

Decomposing the Foreclosure Crisis:
House Price Depreciation versus Bad Underwriting

Kristopher Gerardi, Adam Hale Shapiro, and Paul S. Willen

Working Paper 2009-25
September 2009

Decomposing the Foreclosure Crisis: House Price Depreciation versus Bad Underwriting

Kristopher Gerardi, Adam Hale Shapiro, and Paul S. Willen

Working Paper 2009-25

September 2009

Abstract: We estimate a model of foreclosure using a data set that includes every residential mortgage, purchase-and-sale, and foreclosure transaction in Massachusetts from 1989 to 2008. We address the identification issues related to the estimation of the effects of house prices on residential foreclosures. We then use the model to study the dramatic increase in foreclosures that occurred in Massachusetts between 2005 and 2008 and conclude that the foreclosure crisis was primarily driven by the severe decline in housing prices that began in the latter part of 2005, not by a relaxation of underwriting standards on which much of the prevailing literature has focused. We argue that relaxed underwriting standards did severely aggravate the crisis by creating a class of homeowners who were particularly vulnerable to the decline in prices. But, as we show in our counterfactual analysis, that emergence alone, in the absence of a price collapse, would not have resulted in the substantial foreclosure boom that was experienced.

JEL classification: D11, D12, G21

Key words: foreclosure, mortgage, house prices

Parts of this paper appeared earlier in Boston Fed Working Paper 07-15, titled “Subprime Outcomes: Risky Mortgages, Homeownership and Foreclosure.” The authors thank Chris Foote, who figured out how to identify subprime purchase mortgages in the Warren Group data using the HUD subprime lender list, which greatly simplified this project. They also thank Paul Calem, John Campbell, Jeff Fuhrer, Simon Gilchrist, Lorenz Goette, Robert King, Andy Haughwout, Zach Kimball, David Laibson, Andreas Lehnert, Atif Mian, Paolo Pellegrini, Karen Pence, Anthony Pennington-Cross, Erwan Quintin, Julio Rotemberg, Amir Sufi, and Geoff Tootell; audiences at Boston University, the University of California-Berkeley, Harvard University, MIT, Fordham University, Moody’s KMV, Freddie Mac, George Washington University; the Federal Reserve Banks of Atlanta, Boston, New York, Philadelphia, Kansas City, Richmond, and San Francisco; the Board of Governors of the Federal Reserve System; the Chicago Fed Bank Structure Conference; the Homer Hoyt Institute; the Rodney White Conference at the Wharton School; and participants in the Stanford Institute for Theoretical Economics, the Society for Economic Dynamics, the National Bureau of Economic Research Summer Institute, and the American Real Estate and Urban Economics Association midyear meetings for helpful comments and suggestions. They are grateful to the hospitality of Studienzentrums Gerzensee, where this paper was completed. They also thank Tim Warren and Alan Pasnik of the Warren Group and Dick Howe Jr., the registrar of deeds of North Middlesex County, Mass., for providing them with data, advice, and insight, and they thank Anthony Pennington-Cross for providing computer programs. Finally, they thank Elizabeth Murry for helpful comments and edits. The views expressed here are the authors’ and not necessarily those of the Federal Reserve Banks of Atlanta or Boston or the Federal Reserve System. Any remaining errors are the authors’ responsibility.

Please address questions regarding content to Kristopher Gerardi, Research Department, Federal Reserve Bank of Atlanta, 1000 Peachtree Street, N.E., Atlanta, GA 30309-4470, kristopher.gerardi@atl.frb.org; Adam Hale Shapiro, U.S. Department of Commerce, Office of the Chief Economist, Bureau of Economic Analysis, 1441 L Street, N.W., Washington, DC 20230, 202-606-9562, adam.shapiro@bea.gov; or Paul S. Willen, National Bureau of Economic Research and Research Department, Federal Reserve Bank of Boston, P.O. Box 55882, Boston, MA 02205, 617-973-3149, paul.willen@bos.frb.org.

Federal Reserve Bank of Atlanta working papers, including revised versions, are available on the Atlanta Fed’s Web site at www.frbatlanta.org. Click “Publications” and then “Working Papers.” Use the WebSubscriber Service (at www.frbatlanta.org) to receive e-mail notifications about new papers.

1. Introduction

There are two competing theories to explain the explosion of foreclosures in the United States in 2007–2009. The first focuses almost entirely on underwriting standards, and attributes the crisis to loans that borrowers had trouble repaying, either because they had bad credit and little income to begin with, or because the loans were unrealistically generous at the time of origination and later became unaffordable. There is persuasive *prima facie* evidence that is consistent with this theory. For example, subprime loans, which expanded credit by offering loans to borrowers who previously could not obtain any mortgage credit, and by offering large loans to those who previously could only obtain small loans, account for a disproportionate number of foreclosures. Subprime lending emerged as a major force in mortgage markets shortly before foreclosures began to accelerate, and were concentrated in communities that later became foreclosure hotspots. An alternative explanation focuses on house prices and points to the precipitous decline in prices that started in 2005 or 2006, depending on where in the country one looks, as an explanation for the crisis. There is *prima facie* evidence that is consistent with this theory as well. For example, until house prices began falling, subprime mortgages performed very well. In addition, states where house prices began declining sooner experienced foreclosure waves sooner than in states where prices began falling later.

In this paper, we assess the merits of the “poor underwriting” and “falling house price” theories. We argue that, in a sense, both explanations are true, but that prices have primacy. We find that had prices not fallen, we would simply not have had a major foreclosure crisis, regardless of whether lenders had lowered underwriting standards in 2003 and 2004. By contrast, the observed fall in prices would have generated a substantial increase in foreclosures, even if lenders had retained the underwriting standards that prevailed in 2002. To be sure, the increase in foreclosures would have been substantially smaller without subprime lending because as we show that subprime loans are far more sensitive to a decline in house prices than prime loans, but the foreclosure rate would still have been very large relative to historical levels, and would have been still considered a major public policy problem.

The central issue that any researcher must address in such an analysis is how to identify the effect of house prices on foreclosures. One natural explanation for the observed relationship is that the increase in foreclosures caused prices to fall. If that were the case, then the observed relationship between prices and foreclosures becomes evidence for the bad underwriting theory not the falling prices theory. The first contribution of this paper is to use a novel identification strategy to test this explanation of the crisis. We exploit within-town, within-time variation in household-level measures of housing equity to refute the theory that price falls resulted from differences in foreclosures over time and across communities. If foreclosures caused the fall in prices, we would expect to see little effect of prices on foreclosures within communities and within time periods, and large differences in the effect of prices on foreclosures across communities and over time due to simultaneity bias. In fact, we find that the observed relationship between prices and foreclosures is virtually identical whether we use variation across communities, within communities, within a given time period, or over time.

