SECTION 2

ESTIMATION OF NEIGHBORHOOD EFFECTS IN LOAN-LEVEL DATA

2.1. Data Sources and Key Definitions

The primary data used in this study are samples of FHA-insured loans[20] that were endorsed in 1992 or 1994, or for which applications were submitted in one of these years, and for which the subject property was contained within the 22 MSAs listed in Table 1 below. The choice of these particular MSAs was dictated by several factors. First, each of these MSAs contained sufficient numbers of defaults to support statistical analysis. Second, these MSAs contain a substantial portion of FHA business. Third, many of these MSAs contain substantial populations of blacks and Hispanics, but there is also significant variation in minority representation, thus supporting explorations along race/ethnic lines.

Table 1 provides some key characteristics of the 22 MSAs. The first four columns are based on data from all 1992 and 1994 FHA-insured purchase money loans. The first column shows the number of FHA defaults (defined more precisely below) in these data; there are at least a few hundred defaults (as of April 1996) in each of the 22 MSAs. The second third and fourth columns give the number of black, Hispanic, and total FHA-insured loans. In each case there are at least several hundred loans to at least one of the two minority groups. The next two columns, which are based on the 1990 Census of the Population, show the percentages of the MSA population that were black or Hispanic, respectively, in 1990. Minority representation varies substantially across these MSAs, but in some cases exceeds 40 percent. The final two columns are produced from 1993 and 1994 HMDA data. The second column from the right shows the precentage of national FHA-insured loans arising in each of the 22 MSAs. In the aggregate, these 22 MSAs contain about 44 percent of FHA business in 1993-94. Finally, the last column

[20] Streamline refinances are excluded because the applications for such loans lack critical information, such as loan-to-value ratios.

on the right shows FHA originations as a percentage of total originations within each MSA. These percentages vary widely, but in many cases the FHA share of originations exceeds 25 percent.

The FHA data files contain a variety of loan and borrower characteristics measured at loan origination, as well as information on the status of the loan at the time that the files were constructed (*e.g.*, whether the loan had prepaid or had resulted in a claim) and dates of critical events in the life of the loan (*e.g.*, the date of default, if any). After stripping these loan files of all identifying information that could be used to link a loan to a specific individual, portions of the 1992 and 1994 application data were sent to Trans Union and Equifax to obtain credit scores for borrowers and coborrowers.

For purposes of this analysis, defaults were defined to include only those loans that (a) defaulted on or before April 30, 1996, and (b) for which a claim[21] was recorded as of April 30, 1998.[22] In particular, loans that entered default status but subsequently cured are not included under this definition of default. To maintain consistency with the timing requirements for defaults, loans were recorded as prepayments only if they prepaid on or before April 30, 1996. Prepayments were further subdivided into two categories: (a) those occurring en route to refinancing with FHA, and (b) the remainder of prepayments, which occur for unspecified reasons that could include conventional refinancing, changing residence, or simply paying off the mortgage.[23]

The statistical estimation (described below) follows each of the loans on a monthly basis until it defaults or prepays or, if neither of the latter actions occurs, until April 30, 1996. Given

---

[21] Loans entering the assignment program are treated as defaults even though such loans technically remain active.

[22] We require that default occur at least two years prior to when our observation window closes (on April 30, 1998) so that sufficient time remains for all those who have defaulted to be observed in claim status by the close of the observation window. The concern with including defaults that occur later is that we may treat defaults inconsistently; that is, we may miss some defaults which, because of state foreclosure practices, are not recorded as claims by April 30, 1998.

[23] Inability to separate conventional refinances from other types of prepayment activity is a problem that weakens some of the statistical analyses in this paper.

that we are using 1992 and 1994 applications and originations, we follow loans for a maximum of 52 months, but most loans --- particularly the 1994 loans --- have a much smaller potential period of observability. Because defaults must also occur within at most 52 months, the defaults examined here would probably be considered "early" defaults. By way of comparison, Calem and Wachter (1999) use originations from 1998 through July 31, 1994, of which over 90 percent originated during or after 1990. Because they record delinquency as of November 15, 1994, potential duration in their study ranges from a few months up to a maximum of about 59 months for the 1990-1994 originations. Berkovec, et al (1994) examine loans that have had potential exposure of about 39 to 75 months.

Sampling from the universe of endorsed loans helped reduce the estimation burden to more manageable proportions.[24] Because defaults are relatively rare, they were oversampled. Because prepayment was relatively common, however, prepayment was not used explicitly to structure the sample. Among the 231,583 (non-streamline refinance) loans with 1992 application or endorsement dates, 9801 were defaults by our definition (and 221,782 were nondefaults). Of these, we selected a sample[25] consisting of 3,057 defaulted loans and 6,534 nondefaulted loans for inclusion in the statistical analysis. The set of nondefaulted loans contains 535 loans that refinanced through FHA and 836 that prepaid for other reasons. Among the 235,214 (non-streamline refinance) endorsements with 1994 application or endorsement dates, 5320 were defaults by our definition (and 229,894 were nondefaults). From these, we selected a sample consisting of 2,184 defaulted loans and 6,609 nondefaulted loans for inclusion in the statistical analysis. Of the nondefaulted loans, 402 refinanced with FHA and 264 prepaid for other reasons. Given the stratified nature of the sample, the statistical procedures employed weighting according to sample stratum (application year and default status) from which the loan was

[24] The hazard estimation procedure used below treats each month of loan activity as an observation, thus greatly expanding the effective sample size used in estimation and providing the impetus to subsample from the population of available loans.

[25] A moderate amount of data "cleaning" was performed prior to drawing the samples, and additional cleaning was performed afterwards, mainly to remove cases with doubtful values of relevant variables.

drawn.[26]

A problem that must be faced at the outset is how to define a neighborhood for the purposes of this analysis. One pragmatic approach is to follow area delimitations established by others, while admitting that these definitions may not be most appropriate for the task at hand. Here we follow this procedure, identifying the census tract as the neighborhood. Census tracts, which generally contain between 2,500 and 8,000 people and (when formed) are intended to be relatively homogeneous with regard to economic status, living conditions, and population characteristics, are small enough that they might plausibly be considered neighborhoods. Of course, there is no guarantee that they do indeed represent a correct geographical division from the housing market perspective, and it would be surprising if in fact tract boundaries perfectly coincided with those of local housing markets.

The FHA data were supplemented with monthly BLS data on unemployment rates at the MSA level, quarterly Freddie Mac data on house price indices at the MSA level, monthly data on Treasury rates, monthly data on conventional mortgage rates and points from the Freddie Mac Primary Mortgage Market Survey, and a variety of tract-level and MSA-level measures from Census files. Among the latter are counts of the populace by race or ethnicity, as well as area (MSA and tract) median incomes.

2.2. A Preliminary Look at the Data

Before undertaking a more detailed statistical analysis of the loan-level data from the 22 MSAs, it is of interest to give a brief overview of some relevant and interesting features of the data.

---

[26] The sample selection procedure was complicated by the fact that credit scores had been obtained only for loans for which applications were submitted in 1992 or 1994, not for cases that were endorsed in these years but for which applications were submitted in other years. We structured the sample for this analysis by selecting two different stratified samples, one composed of loans with credit scores and one composed of loans that had not been submitted for scoring. After preliminary statistical analysis that indicated that results for the full sample were very similar to those for the sample with scores, we proceeded to use the combined sample with appropriate weights. An additional complicating factor in calculating weights is that stratification was performed prior to the final determination of default status and was based on claim status as of an earlier date.

