using any other scheme, for that matter), the estimate of c^* will pick it up because the impact of *M* on *L* due to unobservable factors is no longer included in the *L* variable.[12] This approach could easily be extended to consider discrimination in setting θ by including an interaction between $\hat{L}_2$ and *M*, as in equation (17).

This approach requires more data than the second approach as we have presented it (but it does not require more data than the second approach with the *X*'s included in the regression). The challenge for this approach is to have information on all the *X*'s used by the firm in setting its prices. This challenge is analogous to the one widely acknowledged in traditional studies of discrimination using equation (12): without complete data, any estimate of discrimination is likely to be subject to omitted variable bias. As before, however, including the *X*'s directly in the regression does not help. This step would make it possible to determine whether the weights used by firms differ from the weights in the researcher's equation to predict losses, but it would do this by hiding discrimination, either by introducing multicollinearity or by hiding discrimination in the coefficients of the *X*'s. In fact, the estimated value of c^* in equation (18) already picks up both disparate-impact and disparate-treatment discrimination, including statistical discrimination. No additional insight into discrimination can be gained (and some might be lost) by adding the *X* variables.

As indicated earlier, it is appropriate to include additional control variables if they shed light on non-discriminatory reasons for variation in the *P*/*L* ratio. Time dummies to account for variation in this ratio over time would be appropriate, for example. For the reasons just given, however, variables without a clear link to the *P*/*L* ratio should not be included.

This approach could be extended to consider a more complicated model of a firm's profits.[13] Suppose, for example, that an insurance firm's losses depend on the probability of a claim being filed, *F*, multiplied by the expected payment on a claim once it has been filed, *E*. With this specification, *F* might be a function of one set of variables, *Z*, and *E* might be a function of another, *X*. Using the estimated coefficients from a regression of *F* on *Z* and setting minority status at its average value yields a predicted value of *F* that is analogous to L_2 in equation (4), say F_2. A comparable procedure leads to a predicted value of E as a function of X, say E_2. In this model, the expected value for L_2 in equation (18) is the product of F_2 and E_2. One cannot account for this more general model simply including new explanatory variables in equation (18). Not only would this be a misspecification, but, as shown earlier, it might hide discrimination in the coefficients of the new variables.

References

Altonji, Joseph G., and Rebecca M. Blank. 1999. "Race and Gender in the Labor Market," in Orley Ashenfelter and David Card, eds., *Handbook of Labor Economics*, 3C. North-Holland, 3143-3259.

Altonji, Jospeph G., and Charles R. Pierret. 2001. "Employer Learning and Statistical Discrimination," CXVI (1) (February) *The Quarterly Journal of Economics* 313-350.

Ayres, Ian, and Joel Waldfogel. 2001. "A Market Test for Race Discrimination in Bail Setting," in Ian Ayres, ed., *Pervasive Prejudice? Unconventional Evidence of Race and Gender Discrimination.* Chicago: University of Chicago Press.

Becker, Gary S. 1993. "Nobel Lecture: The Economic Way of Looking at Behavior," 101(3) *Journal of Political Economy* 385-409.

Berkovec, James A., Glenn B. Canner, Stuart A. Gabriel, and Timothy H. Hannan. 1994. "Race, Redlining, and Residential Mortgage Loan Performance," 9(3) *Journal of Real Estate Finance and Economics* 263-294.

Berkovec, James A., Glenn B. Canner, Stuart A. Gabriel, and Timothy H. Hannan. 1996. "Response to Critiques of 'Mortgage Discrimination and FHA Loan Performance.'" *Cityscape: A Journal of Policy Development and Research* 2 (1): 48-54.

Berkovec, James A., Glenn B. Canner, Stuart A. Gabriel, and Timothy H. Hannan. 1998. "Discrimination, Competition, and Loan Performance in FHA Mortgage Lending," 80(2) *Review of Economics and Statistics* 241-250.

Courchane, Marsha, and David Nickerson. 1997. "Discrimination Resulting from Overage Practices." *Journal of Financial Services Research* 11 (1-2) (April): 133-151.

Crawford, Gordon W., and Eric Rosenblatt. 1999. "Differences in the Cost of Mortgage Credit: Implications for Discrimination." *Journal of Real Estate Finance and Economics* 19 (2) (September): 147-159.

Galster, George C. 1996. "Comparing Loan Performance between Races as a Test for Discrimination." *Cityscape: A Journal of Policy Development and Research* 2 (1): 33-39.

Galster, George, Stephen Ross, and John Yinger. 1996. "Rejoinder to Berkovec, Canner, Gabriel, and Hannan." *Cityscape: A Journal of Policy Development and Research* 2 (1): 55-58.