We use a dataset based on deed records in the Commonwealth of Massachusetts from 1989 through 2008, which has been previously used to study housing issues, most notably, in an influential paper on forced sales and foreclosures by Campbell, Giglio and Pathak (2009). The dataset was constructed by the Warren Group, and has both advantages and disadvantages over other datasets used to study foreclosure and mortgage performance. The main disadvantage is that we do not observe detailed borrower characteristics such as credit score, or loan characteristics such as whether the interest rate is adjustable or fixed. However, we do observe the identity of the lender and we use this information to match the data with the HUD list of subprime lenders in order to identify loans as subprime or prime. Using a subset of loans for which the data provider recorded more detailed loan information, we show that these loans conform to generally accepted

definitions of subprime loans. In addition, the sharp contrast between the performance of loans we identify as prime and those that we identify as subprime largely validates the HUD measure as a reasonable proxy for underwriting standards. The other significant disadvantage is that the dataset contains information for the state of Massachusetts only. However, we show that the Massachusetts experience in the current crisis is quite typical of the country as a whole and not dissimilar to that of more prominent foreclosure and subprime lending hot-spots like California.

Using the Warren Data confers many advantages over other commonly used datasets in the empirical housing literature. The first is that it contains the universe of all residential property and mortgage transactions in Massachusetts, and thus covers all types of loans. In contrast, the often-used First American, LoanPerformance (LP) dataset, a natural alternative, is almost entirely composed of subprime and reduced-documentation Alt-A loans, and contains only mortgages securitized by private institutions. More than half of the foreclosures in Massachusetts in 2007 and 2008 involved prime loans, and approximately two-thirds involved homes that were purchased with prime mortgages. The second advantage is that the dataset contains a long time series which, importantly, includes two substantial house price downturns. The first occurred between 1989 and 1992, while the second is the current crisis, which began in 2006. This stands in sharp contrast to other datasets which, because of left-censoring, can only be used with relatively recent data.¹ The long time series allows us to address the oft-made claim that this episode is fundamentally different from ones that preceded it. Finally, a third and significant benefit of the Warren Data is that it allows us to observe and match all liens on a given property, providing us with accurate measures of the cumulative loan-to-value (CLTV) ratios that are essential for calculating a borrower's equity position. The LPS dataset, by contrast, does not allow the researcher to match first and second liens, and furthermore, does not even flag the presence of a second lien.²

Another innovation in this paper is the use of what we will refer to as “homeownership experiences” rather than individual loans as the unit of analysis. A homeownership experience is a chain of mortgages that begins with a purchase mortgage and terminates either with the sale of the property and discharge of all outstanding liens, foreclosure, or with right-censoring. The advantage of homeownership experiences for our questions results from the fact that the typical homeowner refinances several times after purchasing a home. Some borrowers do this to take advantage of lower mortgage interest rates, and others to improve their homes, but for subprime refinances, the reason is often related to adverse household financial shocks (such as job loss, divorce, medical problems, etc.). This leads to a potential problem in analyzing the effects of subprime loans on foreclosure rates: borrowers who obtained subprime loans because of financial problems may have otherwise defaulted on their previous mortgages. Thus, we would attribute a foreclosure to a subprime loan, even though the default may have occurred anyway. In other words, both the foreclosure and the subprime loan share the same cause. One possible solution would be to look at loan-level data and exclude refinances, but since most subprime purchase loans end with a refinance and not foreclosure, such a tactic would underestimate the effect of subprime loans on foreclosures.

Our evidence that falling prices are at the center of the crisis does not mean that subprime lending did not play an important part. But, we show that the role of the subprime market is somewhat more subtle than its characterization in the poor underwriting explanation. In past work, we have shown that subprime mortgages were not inherently problematic or “designed to fail,” as some have argued, but in this paper we do find that they end up in foreclosure more frequently because of their higher sensitivity to falling house prices. We experiment with different specifications of a proportional hazard model and conclude that

¹The Lender Processing Services (LPS) dataset, for example, has broad coverage of the U.S. mortgage market (unlike the LP dataset), including a large prime, conforming subset, but is only usable starting with the 2005 vintage.

²The LP dataset has only partial information on the existence and size of second liens at origination

a dummy variable for whether the home was purchased with a subprime mortgage in a single model that includes both prime and subprime homeownerships is the most effective way to model the effect of subprime borrowing.

Our estimates show that a subprime homeownership is approximately six times more likely to end in foreclosure and, because we are using a proportional hazard specification, this implies that the effect of any other risk factor goes up by a factor of six when we compare prime and subprime loans. For example, our estimates imply that a homeowner with 20 percent negative equity is 10 times more likely to default than a borrower with 20 percent positive equity. Combining this with our estimate of the effect of a subprime loan implies that a subprime borrower with 20 percent negative equity is 60 times more likely to default than a prime borrower with 20 percent positive equity. It is easy to see from this how one might compare 100 percent LTV subprime borrowers in Massachusetts from 2005 and 80 percent LTV prime buyers from 2004, which correspond roughly to the former and latter examples in the last sentence, and conclude from the 60-fold difference in foreclosures that subprime was “toxic.” But, in fact, had the two loans been originated in the same year, with the same LTV, the difference would be dramatically smaller.

Overall, we show that one can decompose the causes of the increase in foreclosures in Massachusetts that started in 2007 into the effects of lower underwriting standards and the effects of falling prices, but, because the two interact, the effects are subtle. We start by comparing the 2002 and 2005 vintages of home buyers and show that the quarterly foreclosure hazard for the former never exceeded 0.05 percent in the first 14 quarters after purchase, whereas foreclosure rates for the latter group regularly exceeded 0.6 percent. We take the 2005 vintage and give them the price paths experienced by the 2002 vintage, but leave the distribution of homeownership characteristics the same. This exercise isolates the effect of underwriting because the 2005 vintage had 5 times as many subprime buyers as the 2002 vintage. According to our estimates, foreclosure rates would rise relative to 2002 levels, but still remain almost an order of magnitude below those seen in 2005. We view this as evidence that declining underwriting standards alone could not have caused the crisis; had 2005 buyers received the 2002 price experience, they would have had relatively low foreclosure rates, despite the large share of subprime mortgages. The second counterfactual exercise is to give the 2002 buyers the house price experience of the 2005 vintage of home buyers. In this case, the model generates foreclosure rates that, while are only half of what we actually observed, are still gigantic relative to historical averages, or to the 2002 vintage, and would have constituted a major public policy issue. All that said, the gap between the 2002 vintage with 2005 prices and the 2005 vintage with 2005 prices illustrates how subprime lending aggravated the effects of the price decline, essentially doubling the number of foreclosures.