These will serve to illustrate some important empirical regularities that will be investigated more systematically in the statistical work to follow. This overview utilizes Tables 2 through 4, each of which is based on a large sample of approximately 600,000 purchase money loans drawn from the two application or origination years, 1992 and 1994. For these tables, default and prepayment rates are calculated over the pooled sample from both years.

Table 2 shows default rates by race[27] within census tracts classified by 1989 racial composition; in addition, the table gives the percentage of FHA loans by race falling in each tract classification. Note in this regard that within the universe of loans for these 22 MSAs, overall black and Hispanic default rates are virtually equal at 5.0 percent,[28] while that of others (the predominately white group composed of those who are neither black nor Hispanic) is much lower at about 2.3 percent. The table is split into three sections. The top section classifies tracts by the percentage of the population that is black; the middle section classifies tracts by the percentage of the population that is Hispanic; the bottom section classifies tracts by the percentage of the population that is of races other than black or Hispanic. Each section of the table gives default rates by race within each tract classification. In the first section of the table, we see, for example, that within tracts in which blacks make up no more than 10 percent of the population, the overall default rate is 2.51 percent, the black default rate is 3.79 percent, the Hispanic default rate is 4.31 percent, and the "other" (predominately white) default rate is 2.08 percent. Such tracts contain 69.03 percent of the FHA loans in this sample, 19.36 percent of the black loans, 71.05 percent of the Hispanic loans, and 79.81 percent of the "other" loans.

The first section shows that default rates, both overall and within race, tend generally to rise as the black percentage rises, though the pattern for default rates of Hispanics seems more irregular than the others. Note also that the vast majority of the loans (over 80 percent) fall into tracts in which no more than 20 percent of the population is black. The second section of the table shows a more confusing pattern. For most race/ethnic groups, default rates tend to rise with

---

[27] Here and in what follows we shall often use "race" as a shorthand for "race and ethnicity."

[28] Hispanic default rates over this interval have presumably been affected by the California recession and the concomitant decline in California house prices.

the percentage Hispanic but then to peak out and decline as Hispanic representation continues to rise. Finally, the bottom section of the table generally shows declines in default rates as population composition shifts to heavier representation of "other" (predominately white) persons.

Another interesting feature of these figures is revealed by making cross-race comparisons within each tract classification, *i.e.,* by looking across each row of the table. We see that while the default rate for "others" is always lower than that of blacks and Hispanics within each tract classification, the relationship between black and Hispanic default rates is irregular. The rate for Hispanics never exceeds that of blacks for any of the groups classified by tract percentage Hispanic; at the same time, the rate for Hispanics exceeds that of blacks for all but one of the groups classified by tract percentage black.

Table 3 is intended to show how default rates vary jointly by tract income and by borrower income. Each column of the table corresponds to an approximate decile of the difference between tract median income and MSA median income, as measured in the 1990 Census. Each row of the table corresponds to an approximate decile of the difference between borrower income and MSA median income for the corresponding application year. Each figure within the table gives the default rate for all loans within the relevant tract and borrower income decile. Thus, for example, the upper left-hand cell indicates that the default rate is 5.81 percent for loans in which tract income (less MSA income) is in the bottom decile and individual income (less MSA income) is in the bottom decile. Looking down each column of the table, we see patterns that are generally unclear; that is, there does not appear to be any strong systematic change in the default rate as one moves across borrower income categories within tract income categories. Looking across a row, however, seems to show that default rates decline as one moves towards higher tract income categories within a borrower income category.

Table 4 shows how default rates by race vary with the prepayment rate of the tract. Prepayments here include refinances with FHA as well as prepayments for unspecified purposes. All tracts are classified according to their prepayment rates in the full universe of (FHA-insured purchase money) loans. Next, the default rate is calculated for all loans within those tracts that fall into each class of prepayment rates. Each row in Table 4 corresponds to a different category of tract prepayment rates, and each of the first four columns gives the default rates for a race/ethnic group within the corresponding tracts; the last four columns show the percentage of loans by race falling into tracts within each prepayment class. We see, for example, that in tracts having no prepayments, 5.09 percent of all loans, 6.94 percent of black loans, 5.95 percent of Hispanic loans, and 2.72 percent of "other" (nonblack and nonHispanic) loans defaulted. Moreover, these tracts contained 4.92 percent of all loans, 9.66 percent of black loans, 6.8 percent of Hispanic loans, and 3.56 percent of "other" loans.

Table 4 indicates that default rates for "others" are always less than those of blacks and Hispanics. Hispanic default rates tend to be higher than those of blacks in tracts at the lower end of the prepayment rate categories, but the ordering is reversed at the upper end of the prepayment rate categories. Within each race/ethnic group and overall, default rates tend to be higher in tracts with lower prepayment rates. The latter apparent regularity --- higher prepayment rates accompanying lower default rates --- could be traceable to a common set of underlying factors that lead, on the one hand, to higher default rates and, on the other hand, to lower prepayment probabilities. The empirical work below may shed some light on this possibility.

2.3. Statistical Models and Methods

2.3.1. A Default Model

In this section we turn to statistical analysis designed to isolate the effects of neighborhood on the probability of default of individual loans. We begin by controlling for a variety of default-related factors related to the loan, the economic environment, and the borrower. Included are time-varying factors designed to pick up important changes that evolve over the lifetime of the loan, some of which are preprogrammed and thus known with relative certainty at the outset and others which are highly uncertain at loan origination.

The statistical model used to estimate the impact of default related factors is the proportional hazard model which, for a single risk (say, default), gives the conditional probability of default at each point in time given that a loan has survived to that point in time. The proportional hazard (in continuous time) is represented as

$$\lambda(t) = \lambda(0)\exp(x'(t)\beta) \tag{1}$$

where $\lambda(t)$ is the hazard function (in continuous time), $\lambda(0)$ is the baseline hazard, $x(t)$ is the vector of observables at time t, and β is a vector of unknown parameters. For the empirical implementation here, time-varying data elements are assumed to remain constant over one month intervals.[29]

2.3.1.1 Possible Advantages and Disadvantages of the Empirical Approach

In the current context, there are two principal advantages of the hazard specification over simple dichotomous dependent variable models, such as probit and logit, that recognize only whether or not each loan defaulted over the observation interval. First, the logit and probit approaches cannot properly recognize time-varying characteristics, for such models either evaluate these time-varying characteristics at only a few "representative" points, or they attempt to parameterize the time series with a few parameters that will typically not capture all features of the time series. In contrast, the hazard approach allows arbitrary changes in time-varying characteristics for each period (months in this case) over the observation interval. The hazard approach thus offers more appropriate treatment of dynamic processes in particular. Second, the hazard approach easily accommodates the censoring that occurs when the observation window on a loan closes while the loan remains active, in which case we do not know when, if ever, the loan will default. This feature permits us to use all observations on each active loan in the sample, even if the potential observation window differs wildly from loan to loan.[30]

Although the hazard approach offers potential advantages over simple logit or probit, it does not solve all potential statistical problems in estimating models of default. In particular, some

[29] As noted, weighting is used to accommodate stratification.