Harrington, Scott E., and Greg Niehaus. 1998. "Race, Redlining, and Automobile Insurance Prices." *The Journal of Business* 71 (3) (July): 439-469.

Hellerstein, Judith K., David Neumark, and Kenneth R. Troske. 1999. "Wages, Productivity and Worker Characteristics: Evidence from Plant-Level Production Functions and Wage Equations," 17(3) *Journal of Labor Economics* 409-446.

Munnell, Alicia H., Geoffrey M. B. Tootell, Lynn E. Browne, and James McEneaney. 1996. "Mortgage Lending in Boston: Interpreting HMDA Data," 86(1) *American Economic Review* 25-53.

Ross, Stephen L. 1996a. "Flaws in the Use of Loan Defaults to Test for Mortgage Lending Discrimination." *Cityscape: A Journal of Policy Development and Research* 2 (1): 41-48.

Ross, Stephen L. 1996b. "Mortgage Lending Discrimination and Racial Differences in Loan Default," 7(1) *Journal of Housing Research* 117-126.

Ross, Stephen L., and John Yinger. 2002. *The Color of Credit: Mortgage Discrimination, Research Methodology, and Fair-Lending Enforcement*. Cambridge, MA: MIT Press.

Szymanski, Stefan. 2000. "A Market Test for Discrimination in the English Professional Soccer Leagues," 108(3) *Journal of Political Economy* 590-603.

Yinger, John. 1996. "Why Default Rates Cannot Shed Light on Mortgage Discrimination." *Cityscape: A Journal of Policy Development and Research* 2 (1): 25-31.

Endnotes

---

[1] For lack of a data set and a co-author, this working paper has never been published. If you have a data set and would like to work on this paper with me, please let me know!

[2] The material in this section draws heavily on Ross and Yinger (2002).

[3] This approach is widely used for studies of discrimination in behavior other than pricing. See, for example, the study of discrimination in loan approval by Munnell, et al. (1996).

[4] Berkovec et al. (1996) provide a rejoinder to these critiques; Galster, Ross, and Yinger (1996) then provide a reply.

[5] Other examples include Szymanski (2000), who examines the performance of English soccer teams as a function of the teams' racial composition after controlling for the teams' wage bill, and Ayres and Waldfogle (2001), who study bail charges as a function of race after controlling for flight risk.

[6] If workers in protected classes are relatively recent hires, this result is also consistent with the hypothesis that firms practice statistical discrimination in setting initial wages.

[7] As recognized by Berkovec et al. (1996), this problem also arises in performance-based studies of discrimination in loan approval, and they only claim to look for the presence of prejudiced-based discrimination. Ross (1996b) shows, however, that in the presence of statistical discrimination, the performance-based approach provided by Berkovec et al. (1996) may not even be able to correctly measure prejudice-based discrimination.

[8] In some markets, including the labor market, he pricing decision is updated based on directly observed information on performance. If the researcher can observe performance, then so can the firm, and the firm is entitled, of course, to base pay on

performance. The shift to an on-going relationship makes it possible to differentiate between prejudice-based discrimination and statistical discrimination. Specifically, if a firm practices statistical discrimination in wages at the time of hiring, the coefficients of variables measuring predicted performance at the time of hiring (as in equation (12)), including membership in a protected class, should decline as firms gain information on an employee's performance and no longer have to predict it (Altonji and Pierret, 2001).

[9] These two approaches also make different assumptions about the specification of the random error term. Equation (15) assumes that there is random variation in θ and that there is not an additive error in the P equation (which is the assumption in equation (13)). This difference in assumptions affects the estimates obtained from the two approaches, but does not alter the substantive arguments made in the text.

[10] This assumption has also appeared in the mortgage discrimination literature. Berkovec et al. (1998) test this assumption by determining whether discrimination is higher in markets where lenders are more concentrated. A critique of their study is provided by Ross and Yinger (2002).

[11] This issue is discussed at length in Ross and Yinger (2002) for the case of discrimination in mortgage interest rates. As they put it, a test for discrimination based on predicted performance (the method discussed in the next section) should include "controls for market conditions and perhaps other factors, such as loan characteristics associated with higher origination costs" (p. 344).

[12] As pointed our earlier, Ross and Yinger (2002) develop a test a test for discrimination in loan approval that brings in performance information. That test also begins by

estimating a performance regression, but the second step is different because the loan-approval variable is discrete.

[13] Ross and Yinger (2002, p. 342) suggest that one could include performance and performance squared in a study of mortgage pricing discrimination. This suggestion reflects the fact that the link between the price/performance tradeoff and profits is not straightforward for mortgage interest rates. As discussed earlier, however, insurance profits are closely linked to the difference between prices and performance, so including a squared term does not make sense in an application of equation (18) to insurance.