To motivate our empirical model, we construct a simple model of homeownership termination. Our innovation here is to fix the problem that option theoretic models of default imply no role for household characteristics by embedding one in a portfolio choice framework. In our model, household characteristics affect the stochastic discount factors used to discount the costs and benefits of continuing to make mortgage payments. The set of characteristics includes the household’s financial situation, and thus provides a theoretically robust argument for the heuristic “double-trigger” model used by practitioners. This model posits that foreclosure occurs when a borrower faces negative equity *and* cash-flow problems. But one important point of our model is that we do not distinguish between “strategic” and “non-strategic” default because borrowers have access to unsecured credit and so a mortgage payment is always, strictly speaking, affordable. In our world, a borrower in financial distress with negative equity will choose not to default if they believe, for example, that house prices will rise rapidly and that they will eventually have positive equity. On the other hand, a borrower with high financial wealth may default in the same situation if he or she believes that house prices are going to continue to fall.

The key insight from the model is that negative equity is a necessary condition for default, because selling

dominates defaulting if a borrower has positive equity. However, negative equity is not sufficient for default, because future house price appreciation may make it profitable to continue making mortgage payments. We argue that cash-flow, or the lack thereof, then determines whether it is optimal to default – if a borrower must make extreme sacrifices in terms of current consumption to realize benefits in the distant future, it may not make sense to continue paying the mortgage, but if the sacrifice of current consumption is small, it may make sense to continue making payments.

The model predicts that if HPA is high we will observe few foreclosures, even if people have cash-flow problems. The theory also suggests that a fall in HPA will not generate foreclosures unless we also have cash-flow problems. Thus, the presence of cash flow problems is a necessary condition for negative HPA to induce foreclosures. The bottom panel of Figure 1 shows that some fraction of borrowers (at least two percent) *always* have cash-flow problems. Although cash-flow problems did appear to increase in 2006 and 2007, the scale of the increase in foreclosures during this period is an order of magnitude larger than the increase in delinquencies.

We estimate a competing risks model of homeownership termination which is similar to the mortgage termination models common in the real estate literature except that the two outcomes are foreclosure and sale rather than foreclosure and prepayment. Our results show that foreclosure is highly sensitive to unemployment and a foreclosure hazard that is essentially zero for equity levels greater than 25% but then rising sharply after that. Interestingly the sale hazard displays a similar kink at the same point suggesting that many foreclosures involve borrowers who would have sold if they had positive equity. This relationship, with reduced sales offsetting increased foreclosures, should mitigate the concern that foreclosures increase the supply of homes on the market.

One often-made argument that runs counter to our claims in this paper is that the credit expansion may not have directly caused the foreclosure boom, but did so indirectly by causing the bubble which led to the negative HPA starting in 2006. While intuitively appealing, this theory has many problems. First, while economic theory can explain why a credit expansion will lead to higher prices, no theory exists that links credit expansions robustly to bubbles. Second, there is good reason to think the causality went from rising prices to relaxed underwriting and not the other way around. Gerardi, Lehnert, Sherlund and Willen (2008) use contemporary reports by investment banks to show that optimistic beliefs about house prices motivated investment in subprime ABS. If the investment banks had known that house prices were going to fall as they did, they would never have invested in the securities. In other words, it was beliefs about house price appreciation that led to the observable decline in lending standards not the converse.

We proceed as follows. We describe a simple model of ownership termination in Section 2. We then map out our empirical strategy in Section 3, paying particular attention to our use of the ownership experience as opposed to the loan as the unit of analysis. In Section 4, we describe our dataset. Section 5 provides the basic results of the empirical model. We focus on the key identification issues in Section 6. In Section 7, we make our case that house prices are at the heart of the crisis by conducting counterfactual exercises using our estimated model. Section 8 concludes.

2. A Two-Period Model of the Default Decision

Why would a household default on a home mortgage? According to what we will call the “frictionless option model” (FOM), a borrower should default if the value of the house H falls short of the value of the mortgage V on that house. This statement is often understood to mean that the borrower should default if the value of the house falls short of the unpaid principal balance U , that is, if the homeowner has negative

equity. But, as Kau, Keenan and Kim (1994) explain,³ this interpretation is wrong: V is the present value of future payments discounted using the market stochastic discount factor. Since borrowers have the option to refinance the mortgage and to default in the future, V is always less than U , so negative equity is not a sufficient condition for default.

The FOM is useful in that it provides an explanation for why homeowners tolerate negative equity, and does well in explaining default behavior qualitatively, but not quantitatively. The problem is that the FOM assumes that the “economic environment facing a homeowner is described by two variables: the interest rate and the house value” (Kau, Keenan, and Kim, 1994). Essentially, the FOM implies that two different households who bought similar houses in the same market, at around the same time, on the same terms, will both default at the same time. Such a prediction, however, is not consistent with the data. Moreover, since individual shocks such as unemployment, divorce, and illness are quite important to the household’s default decision, the deviations from the FOM occur in a systematic way.⁴ To rectify this predictive deficiency, researchers have turned to two alternative explanations for variation in default rates among otherwise similar borrowers. The first incorporates “trigger events”—divorce, illness, and spells of unemployment are the typical examples—which make some borrowers more vulnerable to default. The second is the concept of “ruthless” default, which posits that to actually follow the advice of the FOM, a borrower must be cold-blooded and devoid of emotion. Neither of these explanations is satisfying from a theoretical standpoint, the former because it does not explain why the trigger events lead to default, and the latter because it appeals to a parameter termed “ruthlessness,” a concept which is difficult to define formally, let alone quantify.⁵

We develop a simple, two-period model to formally explain the channel by which trigger events may lead to default, and to guide us in our choice of control variables for the empirical analysis. The model retains the basic structure of the FOM but yields the intuitive prediction that financially strapped borrowers are more likely to default, conditional on a given level of house prices and interest rates. To perform this analysis, we embed the financial contracts from the FOM into a portfolio choice model with constraints. As shown by He and Pearson (1991) and others, portfolio constraints imply that households use different stochastic discount factors to value assets. Put differently, individual household valuations of identical assets typically won’t be identical. The upshot of this implication is that while the intuition that a borrower will default if the value of the mortgage exceeds the value of the house still remains true, both the value of the house and the value of the mortgage will differ across households depending on unique characteristics and the individual financial situation of the particular household. For example, possessing less financial wealth reduces the value of the house relative to the value of the mortgage, and makes default more likely.

We now establish our model and derive our basic result about the default decision. At the end of this section, we return to the FOM and show that our characterization of default, suitably adjusted, is the same as the FOM characterization of default.