[30] Despite the theoretical advantages of the hazard approach, we have sometimes found empirically that default predictions using a logit model are generally no worse than those using the hazard.

researchers have pointed out a number of possibly serious difficulties in applying standard single-equation methods to estimate default relationships, and these problems would not be solved by utilizing a hazard approach. Yezer, et al (1994)[31] and Rachlis and Yezer (1993)[32] discuss in particular the issues of simultaneous equations bias and sample selection bias. In the context of estimating default relationships, simultaneous equations bias may arise because borrowers select the terms of their loans with an eye towards what underwriters will accept, and underwriters in turn look to anticipated default probabilities for guidance on what terms to accept from each borrower. Sample selection bias may arise in estimating the default relationship because that relationship is observed only for those who receive loans, and the set of those receiving loans is dictated by anticipated default probabilities.

One might have different views of these critiques of the standard estimation framework. On the one hand, one could argue that this work is too pessimistic, both in general and in the context of estimating default models in particular. With regard to simultaneous equations bias in mortgage terms, it is reasonable to assume that borrowers choose mortgage terms with an eye towards what underwriters will approve, and the result will be that mortgage terms, such as LTV, will vary with borrower characteristics, like asset levels. Simultaneous equations bias arises under the assumption that the default index used by underwriters enters the equation that determines ultimate mortgage terms, while mortgage terms themselves enter the equation for the default index. More generally, the presence of unobservables in the default index that are correlated with mortgage terms will lead to simultaneous equations bias. Notice, however, that if mortgage terms are determined by factors that affect default and that are observed by the researcher, but mortgage terms are not determined by unobserved factors that enter the default index, then there is no simultaneous equations bias. For this reason, the observation that

---

[31] Yezer, Anthony M.J., Phillips, Robert F., and Trost, Robert P., "Bias in Estimates of Discrimination and Default in Mortgage Lending: The Effects of Simultaneity and Self-Selection," *Journal of Real Estate Finance and Economics*, 9: 197-215 (1994).

[32] Rachlis, Mitchell B., and Yezer, Anthony M.J., " Serious Flaws in Statistical Tests for Discrimination in Mortgage Markets," *Journal of Housing Research*, Volume 4, Issue 2, 315-336 (1993).

mortgage terms vary with borrower characteristics does not in itself imply simultaneous equations bias in estimating default probabilities.

The potential sample selection problem is in some ways similar. The selection problem arises if applicants are selected to receive loans partly on the basis of default-related factors that are unobserved to the analyst.[33] Notice, however, that selection of borrowers strictly on the basis of variables that are observed to the analyst and are already included in the default model need not cause a sample selection problem. As with the potential problem of simultaneous equations bias, the importance of sample selection bias will presumably vary with the amount and quality of underwriting data available to the analyst. In particular, if the analyst has most of the information available to the underwriter and incorporates this information properly in estimating the default relationship, problems of simultaneous equations bias and sample selection bias may be less important.

On the other hand, one may take the view that these problems are likely to be serious, in part because there may be numerous factors available to underwriters but not to analysts using even the best of data bases. Even aside from the potential problems of sample selection bias and simultaneous equations bias, there may more generally be omitted variables correlated with the included explanatory variables, and the resulting omitted variable bias could be substantial. Moreover, there may be other problems as well. For example, the appropriate functional form for estimation is unknown; as discussed above, a simple dichotomous dependent variable estimation procedure, such as logit or probit, is highly unlikely to be correct. Many of the variables may be measured with error; heteroskedasticity may be present in the unobservables that enter the default index function; observations may not be independent but instead be correlated within areas; and the list goes on.

Our view is that the approach taken in this paper is likely to be better than in prior work in some important dimensions, but we clearly cannot reject the possibility that serious statistical problems remain. The inclusion of a credit score seems likely to compensate partially for the

---

[33] Sample selection, and possible bias as a result, could also arise at earlier stages of the loan selection and approval processes.

important omission of such data in most prior work, and the use of a hazard model is likely to do a better job of estimating dynamics than would the traditional logit or probit approach. Nonetheless, there are surely other omitted variables; we cannot be sure that the proportional hazard specification is correct; and so forth. Hence, while we hope to offer improvement, we cannot claim perfection.

2.3.1.2. A Motivation for the Structure of the Hazard Model

An option-based model of default provides a useful starting point for parameterization within this statistical framework, though we later deviate from this approach in several ways. Letting time t be measured relative to the date of mortgage origination and assuming that the prepayment option is costless to pursue, default occurs when V_t, the market value of the mortgage at time t, exceeds the sum of H_t, the value of the home at time t, and C_t, the costs of mortgage default at t (expressed as a fraction of house price). That is, default occurs when (in discrete time notation)

$$V_t > H_t(1+C_t)$$

or, in logarithmic terms,

$$\ln(V_t) > \ln(H_t) + \ln(1+C_t) \qquad (2)$$

Ignoring uncertainty over interest rates, future prepayment possibilities, and the like, the current value of a fixed-rate mortgage may be expressed as the discounted value of future mortgage payments using the current market rate, or (with continuous discounting):

$$V_t = (M/r_t)[1-\exp(-r_t(T-t))] \qquad (3)$$

where M is the monthly mortgage payment, r_t is the current market rate at t, and T is the term of the loan.[34]

Optimal and costless refinancing of the loan (if available) will prevent the value of the loan from exceeding the principal balance at t, B_t, but the valuation equation above continues to hold for loans that remain active whether or not refinancing would be optimal. Because the principal balance B_t is the value of the loan evaluated using the note rate (presumably the market rate at loan origination), r_0, we may rewrite the valuation equation as

$$V_t = B_t(r_0/r_t)[(1-\exp(-r_t(T-t)))/((1-\exp(-r_0(T-t))] \qquad (4)$$

or in logarithmic terms as

$$\ln(V_t) = \ln(B_t) + \ln(r_0/r_t) + \ln[(1-\exp(-r_t(T-t)))/((1-\exp(-r_0(T-t))] \qquad (5)$$

For simplicity in estimation, we represent the final term in the latter equation as simply proportional to ln (T-t).

For ARMs, the mortgage rate and mortgage payment adjust to keep the current-rate-discounted value of remaining mortgage payments equal to the current principal balance, aside from lags and limits on the size of adjustment permitted. Thus, for ARMs, the latter valuation equation reduces to

[34] An implication of this and other option-based default models is that a reduction in interest rates after loan origination should lead to an increase in the market value of the loan and, absent refinancing possibilities, to an increase in the incentive to default. Not all would agree that, as a practical matter, such a reduction in interest rates is an important determinant of default behavior. Additional discussion of this point is offered below.

$$\ln(V_t) = \ln(B_t)$$

The components of the proportionate costs of default (C_t) are unclear, but possible proxies include income, age, and credit scores, to the extent that the latter reveal past willingness to engage in credit-damaging behavior. Elements of affordability, such as the front-end ratio, are entered into the hazard specification in ad hoc fashion. Trigger events — such as loss in income due to unemployment — are often alleged to greatly increase the probability that an "in-the-money" default option will be exercised; proxies for these events may also be entered into the specification in ad hoc fashion.

If refinancing is not freely available to holders of fixed rate mortgages, then default need not necessarily occur even when inequality (2) holds. In particular, even if refinancing is completely prohibited for whatever reason, the holder of a mortgage valued at more than the market value of the home would wish to sell rather than default as long as the market value of the home (net of sales expenses) exceeds the principal balance of the existing mortgage. The existence of additional (monetary or nonmonetary) costs of default would further reinforce the incentive to sell rather than default.