We consider a two-period model, in which the household enters period 0 living in a house with a market price H_0 , and holds an interest-only mortgage with unpaid principal balance U , interest rate r_m , and with interest due in period zero. The household receives labor income y_0 , and is assumed to have wealth w . For simplicity, we assume that there is only one type of house in the model, and that the house is also available

³The FOM model of Kau, Keenan and Kim applied the contingent claims framework of Black and Scholes (1973), Cox, Ingersoll, and Ross (1985), and Epperson et al. (1995). See Kau and Keenan (1995) for a survey of the mortgage pricing literature.

⁴See Vandell (1995) for a survey of the empirical default literature, and a discussion about the importance of individual variables in the default decision.

⁵The literature does stress the potential importance of transaction costs in modeling default. However, the concept of transaction costs is somewhat vague in the literature. Often transaction costs refer to the direct costs of defaulting, in terms of default penalties and mobility costs. In other cases such costs refer more broadly to frictions in the default decision.

for rent at price ρ . In period 0, the household may choose to 1) sell the house and rent, 2) default on the mortgage and rent, or 3) continue paying the mortgage and keep the house. In addition, the household simultaneously chooses an optimal level of consumption and saving, where it can borrow $\theta_B > 0$ dollars at unsecured interest rate R_B , or save $\theta_L \geq 0$ dollars at interest rate R_L . We assume that the household cannot refinance the mortgage. In period 1, the household receives income y_1 and the market value of the house is H_1 . For simplicity, we assume that the household knows the value of the house in period 1 with certainty in period 0. This assumption does not change the main results derived from the model. If the price of the house is stochastic, then expectations of future house prices, instead of the actual present value, will matter in the household's default decision.⁶ In period 1, the household can either sell the house or default on the mortgage.⁷ After making this choice, the household consumes its net savings and income, and the model ends. Formally, we consider three different budget constraints. If the household keeps the house, it faces this set of constraints:

$$c_0 = y_0 + w - r_m U + \theta_B - \theta_L \quad (1)$$

$$c_1 = y_1 + (H_1 - U)^+ - \theta_B R_B + \theta_L R_L. \quad (2)$$

If the household elects to sell, it faces this set of constraints:

$$c_0 = y_0 + w + H_0 - U - \rho + \theta_B - \theta_L \quad (3)$$

$$c_1 = y_1 - \theta_B R_B + \theta_L R_L. \quad (4)$$

Finally, if it elects to default, the household's constraints are:

$$c_0 = y_0 + w - \rho + \theta_B - \theta_L \quad (5)$$

$$c_1 = y_1 - \theta_B R_B + \theta_L R_L. \quad (6)$$

We assume a utility function that is additively separable across housing consumption and non-housing consumption as well as across periods. We assume a log specification, in which housing enters the period 0 utility function, but does not enter the period 1 utility function. Formally, the borrower's problem is to maximize

$$\log(c_0^\alpha H^{1-\alpha}) + \beta * \log(c_1^\alpha H^{1-\alpha}) \quad (7)$$

For the case in which the household elects to keep the house in period 0, it will choose to sell in the final period if the house is worth more than the mortgage, and, if not, it will default. If instead, the household elects to default or sell the house in period 0, then we assume that it must rent a house of the same size (or value) at a price ρ (in order to obtain shelter services).

The first result here is that a borrower will default in the first period only if $H_0 < U$. In other words, negative equity is a necessary condition for default. However, by no means is negative equity a sufficient condition for default. The top panel of Figure 1, for example, shows that even in years in which house prices appreciated substantially, some borrowers did in fact default on home mortgages. High house price appreciation means that, in general, household equity goes up, but mortgage equity withdrawal, judgments,

⁶In the case of stochastic house prices, the household's default decision will also depend on the variance of the house price shock.

⁷For simplicity, we assume that there is no penalty associated with default. In reality, default may have a significant negative impact on the ability of a household to borrow in the future. If we impose a penalty in the model, the decision to default will be a function of the magnitude of the penalty.

tax delinquency, and arrears on the mortgage can all lead to default. Overall, in the years with the highest house price appreciation, we see historically low levels of default.

We now focus on the interesting special case, where $H_0 < U$, $H_1 > U$, and $r_m U > \rho$. The first condition means that the household will never sell in the first period, since the house is worth less than the mortgage. The second two conditions imply that it is costly to keep the house, but there is some future benefit to doing so in the form of future price appreciation. We define the return to keeping the house as:

$$R_K = \frac{(H_1 - U)^+}{r_m U - \rho}.$$

The interpretation of R_K is straightforward. The dividend on keeping the house is the payoff from selling the house in the next period, and the price is the cost of paying the mortgage less the rent the household would have to pay if it sold the house. Under these conditions, we advance and prove the following proposition:

Proposition 1 *If*

$$R_B > R_K > R_L, \tag{8}$$

then there exists w^ such that if $w > w^*$, the borrower keeps the house and if $w < w^*$, the borrower defaults. The implicit function $w^*(y_0, y_1, \beta, r_m, H, U)$ has the following properties: $w_{y_0}^* < 0$, $w_{y_1}^* > 0$, $w_{\beta}^* < 0$, $w_H^* < 0$, $w_U^* > 0$, and $w_{r_m}^* > 0$.*

Proof: See Appendix.

Proposition 1 shows how the incidence of foreclosure depends on household-level parameters. If we assume some distribution of households with respect to the exogenous parameters, then we can now say something about what determines the default rate. First, if we lower individual household wealth, we get more defaults. Second, anything that reduces the relative value of future consumption (higher future income, lower current income, less patience) tends to increase the likelihood of a default decision that leads to a foreclosure. Third, as one would expect, increasing the mortgage interest rate r_m makes default and thus foreclosure more likely. Finally, reductions in rental prices make holding on to the house more expensive and increase the likelihood of default.

In the model, the effect of changes in house prices on foreclosures is somewhat more subtle. Holding H_1 constant, changes in H_0 actually have no effect on foreclosures. The logic for this is simple: a reduction in the value of the house makes no difference if the household is already out of the money (that is, if $H_0 < U$). What matters, in fact, is H_1 : the key here is beliefs about future prices rather than the current market price. Essentially, holding onto the house makes sense in the model only if the household believes that somewhere down the line, the house will be worth more than the outstanding mortgage, because that is what justifies the premium paid over renting. In the end, H_0 matters only if one believes that there is a relationship between current and future housing prices. As most buyers pay close attention to recent trends in house prices, a reduction in H_0 , conditional on $H_0 < U$, will likely increase the foreclosure rate.

Above, we argued that our model was closely connected to the FOM. According to the FOM, we need to compare the value of the mortgage with the current market value of the house. In the FOM, we can establish those values by looking at market prices, but in our model it is more complicated. Take the value of the house: the house is an asset that pays ρ this period and H_1 next period when the borrower sells the house. If the household can finance the house purchase out of savings, then the value is:

$$H = \rho + \frac{H_1}{R_L}.$$

Note that the value of the house to the borrower is *not* necessarily equal to the current market price of the house. Such an outcome would be an arbitrage opportunity in a frictionless world. If $H > H_0$, for example, one could borrow H_0 dollars at interest rate R_L , buy a house, and make a certain profit of $R_L(H - H_0)$. Another household that has to finance the house purchase by borrowing would value the house at:

$$H = \rho + \frac{H_1}{R_B}.$$

Using our definitions of the value of the house and the mortgage, we can show that condition (8) is equivalent to the FOM criterion that a sufficient condition for default is that $V > H$. For a household that can finance a house with accumulated savings, the condition $R_K > R_L$ and $V > H$ are equivalent. To see why, we can use the definitions of V and H to obtain:

$$r_m U + \frac{\min(U, H_1)}{R_L} < \rho + \frac{H_1}{R_L}.$$

Re-arranging and using the fact that $H_1 - \min(U, H_1) = (H_1 - U)^+$, yields:

$$R_L < \frac{(H_1 - U)^+}{r_m U - \rho} = R_K.$$

3. Empirical Model

3.1. Ownership Experiences

Our analysis is unique in that, to the best of our knowledge, this is the first study capable of tracking the same borrowers across different mortgage instruments for the same residential property. The empirical default literature has primarily used loan-level data to simultaneously model the decision either to default or to prepay a mortgage.⁸ Instead of characterizing the prepayment and default probabilities of a single loan, in this paper we are able to characterize sale and default probabilities across the time horizons of entire “ownership experiences.” In using this term, we are referring to the time that an individual household lives in a particular house. We believe this is a significant methodological contribution to modeling default behavior for at least two reasons.

First, it is unlikely that the probability of defaulting on a subsequent mortgage is independent of the risk associated with prior mortgages purchased by the household. For example, many borrowers choose to extract equity from their homes to smooth consumption by refinancing their mortgages; this is a practice called mortgage equity withdrawal. While undoubtedly, many borrowers doing so are simply consuming the returns to their housing investment as part of their optimal consumption plan, other households may have experienced a recent adverse transitory income shock, and are extracting their housing equity as a precautionary buffer. This latter scenario suggests that, in some cases, mortgage loans that end “successfully” in the form of prepayment may actually be a signal of financial distress. In such a case, the subsequent mortgage actually has a high probability of default. Thus, it is more informative to observe all of the mortgages issued to a given borrower on a given property to calculate unbiased probabilities of default.

In addition, using loan-level data could also overestimate default probabilities for borrowers who purchase their homes with subprime mortgages, if the data do not distinguish between purchase and refinance loans. For instance, if many prime borrowers refinance into subprime mortgages as a result of financial duress,

⁸The term prepay is used in the literature to describe a situation when a borrower prematurely pays off a mortgage. Reasons for prepayment include the sale of the home, a refinance into a new loan, or early repayment of a loan.

then subprime loans will overestimate default probabilities of borrowers who purchase homes with subprime mortgages.

Thus, for these reasons, we conduct our analysis over entire ownership experiences, rather than individual loans taken out at one point in the ownership cycle. In our unique data, discussed in detail below, we find that the average number of mortgages over the life of completed ownerships is slightly less than 3, which suggests the potential importance of this differentiation.

3.2. A Competing Risks Model of Home Ownership Termination

Most of the recent literature on the determinants of mortgage default has used loan-level data to simultaneously model default and prepayment decisions in a reduced-form framework. Deng, Quigley, and Van Order (2000), Deng and Quigley (2004), and Pennington-Cross and Ho (2006) are some recent examples of studies that have analyzed the decisions to default and prepay jointly. These studies emphasize that the default option and prepayment option are linked by mutual exclusivity, since by exercising one choice, the borrower forfeits the opportunity to exercise the other one. In actuality, however, the prepayment option is exercised for one of two reasons, either because the borrower is refinancing or because he or she is selling the property and moving. However, because of data limitations, the vast majority of studies do not distinguish between these two possibilities.

We have argued above that it is problematic to use loan-level data to answer the questions that we pose in this study. Thus, we use data on homeownership experiences instead, and estimate the joint probability of sale and foreclosure using a competing risks, proportional hazard, duration model. Sale and foreclosure are competing risks in the sense that these are the only possibilities (other than death) by which an ownership experience can end, and exercising one of these choices precludes the possibility of exercising the other. The model specification that we use is taken from Meyer (1995), and is very similar to the specification of Pennington-Cross and Ho (2006) and Yu (2006). It is a competing risks, proportional hazard, duration model that allows for time-varying covariates and unobserved heterogeneity, and is based on the studies of Han and Hausman (1990), Sueyoshi (1992), and McCall (1996). As in these papers, our data are observed in discrete intervals. In principle, we could estimate the model using monthly intervals, but in order to obtain precise house price indexes, it is necessary to use quarterly intervals.⁹

The hazard function is given by the probability that homeowner i terminates the ownership experience at time t , conditional on the fact that he or she has lived in the home until time t . In our model, there are two hazard functions corresponding to the two competing actions that the homeowner can take to end the ownership experience: foreclosure and sale. Formally, the hazard function for the i th action is given by

$$\lambda_i^r(t) = \lim_{\Delta t \rightarrow 0} \frac{P(t < T_i^r < t + \Delta t | T_i^r \geq t)}{\Delta t}, \quad (9)$$

where homeowner i can terminate the ownership by selling the home, $r = S$, or defaulting on the mortgage and allowing the lender to foreclose, $r = F$.