Although the possibility of selling the home may serve an important role in reducing default probabilities for borrowers facing difficulties in refinancing, it is surely possible to imagine cases in which difficulties in refinancing lead to defaults that would not otherwise occur. Consider, for example, a case in which the borrower wishes to take advantage of lower interest rates by refinancing his or her mortgage and but is prevented from doing so by reason of discrimination. Sales expenses may exceed the owner's equity (plus costs of default), and the owner may thus find it optimal to default. Perhaps more importantly, homeowners wishing to refinance, but unable to do so when they would otherwise qualify for refinancing, may later face a decline in home value that induces them to default. Had they been able to refinance earlier, they may have been willing to hold onto their home through the subsequent price decline; because they were unable to refinance earlier, however, they later face a situation in which they would no longer qualify for refinancing and default is optimal. Whether cases like these arise frequently enough to make default probabilities rise as a result of reduced refinancing opportunities is, of course, an empirical question.

2.3.2. A Prepayment Model

Because one purpose of the analysis is to examine a possible link between default behavior and refinancing probabilities, specification of the probability of refinancing is required as well. As noted, the FHA data used here permit FHA refinancing to be separated from other types of prepayment, and thus we specify and estimate two prepayment probabilities, one for each of the two observable types of prepayment activity.[35] We assume that each of the two conditional prepayment probabilities follows a proportional hazard model. That is, the form of each conditional prepayment probability is that given in Eq. (1) above.[36]

Prepayments can occur for a variety of purposes, all of which should be considered in developing a specification. Prepayment may occur, for example, because the borrower has found a job elsewhere and decides to change residence. The FHA data contain a few correlates of job-related geographic mobility, such as age of the borrower. One might also expect income to be positively related to the probability of long distance job-related changes of residence.

Prepayments may also occur because the borrower faces anticipated or unanticipated (positive or negative) changes in wealth or in the cost of homeownership. Unanticipated increases in cost seem especially likely when mortgage rates rise to holders of ARMs; this effect

---

[35] As noted above, it is unfortunate that we are unable to distinguish refinancing with conventional sources from other kinds of prepayment activity within the set of loans that prepay for purposes other than for FHA refinancing.

[36] A study by Deng, Quigley, and Van Order uses a minimum distance estimator to estimate a competing risks proportional hazard model of default and prepayment. See Deng, Yong-Heng, Quigley, John, and Van Order, Robert. "Mortgage Default and Low Downpayment Loans: The Costs of Public Subsidy," National Bureau of Economic Research, October 1994. Here we estimate separate proportional hazard models for each prepayment type and for default; there is no guarantee that the sum of the predicted probabilities does not exceed one. To accommodate the dependence of each choice on the characteristics of the other choices, variable lists are identical for all models. We do not, however, build in any covariance across unobservables in the three models. For estimating each model, the at-risk set in each month is all loans that have not yet defaulted or prepaid.

may be represented by $\ln(M_t / M_0)$, where M_t is the monthly mortgage payment at time t and M_0 is the mortgage payment at mortgage origination. Although we have no direct measures of post-origination changes in wealth, higher unemployment rates may be a weak proxy if higher unemployment rates result in higher unemployment probabilities for individual borrowers, and this change affects total wealth as perceived by the borrower.

Prepayments for the purpose of refinancing a loan --- FHA or otherwise --- might be expected to occur among holders of fixed-rate mortgages when mortgage rates decline. Under such conditions, the incentive to refinance is increasing in the size of the existing mortgage balance (B_t), the ratio of mortgage rates at origination to current mortgage rates (r_0 / r_t), and the length of the remaining loan term, T-t. Although holders of ARMs automatically benefit (with a lag) from declines in mortgage rates, there may still be incentives to refinance ARMs in order to lock in lower rates. If so, the incentive to refinance might be again expected to vary with the current mortgage balance, the ratio of mortgage rates at origination to current mortgage rates, and the length of remaining loan term. Even if rates do not decline after origination, however, a steeper yield curve may reveal higher anticipated rates in the future, which may lead holders of ARMs to refinance with a fixed-rate loan so as to lock in a currently lower rate.

In practice, prepayment for refinancing purposes can occur only when the borrower wishes to refinance and the borrower is found qualified to do so. Thus, we might expect prepayments to depend upon variables that affect the ability to obtain refinancing, but these variables would include many if not all of the same variables that are included in the default hazard. Moreover, because default probabilities may depend on refinancing incentives that remain unexercised, variables appearing in the prepayment hazards should appear in the default hazard as well. In addition, the choice between the two kinds of prepayment presumably depends on the characteristics of both choices. Thus, we arrive at empirical prepayment hazard models containing the same set of variables as those in the default hazard model.

2.4. Variable Definitions and Concepts

Table 5 lists variable names and provides brief definitions. The variables fall into five categories that are distinguished by the nature and source of information: first, "neighborhood controls," which characterize the tract or the MSA; second, "demographic and credit characteristics" (race, age, credit (FICO) score); third, "measures of financial resources and costs," consisting of income and asset information, as well as the (log of the) front-end ratio; fourth, "characteristics of the property, the mortgage, and the interest rate environment," including changes therein; and, fifth, a set of "miscellaneous variables" that include various functions of mortgage duration.

The first group of variables --- "neighborhood controls" --- will be a major source of interest, for this group contains the tract characteristics that may be reflected in default rates. All of the tract-level variables are obtained from the 1990 Census and represent a single point-in-time snapshot of the tract. The monthly MSA unemployment rate is produced by BLS.

The second group — "demographic and credit characteristics" --- includes only variables measured at the time of origination. Age and race are obtained from FHA files. The FICO score, which characterizes an individual's credit history, is obtained from either Equifax or Trans Union. More specifically, when an individual (borrower or coborrower) had more than one FICO score present, we arrived at a single "operational" score for that person by taking the minimum of the two credit bureau scores. When both the borrower and coborrrower had one or more FICO score readings, we averaged the operational scores for borrower and coborrower.[37]

The third group of variables --- "measures of financial resources and costs" — consists of FHA variables measured at the time of loan application. Note that the effect of assets, measured as assets after closing relative to the monthly mortgage payment, is entered as a spline — a series of linear segments (only two segments in this case) joined at the endpoints. The particular breakpoint for the spline was chosen with the aid of plots of log odds ratios,[38] plots of probabilities of default, and/or experimentation with alternative breakpoints.

The fourth group — "characteristics of the property, the mortgage, and the interest rate environment" — includes a variety of measures that arose in the discussion of default and

[37] This procedure is fairly similar to that developed for a study of mortgage scoring.

[38] This exercise was performed in a different study using similar data. The spline was not reestimated for the two kinds of prepayment hazards.

prepayment models above. The log of the mortgage balance (lnbal) is calculated from the FHA data on the initial principal balance, the note rate, and the term of the loan.[39] Home values in each month (used in calculating lnhval) are obtained by updating the initial value of the home using post-origination MSA-level house price growth calculated from the quarterly Freddie Mac house price index series. Note that the difference between lnbal and lnhval is the (log of the) contemporaneous loan-to-value ratio; permitting these variables to enter the estimation equation separately allows the numerator and denominator of LTV to have separate effects. Market mortgage rates (for use in computing lintrat) are calculated using internal rates of return on 30-year conventional fixed-rate mortgages; the calculation scheme includes both the note rate and points and assumes that prepayment occurs at 10 years. The variable rtdiff is the slope of the yield curve, measured using monthly values of 30-year and 1-year constant maturity Treasuries. The variable logpirto is the (log of the) ratio of the current mortgage payment to the original mortgage payment (principal and interest); as such, the variable assumes the value zero for fixed-rate loans.