In our sample, the first realized termination time of owner experience i is when the homeowner either sells the home or defaults on the loan, or when the end of the sample is reached. Letting T^F and T^S , denote discrete random variables representing the duration to foreclosure and sale, respectively, and T^C

⁹An important difference between our specification and the specification in Deng, Quigley, and Van Order (2000) is that in our present analysis we do not model unobserved heterogeneity. Allowing for unobserved heterogeneity using a multi-dimensional, discrete distribution, such as the one used in Deng, Quigley, Van Order (2000), with a data set that is as large as ours is extremely computationally burdensome. The time required to estimate the current specification without unobserved heterogeneity is already significant. We have experimented with allowing unobserved heterogeneity for very small samples of the data, and have not found substantially different results.

denote censoring due to the end of the sample, the first realized termination is $T_i^* = \min\{T_i^F, T_i^S, T_i^C\}$. The probability that homeowner i reaches T_i^* conditional on a vector of observed covariates, $X_i(t)$, is a function of both the foreclosure hazard, $\lambda_i^F(t|X_i(t))$, and the sale hazard, $\lambda_i^S(t|X_i(t))$, and is called the survival function:

$$V(T_i^*|X_i(t)) = \exp \left[- \int_{t=1}^{T_i^*} (\lambda_i^S(t|X_i(t)) + \lambda_i^F(t|X_i(t))) dt \right]. \quad (10)$$

where

$$\lambda^r(t|X_i(t)) = \lambda_0^r(t) \exp(X_i(t)' * \beta^r), \quad r \in \{S, F\}. \quad (11)$$

This specification for the hazards, which is common in the literature, implicitly assumes that the foreclosure hazard $\lambda_i^F(t|X_i(t))$ and the sale hazard $\lambda_i^S(t|X_i(t))$ each take the same form as a function of the duration of the ownership period for all N ownership experiences in our sample. In other words, the analysis rests on the assumption that there is a hazard common to all homeownership experiences, $i \in N$ for default $\lambda_0^F(t)$, and for sale $\lambda_0^S(t)$. These hazards are referred to as the baseline hazards for foreclosure and sale, respectively. Our analysis estimates the effect of the covariates on the joint decision to sell and default as a proportion of the two respective baseline hazards. We specify the baseline hazards in a non-parametric manner, which basically entails estimating a set of dichotomous variables (one for each discrete quartly interval).¹⁰

Since we do not observe the data continuously, but only in discrete intervals, we must account for this when forming the likelihood function. In what follows, we use the methods employed in Meyer (1995). Suppose we only observe T_i^* in the interval I_t , where $I_t = [t, t+1)$ for $t = 0, 1, \dots, T^C - 1$ and $I_{T^C} = [T^C, \infty)$. If $T_i^* \in I_t$, then we let $k_i = t$. We follow Sueyoshi (1992) and assume that any time-varying covariates are constant within the intervals I . This implies that the data available to the econometrician for the i th homeowner include $(k_i, \delta_i, X_i(t))$ where k_i is equal to the integer part of T_i^* , $X_i(t)$ denotes the path of the explanatory variables, and $\delta_i^r = 1$ when individual i is observed to fail due to the r th risk (foreclosure or sale), and $\delta_i^r = 0$ otherwise.

Under these assumptions, the probability of homeowner i surviving the r th risk in the interval $(t, t+1)$ is

$$\begin{aligned} P[T_i^{(r)} \geq t+1 | T_i^{(r)} \geq t] &= \exp \left[- \int_t^{t+1} \lambda_i^r(s|X_i(s)) ds \right] \\ &= \exp \{ - \exp[\gamma^r(t) + X_i(t)' \beta^r] \}, \end{aligned} \quad (12)$$

where $\gamma^r(t) = \ln \{ \int_t^{t+1} \lambda_0^r(s) ds \}$.

This implies that the likelihood of the data is given by

$$L(\gamma, \beta) = \prod_{i=1}^N ([1 - \exp\{-\exp[\gamma^F(k_i) + X(k_i)' \beta^F]\}]^{\delta_i^F} * [1 - \exp\{-\exp[\gamma^S(k_i) + X(k_i)' \beta^S]\}]^{\delta_i^S} * H(k_i|X_i)) \quad (13)$$

where

$$H(k_i|X_i) = \prod_{t=0}^{k_i-1} \exp \{ -(\exp[\gamma^F(t) + X_i(t)' \beta^F] + \exp[\gamma^S(t) + X_i(t)' \beta^S]) \}. \quad (14)$$

¹⁰We also estimate a specification of the model in which we restrict the shape of the baseline hazard to a third-degree polynomial in the age of the ownership, such that $\lambda_0^r(t) = \exp(\alpha_0 + \alpha_{1,i}t + \alpha_{2,i}t^2 + \alpha_{3,i}t^3)$. The results do not materially change.

4. Dataset

Our data come from the The Warren Group, and are comprised of historical registry of deeds records from January 1987 through December 2008 for the entire state of Massachusetts, as well as 2007 and 2008 Massachusetts assessor data. The registry of deeds records contain information on all residential home sales and mortgage originations, allowing us to track every mortgage issued on a given residential property over our sample period,¹¹ while the assessor data contain information regarding characteristics of the property. In the data, we see transaction amounts and dates for mortgages and property sales, but we do not have information on mortgage characteristics such as the type of mortgage or the contracted interest rate. The data do contain information about the identity of the mortgage lender that we use to construct indicators for homes purchased with subprime mortgages, which we discuss in detail below. Information regarding the type of sale is also found in the data. Thus, we can distinguish between normal sale transactions, nominal sale transactions, such as transfers among family members, and foreclosure sales.

In this paper we use foreclosure sales as a proxy for default. Foreclosure sales in the data are final transactions, in the sense that they signify the eviction of the household or homeowner from the property. The data also cover foreclosure petitions going back to 2004. Foreclosure petitions are public notices declaring the initiation of foreclosure proceedings that by law mortgage lenders are required to make in Massachusetts. Petitions are usually filed once the borrower has become delinquent on three monthly mortgage payments. While a foreclosure petition is certainly a sign of serious delinquency, it is not a good indication of default, as there are many instances in the data in which a borrower receives a foreclosure petition, but does not end up in foreclosure proceedings. The typical time between the filing of a petition and a foreclosure deed (if it occurs) is about 2 months. Thus, a borrower who receives a petition has some time to decide whether or not to catch up on the late mortgage payments and forestall default. For this reason, we believe that the appropriate proxies for default in our data are foreclosure sales.

The fact that we see every sale transaction in the data allows us to construct the ownership experiences defined and discussed above, by identifying sales of the same property in the data. Figure 2 contains an example of how we construct ownership experiences.

Between January 1987 and December 2008 we observe more than 6 million mortgages and almost 1.6 million ownership experiences in Massachusetts. Table 1 lists the number of sales and foreclosures by year as well as by vintage, beginning in 1989.¹² Our data encompasses the housing bust of the early 1990s in the Northeast, and this event is apparent from increased foreclosure numbers. There is a large increase in foreclosures beginning in 1991 and peaking in 1992, with over 9,000 foreclosures statewide. Data from the housing boom that took place in the early 2000s display the opposite extreme. In this period foreclosures dropped significantly, reaching a low point of fewer than 600 in 2003. We see evidence of the current foreclosure crisis at the very end of our sample. Foreclosures in 2008 surpass the levels witnessed in the early 1990s.

The top panel of Figure 1 displays a time series plot of the foreclosure rate.¹³ The top panel of Figure 1 also contains a graph of annualized house price growth over our sample period.¹⁴ This figure clearly shows

¹¹Residential properties include condominiums, single-family homes, and multi-family homes.