The relative house price variables merit additional explanation. The variable HPrelPW is the ratio of the sales price of the home relative to the reference home price in the area, as given by the PricewaterhouseCoopers median home price series.[40] When the Price-Waterhouse median price series is unavailable, we set HPrelPW to zero and instead measure relative house prices with the variable HPrelLL.[41] The latter construction measures the area reference house price by dividing the area FHA loan limit by 0.95. Because FHA loan limits are intended to be 95 percent of the area median house price, HPrelLL is effectively sales price divided by the area median house price. For those relatively rare loans in areas for which (a) the Price-Waterhouse series is

[39] For ARMs, we calculate the current principal balance by annually updating the original note rate using the contemporaneous one-year constant maturity Treasury rate, recalculating mortgage payments accordingly, and calculating the implied mortgage balance in each month.

[40] The PricewaterhouseCoopers median home price series is briefly described in the *MMI Fund Analysis FY 1998,* an actuarial review by PricewaterhouseCoopers LLP.

[41] More precisely, HPrelPW is set to zero when HPrelLL is used to measure relative price, and HPrelLL is set to zero when HPrelPW is used to measure relative house price.

unavailable, and (b) the FHA loan limit is at the legislative maximum or minimum, and is thus constrained so that it may no longer accurately measure median area house prices, we set both HPrelPW and HPrelLL to zero; in such a case, we also set an indicator (LLmax or LLmin) to unity.

The fifth group --- "miscellaneous variables" --- includes a spline in mortgage duration at six months, as well as the square of duration. These controls attempt to capture unexplained changes associated with mortgage duration. An indicator for the applications or closings in 1994 permits a shift in the intercept for the more recent endorsements. Finally, MSA indicators are introduced to permit differences in intercepts across MSAs.[42]

The discussion of the parameter estimates below provides additional explanation and justification for individual variables.

Sample means by point-in-time status and by race are presented in Table 6. That is, for each loan in the estimation sample, each month is classified as one in which the loan remains active, enters default, or enters prepayment (of either kind); these are the "point-in-time status" classifications. Notice that each defaulted loan has only a single month classified as a point-in-time status of default, and each prepaid loan has a single prepaid point-in-time status; all other months are classified as active point-in-time status. The columns in Table 6 give the means of each variable for each race group calculated over all months falling in the appropriate point-in-time status.[43] For variables that do not change or are not re-measured after loan origination, the focus on point-in-time status is no better than a focus on ultimate loan status — default or prepayment. For variables that change over the course of loan duration, the focus on values in the month of prepayment or default could be more revealing.

[42] Note that the inclusion of MSA dummies makes it unnecessary to deduct MSA means or other MSA-specific values from tract-level measures. That is, estimates (other than the intercept) would be unaffected by doing so. The presence of MSA dummies also removes the need to introduce additional indicators for areas that were impacted by the California recession.

[43] Note that the means for FICO scores (fico) appear low because the zeros (which replace missing values) are included in the calculation. The mean for the variable NOfico can be used to recover the mean for loans with nonmissing FICOs.

2.5. Statistical Estimates of Default and Prepayment Hazard Models

2.5.1. Estimates of Basic Hazard Models of Default and Prepayment

Panels A, B, and C of Table 7 present estimates of default and both prepayment hazard models; the estimated models use sample data from both 1992 and 1994 loans. The two kinds of prepayment hazards are distinguished as "FHA refinance" and as "other prepayment." As noted, the latter includes all kinds of prepayment activity other than refinancing through FHA. The specification used in these models excludes most neighborhood characteristics, the individual race indicators, and the credit scores. It does, however, include many time-varying characteristics, and these distinguish these models from many of the typical default and prepayment models relying solely on data measured at loan application or origination. A brief discussion will illustrate that most effects are as anticipated.

2.5.1.1. Default

Turning first to the default model in Panel A, neither increases in the MSA unemployment rate (unemprt) nor differences in the age of the borrower (borage) appear to matter in default behavior. Holding constant the remaining factors (including the log of the front-end ratio), the effects of individual income (loginc) cannot be statistically distinguished from zero.[44] Analogously, the effect of changing the front-end ratio cannot be statistically distinguished from zero. Additional assets after closing (RSVpmts) reduce the default rate until one has four monthly payments in reserve, after which there is essentially no marginal effect of additional asset holdings (the sum of the coefficients on RSVpmts and RSVpmt4 is approximately zero). Smaller principal balances (lnbal) are associated with lower default rates. Note that differences

---

[44] Notice that because the front-end ratio contains income in the denominator, the total effect of log income is the coefficient on loginc minus the coefficient on logfront. Because loginc and logfront are both included, however, the total effects of log income and of log monthly mortgage payments (the numerator of the front-end ratio) are unconstrained.

in principal balances after loan origination may result from different mortgage terms at origination, different note rates at origination, variation in ARM rates after origination, differences in seasoning, or differences in initial mortgage amounts. Notice also that the estimated impact of lnbal holds fixed contemporaneous house values (and the remaining explanatory variables), and thus increases in lnbal should be viewed as increases in the numerator of contemporaneous LTV. Higher contemporaneous house values, holding constant the contemporaneous principal balance (and thus reducing contemporaneous LTV), are associated with lower default rates, not surprisingly.[45] Default rates do not appear to be higher on condominiums.

The estimates show that by either measure (HPrelPW or HPrelLL) higher house prices relative to the area reference price are associated with lower probabilities of default. This finding may reflect differences in wealth of those buying relatively more expensive homes — beyond the effects captured by the other wealth-related measures. The relative house price variables may convey other information as well, however. Given that FHA-insured homes tend to be lower priced than conventionally insured homes, the more expensive FHA-insured homes in an area may be closer to the heart of the overall house price distribution than are less expensive FHA-insured homes. A thinner market for lower priced homes may increase time spent on the market, and the resulting rise in the implicit cost of selling may increase the likelihood of default.

Increases in the length of time remaining on the mortgage (lnhorizn), conditional on duration as reflected in the spline and quadratic in duration (t, t6, tt), and conditional on mortgage balance (lnbal), appear to increase default, but the effect is not measured very precisely. Given the remaining controls, this variable may be picking up in part the sorting induced by borrowers opting for mortgage terms of different lengths; more specifically those choosing mortgages with shorter terms are less likely to default. The failure to obtain precise estimates may in part reflect the paucity of short-term FHA mortgages.

The remaining substantive variables that measure movements in mortgage rates from the time

[45] The fact that the coefficient estimates on lnbal and lnhval are nearly equal in magnitude and opposite in sign suggests that it may be appropriate to use (the log of) LTV in the default model, thus imposing the usual restriction. We did not conduct the appropriate test.

of origination (lnintrat and lnintarm), the slope of the yield curve (rtdiff and armrtd), and the post-origination change in mortgage payments for ARMs (logpirto) cannot be statistically distinguished from zero.

Increases in duration (t, tt, t6) lead to conditional default rates that rise rapidly with duration initially, but the rate of increase tapers off substantially at six months. The conditional default probability begins to decline at about 35 months, other things the same.[46] Finally, there does not appear to be a significant shift in the intercept in the more recent data (year94).

2.5.1.2. Prepayment Other Than FHA Refinancing

Turning briefly to the estimated effects in the model for other prepayments (*i.e.*, for all purposes other than for FHA refinancing) in Panel B, we see first that variation in the MSA unemployment rate (unemprt) does not appear to matter in this kind of prepayment behavior. Increases in the age of the borrower (borage), however, are associated with lower prepayment possibilities, perhaps a reflection of age-related reductions in geographic mobility in general. Holding constant the remaining factors (including the log of the front-end ratio), higher incomes (loginc) lead to higher prepayment probabilities.[47] It is unclear whether this income impact is associated with increased mobility of higher income borrowers, increased probabilities of qualifying for conventional refinancing, or some other source. Higher front-end ratios (logfront) also lead to higher prepayment probabilities, perhaps because those with additional financial burdens at the outset are more likely to find themselves too burdened to continue with home ownership in the event of any unanticipated shock to income. Under this interpretation, however, it is somewhat surprising that additional financial reserves after closing (RSVpmts) have no discernible effect

[46] Note that other explanatory variables are held fixed when examining the estimated default effects of duration. For this reason, this estimated path of conditional default probabilities over loan duration may not be identical to the path of empirical conditional default probabilities when nothing else is held fixed.

[47] Again, because the front-end ratio contains income in the denominator, the total effect of log income (holding constant monthly payments) is the coefficient on loginc minus the coefficient on logfront; the result is again a positive, but much smaller, effect.

on prepayment probabilities.[48] Higher principal balances (lnbal) are associated with lower prepayment probabilities. One possible explanation for the latter finding is that the negative impact of higher contemporaneous principal balances (holding fixed contemporaneous home value) on the probability of qualifying for conventional refinancing swamps a positive impact of higher balances on the gains from conventional refinancing. Higher contemporaneous house values are associated with higher prepayment probabilities, again perhaps reflecting their impact on the probability of qualifying for refinancing, or perhaps reflecting increased wealth that induces borrowers to move up to higher priced homes. Prepayment probabilities appear to be somewhat higher on condominiums.

The estimates show that higher house prices relative to the area reference price generally matter little in prepayment behavior. Increases in the length of time remaining on the mortgage (lnhorizn) lead to higher prepayment probabilities, as would be expected given that gains to refinancing are higher when there is a longer time left over which to reap the benefits. Also not surprising is the effect of changes in mortgage rates. Reductions in mortgage rates after origination (increases in lnintrat) have a precisely measured positive effect on prepayment probabilities. It may also be the case that rate reductions induce some homeowners to move into more expensive homes. As anticipated, for holders of ARMs, the total effect is smaller (lintarm is negative, though imprecisely measured) but still positive (the sum of the coefficients on lnintrat and lnintarm).[49] Also as expected, larger increases in ARM mortgage payments after origination (logpirto) increase the probability of prepayment. A steeper yield curve (rtdiff) leads to higher prepayment probabilities, and that effect is stronger (though imprecisely measured) for holders of ARMs (armrtd).

The coefficients on the duration spline and quadratic (t, t6, tt) imply that conditional prepayment probabilities initially rise with duration, but the rate of growth tapers off at six months. Conditional probabilities eventually decline with duration (after about 38 months),

---

[48] Part of the problem here may be incorrectly chosen breakpoints on the spline.

[49] We had no success in permitting ARM effects to differ depending on whether contemporaneous mortgage rates were higher or lower than those at mortgage origination.

other things the same. There appears to be a significant shift in the intercept in the more recent data (year94).

2.5.1.3. FHA Refinancing

Turning finally to the estimated effects in the model for FHA refinancing (Panel C), we see that variation in the MSA unemployment rate (unemprt) does not appear to matter in FHA refinancing. Increases in the age of the borrower (borage) reduce refinancing probabilities. Holding constant the remaining factors (including the log of the front-end ratio), higher incomes (loginc) lead to higher refinancing probabilities, but the total effect of income conditional on the mortgage payment (*i.e.*, the coefficient on loginc minus that on logfront) is slightly negative. One possibility is that while higher incomes increase the probability of qualifying for an FHA-insured loan, higher incomes also increase the probability of qualifying for a preferable conventional loan. Higher front-end ratios (logfront) lead to higher refinancing probabilities, perhaps because the demonstrated ability to contend successfully with greater financial burdens is a positive signal to lenders. Additional financial reserves after closing (RSVpmts) at first increase and later decrease FHA refinancing probabilities, which may again reflect the ability of wealthier borrowers to qualify for, and shift to, conventional loans.[50] Higher principal balances (lnbal) seem to be associated with higher FHA refinancing probabilities, as might be expected, but the effect is not statistically significant by conventional standards. There appears to be no statistically significant effect of greater contemporaneous house values, with perhaps the increased ability to qualify for preferable conventional loans dominating the increased ability to qualify for FHA loans. FHA refinancing probabilities appear to be unrelated to condominium ownership.

The estimates show that higher house prices relative to the area reference price generally reduce FHA refinancing probabilities; once again a shift to conventional refinancing is one possible logical explanation. Increases in the length of time remaining on the mortgage

---

[50] Here again breakpoints on the spline may be incorrect.

(lnhorizn) lead to higher refinancing probabilities, as expected. Again conforming to expectations, a reduction in mortgage rates after origination (increases in lnintrat) increases FHA refinancing probabilities. As anticipated, the latter effect is reduced for holders of ARMs (lintarm is negative), but the estimated differential is very small.[51] More sizable increases in ARM mortgage payments after origination (logpirto) increase the probability of FHA refinancing. Unexpectedly, a steeper yield curve (rtdiff) leads to lower FHA refinancing probabilities; that effect is more pronounced for holders of ARMs (armrtd). Shifts to conventional financing may again be part of the story.

The coefficients on the duration spline and quadratic (t, t6, tt) imply that conditional FHA refinancing probabilities rise with duration, but the rate of increase tapers off substantially at six months; conditional probabilities decline with duration after about 12 months, other things the same. There does not appear to be a significant shift in the intercept in the more recent data (year94).

2.5.2. Adding Neighborhood Characteristics to Default and Prepayment Models

In this section we modify the basic models introduced in the last section by adding neighborhood characteristics obtained from 1990 Census data: the fraction of the tract population that is black (trtblk), the fraction of the population that is Hispanic (trthisp), and the (log of the) median income of the tract (lnincmed). Panels A, B, and C of Table 8 show the hazard estimates obtained after this modification to the list of explanatory variables. According to the estimates in Panel A, loans in neighborhoods (tracts) with greater black representation have higher default probabilities, while representation of Hispanics appears to have no statistically significant impact. Increases in tract income appear to reduce the probability of default. Other coefficients appear to be only modestly affected by the inclusion of these tract characteristics.

A finding that the racial composition and income of the neighborhood matter in individual

[51] As noted, we had no success in allowing ARM effects to depend on whether contemporaneous mortgage rates were higher or lower than rates at origination.

default behavior is consistent with the findings of Van Order and Zorn (1995) in the context of conventional loans. As noted there, tract income may appear to be more important than individual income because individual income at loan qualification may contain a substantial transitory component not present in the tract level measure.