¹²The data on sales are fairly reliable before 1989; however, the number of foreclosures in the data before 1989 seems implausibly low, so we restrict our analysis of foreclosures to 1989 and later.

¹³Here we have simply divided the foreclosure numbers from Table 1 by the number of residential parcels in that year. A parcel is defined as a real unit of property used for the assessment of property taxes, and a typical parcel consists of a plot of land defined by a deed and any buildings located on the land. Information on parcel counts is obtained from the Massachusetts Department of Revenue.

¹⁴Our detailed data of repeat sales allow us to use the Case-Shiller weighted-repeat-sales (WRS) methodology to calculate house price indexes at various levels of aggregation. Figure 1 displays house prices at the state level, but for the purposes

the two housing market cycles that the Massachusetts economy has experienced over the past 20 years. The two series are negatively correlated, with the foreclosure rates moving inversely, at a slight lag with house price growth. In 1991, house price growth reached a low of almost -9 percent, and the foreclosure rate peaked one year later in 1992 at more than 0.6 percent of active ownership experiences. The latest figures for 2008 show the state’s house price growth and the foreclosure rate approaching the levels of the early 1990s, as growth is approaching -5 percent, while the foreclosure rate is approaching 0.45 percent.

4.1. Explanatory Variables

While the theoretical model discussed above is too simple to yield formal, testable predictions, it is useful in identifying the types of variables that should be expected to impact an individual borrower’s decision to default. It tells us, for example, that a borrower’s current wealth, income, and housing equity should all be important determinants of default risk. The wealth threshold derived from the model is a function of both current and future income. This threshold also suggests that income variability, although formally left out of our model, will also influence the decision to default.¹⁵ Unfortunately our current dataset does not contain borrower-specific demographic or financial information, nor does it contain detailed information regarding mortgage characteristics. We attempt to proxy for borrower-specific variables with more aggregated information at either the town level or the zip code level. For example, we use information on income at the zip code level and unemployment at the town level to proxy for differences in financial conditions of households across locations, and we use the racial makeup of zip codes to proxy for differences in demographics. In order to proxy for the credit risk of a home buyer, we distinguish between borrowers who purchased a home with a subprime mortgage versus those who did so with a prime mortgage. We provide a detailed discussion of our definition of subprime, and how it is related to other definitions in the literature below.

4.1.1. Initial Owner Equity and Cumulative House Price Appreciation

Previous studies in the literature, such as Deng, Quigley, and Van Order (2000), use house price indexes at the MSA-level and the unpaid mortgage balance computed from the contract terms of each loan to estimate the probability of negative homeowner equity, which proxies for whether or not the default option is “in the money.” While our dataset does not include the necessary contract information to calculate the remaining mortgage balance each period, we argue that even if it did, including such information in our estimation would introduce endogeneity issues. Since we perform our estimation over entire ownership experiences, to compute net homeowner equity in each period would require including mortgage balance information for every loan obtained over the duration of each ownership. The choice to refinance, as well as the choice of the refinanced mortgage terms, are endogenous decisions. Many borrowers, for example, who suffer negative income shocks or other adverse life events may choose to extract equity and obtain larger subsequent mortgages. In these cases the mortgage balances would simply be conveying information about negative income shocks or other adverse shocks affecting the individual borrower. Thus, including the time profile of net equity in our estimation would result in biased estimates.¹⁶ Instead, we choose to include initial homeowner equity, and the true exogenous component of net equity over time, house price appreciation.¹⁷

of estimation, we are able to calculate house price indexes at the city level for approximately two-thirds of the cities in Massachusetts. A detailed discussion of the WRS methodology and our aggregation assumptions can be found in Appendix A.

¹⁵Income variability has been identified in the literature as having an important impact on the probability of default. Herzog and Early (1970) found that borrowers in occupations with greater income volatility were more likely to be delinquent than other borrowers.

¹⁶In addition, while the Warren dataset does include information on every mortgage obtained over the life of the homeownership, it does not include the date of mortgage terminations, it is not possible to obtain a precise estimate of a borrower’s overall mortgage balance over time.

¹⁷See section 6 below for a detailed discussion of the exogeneity of our measure of house price appreciation.

Our dataset has enough information to construct reasonable proxies for initial net housing equity and cumulative house price appreciation. Since we see dates and amounts in our data, we are able to calculate initial loan-to-value ratios (LTV) for each ownership experience in which we observe a purchase transaction (meaning the ownership experiences that are not left-censored). Another novel aspect of the dataset, is that unlike virtually all loan-level datasets, it allows for a very precise measurement of cumulative initial LTV at origination.¹⁸ The LTVs are cumulative in the sense that we observe all of the originating mortgages issued at the time of purchase, so if there is a second or even a third mortgage that accompanies the first mortgage, we will add these amounts to our LTV calculation.¹⁹ In addition to initial cumulative LTVs, we use repeat-sale price indexes to calculate from the date of purchase the average cumulative price appreciation at each quarter in the town containing each property in our data, to construct a measure of nominal equity, E_{it} , given by

$$E_{it} = \frac{V_{i0} \cdot (1 + C_{jt}^{HPPA}) - L_{i0}}{L_{i0}}, \quad (15)$$

In constructing this variable we use only information conveyed by the initial LTV ratio, $\frac{L_{i0}}{V_{i0}}$ and the cumulative amount of house price appreciation experienced (at the town-level), C_{jt}^{HPPA} .

Theoretical models of default predict a non-linear relationship between equity and foreclosure. Specifically, an increase in equity for a borrower in a position of negative equity should have a significantly different effect from an increase in equity for a borrower who has positive equity in his or her home. For this reason, we allow for the effect of equity on default to change depending on the equity level of the borrower. To do this we specify equity as a linear spline in the model, with five intervals: $(-\infty, -10\%)$, $[-10\%, 0\%)$, $[0\%, 10\%)$, $[10\%, 25\%)$, and $[25\%, \infty)$.²⁰

Table 1 displays summary statistics for initial LTVs by year. The table contains means and medians for all ownership experiences initiated in a given year, as well as for ownership experiences that ended in a foreclosure. The differences are substantial. For example, for ownership experiences beginning in 2003, the average initial LTV ratio was 0.81, and the median was 0.85. In contrast, for the ownership experiences that began in 2003 and subsequently defaulted, the average initial LTV ratio was 0.93, and the median was 0.95.