Turning to prepayments for purposes other than FHA refinancing, the estimates in Panel B indicate that an increase in the percentage of the tract that is black is associated with reduced prepayment probabilities, while higher tract income may increase prepayment probabilities. Hispanic representation in the tract appears not to matter.

The estimates in Panel C show that tract characteristics have no statistically significant effect on FHA refinancing probabilities.

Although we do not think it especially enlightening to include additional tract characteristics that have no well-founded causal role in default or prepayment behavior, it is nonetheless of some interest to examine the effects of including some additional tract characteristics. Panels A, B, and C of Table 9 modify the last specification by adding more neighborhood characteristics measured in the 1990 Census: the tract proportion of owner-occupied units that are not mortgaged (misleadingly named propmort), the fraction of the population that did not move within the previous five years (*i.e.*, from 1985 to 1990) (misleadingly named propmove), and the fraction of owner- and renter-occupied units that are in one-unit structures (prop1unt). In Panel A we see that only propmove has an estimated effect that passes typical standards of statistical significance; it indicates that higher fractions of nonmovers are associated with lower default rates. Perhaps tracts with more stable populations are wealthier and therefore less likely to default, or perhaps additional turnover of homeowners itself generates more default activity. We shall return to this point later. The remaining coefficient estimates are largely unchanged when these additional tract characteristics are included.

Turning to prepayment for purposes other than FHA refinancing (Panel B), we see again that only propmove is statistically significant at conventional levels. Perhaps not surprisingly, higher proportions of nonmovers are associated with lower prepayment probabilities. The inclusion of the additional area measures reduces somewhat the estimated impact of trtblk and increases the estimated effect of tract income. There is little impact on most of the remaining estimates.

Panel C shows that among the additional neighborhood characteristics, only prop1unt is statistically significant at conventional levels. A larger share of one-unit structures is associated with higher FHA refinancing probabilities, a finding for which we have no ready explanation. Tract race and income characteristics remain statistically insignificant.

2.5.3. Adding Individual Characteristics to Default and Prepayment Models

Models in the last section control for those factors that might be available to researchers dealing with data on conventional mortgages. In this section we introduce important individual characteristics that are not often available to researchers. In particular, race information is not generally present in databases on conventional loans, and credit scores are typically not available for either conventional or FHA loans.

Panels A, B, and C of Table 10 add only the information on the race of the borrower — black and hisp. In Panel A (default) we see that, perhaps surprisingly, neither race effect on default is statistically significant at conventional levels; the impacts of tract racial composition and income remain, though the former is reduced somewhat in magnitude. According to this specification, race and income appear to operate at the tract level rather than the individual level.

Examining effects on prepayment other than for the purpose of FHA refinancing, Panel B of Table 10 shows, even more surprisingly, that adding individual race indicators changes the estimated tract black (trtblk) effects on prepayment from negative and significant to positive and insignificant. The effects of Hispanic representation in the tract and tract income increase in magnitude and become statistically significant. The individual race indicators themselves are highly significant and point to reduced prepayment activity for individual members of both groups.

Turning next to FHA refinancing, Panel C of Table 10 shows that adding individual race indicators leaves estimated tract race and income effects statistically insignificant. Of the two individual race impacts, only the negative Hispanic effect is statistically significant.

Next we introduce a measure of past credit performance, the FICO score. Panel A of Table 11 shows that the FICO score itself has a precisely measured negative impact on default probabilities. Individual race effects are absent by any reasonable standard of statistical significance, and the presence of the FICO reduces the estimated impact of trtblk, rendering it of questionable statistical significance. The latter change upon the introduction of the FICO readings suggests that one "explanation" that loans default more frequently in tracts that are more heavily black is that borrowers in these tracts have had, on average, poorer credit performance in the past.

One might question whether the latter finding constitutes an "explanation." That is, one might interpret the estimated FICO effect itself as saying that poor performance on credit in the past is associated with poor performance on paying off mortgages in the future. If differences across tracts in default behavior of individual loans are "explained" by the past credit performance of individual borrowers within those tracts, one might then ask why there are differences across tracts in the past credit performance of individual borrowers. We have no ready answer for this question. nor do we know why past credit performance is predictive of future credit performance.

Panel B of Table 11 introduces the FICO score in the prepayment (other than for FHA refinancing purposes) hazard. We see that higher FICO scores lead to higher prepayment probabilities, presumably in part because those with higher FICOs find it easier to qualify for conventional refinancing. In addition, we see that all other neighborhood effects and individual race effects remain largely unchanged. Individuals in more heavily Hispanic and higher income tracts still appear to have higher prepayment probabilities, while minority members themselves have significantly lower prepayment probabilities.

Panel C of Table 11 enters the FICO score in the FHA refinancing hazard. Higher FICO scores lead to higher prepayment probabilities. Again, higher FICO scores increase the probability of qualifying for refinancing, though presumably some of this enhanced ability to qualify permits borrowers to transfer to conventional loans, thus reducing the net impact on FHA refinancing probabilities. We see that all other neighborhood effects and individual race effects are little influenced by the introduction of FICO scores. Only the reduced FHA refinancing probability for individual Hispanics passes conventional statistical significance.

Before proceeding, it is worth noting that the default model presented in Panel A of Table 11 was explored in a few other dimensions. First, one of the salient features of the estimated default hazard is that individual race seems not to matter while the proportion of the tract that is black may matter. As in most data files, variables in the FHA data files are subject to error, and it would be shocking if race information were measured perfectly. Although we have no way of knowing the extent of error in race information, there is some possibility that tract racial composition is a better proxy for race of borrower than is the individual loan-level information. If so, there is some chance that individual race is important, but its effect is masked by measurement or reporting error. While there is no foolproof way to determine whether this possibility has been realized, we conducted one test that should provide a bit more information. The idea behind this test is as follows. Individual race information seems more likely to be in error when the reported race of the individual differs from that of the overwhelming majority of the population in the same tract; similarly, race information seems unlikely to be in error when the reported race is identical with that of the vast majority of the population within the same tract. Thus, we introduced two new interaction indicator variables for blacks and two for Hispanics. One of the black indicators was activated when reported race was black and the property was located in a tract that was at least 80 percent black; the other was activated when reported race was black and the property was located in a tract that was less than 10 percent black. Similarly, the two Hispanic indicators were activated (a) when reported race was Hispanic and tract representation was at least 80 percent Hispanic, and (b) when reported race was Hispanic and tract representation was less than 10 percent Hispanic. These additional four indicators, when used in the default model underlying Panel A of Table 11, yielded coefficients that could not be statistically distinguished from zero. Hence, on the basis of these tests, which are surely not definitive, we find no indication that measurement error lies behind the absence of individual race effects in the default model reported in Table 11.

We also explored in piecemeal fashion the possibility that race effects, at both the tract and individual level, differ across MSAs. Because our main focus is on default behavior, these explorations were again restricted to the default hazard. The ideal procedure would be to estimate a fully interactive model that would permit all effects to differ across MSAs. This rather cumbersome procedure was not attempted. Instead we ran separate estimation procedures in which one kind of tract or individual effect for a minority group was allowed to vary across MSAs while all remaining effects (other than the intercept) were constrained to be equal across MSAs. Thus, for example, one estimation procedure permitted trtblk to differ across MSAs while effects of trthsp, hisp, and black were constrained to be equal across MSAs. These experiments revealed no significant differences across MSAs in tract or individual race effects on default.[52]

Finally, earlier versions of the models were run with interactions between each individual and tract race variable, on the one hand, and duration, on the other hand. These experiments revealed no important changes in race impacts with loan duration.[53]

2.6. Interpretation and Conclusions

The evidence thus far permits answers to three questions raised at the beginning of the paper. First, once one controls for a variety of borrower- and loan-related factors in an appropriate econometric model of default, neighborhood effects related to income and, less clearly, to race do persist. In particular, even though some estimated neighborhood effects in earlier studies could be traceable to changes over time in house prices, principal balances, or unemployment rates that are inadequately accounted for in many studies, these controls do not completely remove the effects of tract income and (perhaps) tract racial composition on default probabilities of

---

[52] Using a single hazard that included both FHA refinancing and all other prepayments, we found that the individual black effect showed no significant differences across tracts, but trthsp effects differed in three MSAs, two MSAs had significantly different trtblk effects, and individual Hispanic effects differed in two MSAs.

[53] In earlier specifications, we also carried out some experiments to try to identify default effects flowing from neighboring tracts. After all, it is not clear that the census tract is the relevant neighborhood, and even if the tract appropriately delimits the neighborhood, characteristics of nearby neighborhoods may matter as well. We used the longitude and latitude of the centroid of each census tract in the Chicago MSA to calculate the distance between the centroids of each pair of tracts. Using a few different functions to weight characteristics of other tracts by distance from the own tract, we examined whether default probabilities at the loan level depend on the characteristics of not only one's own tract, but also the characteristics of other tracts. In limited experimentation, we were unable to find any impacts of neighboring tracts.

individual loans.

Although neighborhood impacts remain, their importance is dramatically affected by introducing controls for characteristics of individual loans, borrowers, and the economic environment. To illustrate this point, we provide estimated tract effects in the absence of all controls other than the duration-related variables, the year 1994 indicator, and the MSA indicators. Table 12 gives these estimates. Panel A contains the default model; Panels B and C provide the two prepayment models for completeness. Comparing Panel A of Table 12 with Panel A of Table 11, we see that the estimated default impact of trtblk falls substantially and becomes of marginal significance with the introduction of a host of controls (in Table 11), while the estimated impact of trthsp remains insignificant. The estimated effect of tract income falls in magnitude but remains significantly different from zero.

It is also interesting to note that the prepayment results in Panel B of Tables 11 and 12 show that the introduction of additional controls causes the estimated effect of trtblk to change signs, reduces the estimated effect of tract income, and increases the positive estimated impact of trthsp. A comparison of Panel C in Tables 11 and 12 reveals that with the introduction of additional controls, the FHA refinancing effects of tract racial composition remain insignificant, and the effect of tract income is rendered insignificant as well.

At the same time, we emphasize that our ability to account for time-varying explanatory variables is severely limited by data availability. Notice in particular that the contemporaneous house value lnhval is based, in part, on house price growth within the MSA. Tract level price change would, of course, be preferable but is not available in these data. This issue is revisited below.

The second question raised at the start of the paper was whether neighborhood characteristics, such as race and income, have effects on default that are separate and distinct from the effects of these same characteristics at the individual level. We find that greater tract representation of blacks is probably associated with higher default probabilities, but individual race effects --- black or Hispanic --- are absent. Hispanic representation at the tract level does not seem to matter in default behavior. Tract income does seem related to default behavior, while individual income has no discernible effect.

In this context, it is of interest to consider the arguments raised by Schill and Wachter (1993) in the context of discrimination in mortgage lending. They note that, given residential segregation of races and ethnicities, racial or ethnic discrimination at the level of the individual may have effects that are virtually identical to discrimination at the level of the neighborhood. Although it is true that residential segregation may make individual-level and neighborhood-level discrimination approximately equivalent in their effects, these two forms of discrimination are still empirically distinguishable as long as residential segregation is incomplete. That is, as long as races and ethnicities are not completely isolated (and as long as neighborhoods are properly identified), discrimination at the neighborhood level can be distinguished empirically from discrimination at the individual level. Partial but incomplete segregation makes the job more difficult in the same way that lack of orthogonality reduces the effective informational content of a given number of observations on other correlated explanatory variables, but empirical identification of separate effects is still possible in principle. In particular, incomplete segregation does not cause bias in estimated effects. In the case at hand, the fact of partial residential segregation does not introduce bias in distinguishing individual from neighborhood differences in default, but it does make the task more difficult.

The third question was whether there is evidence that differences in default probabilities reflect differences in the probability of refinancing, which may in turn be indicative of unequal access to refinancing funds. We find that default probabilities may be higher in tracts with heavier black representation, but even this differential is of questionable significance. If race-based redlining in refinancing were the culprit in generating higher default rates in more heavily black neighborhoods, we would expect to find that either FHA refinancing probabilities or other prepayment probabilities (to the extent they are indicative of conventional refinancing probabilities) would be lower for blacks. Instead we find that FHA refinancing probabilities, as well as other prepayment probabilities (which include conventional refinancing activities) tend to be higher, if anything, among loans in tracts with heavier minority (black or Hispanic) representation; only the Hispanic effect for other prepayment is statistically significant at conventional levels. The results here do not support the notion that race-based redlining or racial discrimination in refinancing lie behind race-related neighborhood differences in default probabilities,[54] but of course, these results should only be viewed as suggestive.[55]

We emphasize that this assessment of redlining pertains only to its potential role in explaining differential default behavior, and the evidence is essentially indirect. Moreover, this investigation is restricted to effects arising among holders of FHA mortgages; holders of conventional mortgages are excluded from the analysis. In addition, we do not account in any way for possible differences in the rate at which groups actually apply for refinancing. For these reasons, this discussion is not intended to comment on the possible existence of redlining in general.

We do find that individual Hispanic ethnicity is associated with lower FHA refinancing probabilities, and individual blacks and Hispanics have lower probabilities of other types of prepayment, which may or may not be the result of discrimination in refinancing against individual borrowers on the basis of race or ethnicity. Yet there are no corresponding default effects traceable to the race of individuals.

These findings suggest that we look elsewhere for possible explanations for the effect of tract income and tract racial (black) composition on default probabilities. We next consider some alternatives.

---

[54] Income-based "redlining" is a possibility, but of course there are other possible explanations for tract income effects on default and prepayment, including the potentially important role of missing variables. Note also that there appears to be no tract income impact on FHA refinancing probabilities; thus, any such refinancing impacts must apparently arise from conventional lending.

[55] One difficulty is that, for reasons given above, one might expect discrimination in refinancing to result in increased probabilities of prepayment for purposes of changing residence. This effect will tend to blur the impact of discrimination on other refinancing probabilities when, as here, data on other refinancing includes both refinancing (by other than FHA) and prepayments for other purposes. In addition, increased difficulty in refinancing in one sector (say, FHA) might also lead to increased refinancing activity in the other (say, non-FHA) sector as borrowers attempt to circumvent difficulties in (say, FHA) refinancing. However, because the increase in prepayment probabilities is derivative of the declining ability to refinance --- serving as one possible outcome for a borrower facing increased difficulties in refinancing --- as is the substitution of one avenue for refinancing for another, it seems very unlikely that difficulties in refinancing would yield increases in both FHA refinancing and in other prepayment probabilities.