4.1.2. Investors

In addition to the financial characteristics associated with an individual borrower, the reason for purchasing the residential property should also impact the default decision. Borrowers who are not owner-occupants, but who purchased the property strictly for investment purposes, are likely greater default risks, *ceteris paribus*. Since non-owner occupants do not face mobility costs and do not have an emotional stake in the property, their cost of default is likely lower relative to the cost to owner-occupants. This is also apparent from the model, as an investor who defaults would not have to pay ρ to rent a home. We do not have direct information in the data to distinguish properties bought to serve as primary residences or investments; however, we can proxy for these different purposes using property characteristics. The Warren Group data contains a limited amount of information on the characteristics of each property, allowing us to distinguish between single-family homes, multi-family homes, and condominiums. We hypothesize that owners of multi-family properties are more likely to be investors, as multi-family units provide a stream of rental income.

¹⁸For example the popular LoanPerformance dataset has only partial information on the existence of second or third mortgages, while data from Loan Processing Services, another common source of loan-level data in the literature, has no information on additional mortgages.

¹⁹Many studies have found evidence that initial LTV ratios alone are significant predictors of default. Early examples in the literature include von Furstenberg (1969) and Campbell, Dietrich, (1983). These studies interpret LTV ratios as acting as a proxy for other, unobservable borrower-specific information, such as wealth or income. This is also our interpretation.

²⁰The intervals are chosen somewhat arbitrarily, but the results are not significantly affected by assuming different intervals.

We also believe that condominiums may proxy in part for investors and real estate speculators, since the condominium market is often hit hardest during a housing bust.

4.1.3. Labor Market Conditions

Besides including initial LTVs, cumulative price appreciation, and property characteristics, we obtain information regarding certain characteristics of the town or zip code where each property is located. Since we know the exact location of each property, we can group properties by town/city as well as by zip code, allowing us to merge our data with data from the Census Bureau and the Bureau of Labor Statistics (BLS). We obtain town-level data on monthly unemployment rates going back to 1990 from the BLS. Unemployment rates have been used previously in the literature as a proxy for income volatility.²¹ In addition, for some households, periods of unemployment turn out to be permanent income shocks, which our portfolio-choice model tells us will affect the default decision. Instead of including the level of unemployment in our empirical model, we choose to use the cumulative change in unemployment experienced in the town of residence since the time of home purchase. This choice is based on our belief that flows into and out of unemployment are the relevant factors that affect foreclosure. We choose to include a cumulative measure rather than quarterly growth rates in order to smooth out some of the noise that is present in the quarterly unemployment measures at the town-level.²²

4.1.4. Demographics

We obtain demographic information at the zip code level from the 2000 U.S. Census, including median household income and the percentage of minority households. Previous studies in the literature have found evidence of racial discrimination in the mortgage lending business. These studies show that minority households are less likely than others to obtain approval for a mortgage, and those that are successful in obtaining a mortgage are often given a loan with inferior terms relative to an equivalent non-minority household.²³ Thus, using the intuition from our simple model, we would expect to observe higher default rates in zip codes with larger percentages of minority households, *ceteris paribus*.

4.1.5. Interest Rates

The mortgage default literature has also identified the difference between the par value and the market value of a loan as a significant determinant of default risk. When the value of the mortgage rises above the outstanding mortgage balance, the probability of default increases.²⁴ The difference between par and market value can be negative only when the prevailing market interest rate that the borrower is qualified to receive falls below the original contract rate of the mortgage. Since we lack information about contracted interest rates, we proxy for this effect using prevailing aggregate interest rates. In the estimation, we include the 6-month LIBOR rate, which is a short-term interest rate that has become a very popular index for adjustable-rate mortgages, especially in the subprime mortgage market.²⁵

4.1.6. Subprime Mortgages

There is no universally accepted definition of a subprime mortgage, so in order to analyze the subprime mortgage market, we must first choose an appropriate definition. The terms subprime borrower, subprime

²¹Williams, Beranek, and Kenkel (1974) found that unemployment rates in Pittsburgh had a positive effect on default. Campbell, Dietrich, (1983) and Deng, Quigley, and van Order (2000) also found evidence of a positive effect of the unemployment rate on default.

²²The estimation results are not sensitive to the additional inclusion of unemployment levels at the town-level.

²³See Munnell et al. (1996).

²⁴See Quercia and Stegman (1992) for a discussion of this effect.

²⁵We have also estimated the model with other aggregate interest rates, such as the 10-year T-bill rate to capture changes in the market values of fixed-rate mortgages, but the results are largely unchanged.

lender, and subprime mortgage, are often used interchangeably by analysts and researchers, but for the purposes of this paper, it is important to make a distinction.

The term subprime borrower has traditionally been applied to a borrower that is perceived to be a higher lending risk relative to the average borrower, usually because of a poor credit history. In the United States, a subprime borrower today typically refers to an individual with a FICO score below 620, who has become delinquent on some form of debt repayment in the previous 12 to 24 months, or who has even filed for bankruptcy in the last few years. There have always been small-scale venues for subprime borrowers; for instance, pawn shops and payday lenders have both existed for a long time. However, until recently, subprime borrowers were unable to systematically obtain large-scale loans such as mortgages. With reformed lending laws and increasingly sophisticated financial markets and instruments, a new mortgage lending channel that was based on the concept of risk-based pricing emerged that serviced subprime borrowers in particular. It is unclear exactly when the subprime mortgage market truly began, but in 1993 the Department of Housing and Urban Development (HUD) began tracking the subprime mortgage market, and developed an annual list of subprime lenders, which were defined as mortgage lenders that specialized in lending to subprime borrowers.²⁶ This list is calculated each year by identifying mortgage lenders that originate a large percentage of “high-cost” loans.²⁷ The Warren dataset has very accurate information about the identity of each mortgage lender. Thus, we are able to use the HUD list of subprime mortgage lenders (by year) to identify subprime mortgages in the data. While this list is not a perfect representation of the subprime mortgage market, since some of the lenders identified as subprime also originate a small percentage of lower cost loans, it does provide a very good representation of subprime lenders.²⁸ In Appendix B we take a closer look at this issue and perform a robustness check, where we use the limited interest rate information that is available in our data for adjustable-rate mortgages, to construct a definition of a subprime mortgage based on interest rate spreads, similar to the methods used by HMDA analysts. We find that the mortgages identified as subprime using the HUD list correspond very well to mortgages that we define as subprime based on the spread between initial contract interest rates and a common benchmark index of prime mortgage rates. Table 2 displays a list of the top ten subprime lenders in our data, in terms of the number of subprime purchase mortgages originated. These ten lenders accounted for almost 70 percent of the total number of subprime purchase mortgages that were originated in Massachusetts. As of October 2008, all ten lenders have shut down their operations, which confirms in our data what has been obvious in reality: the fact that the current foreclosure crisis has taken the hardest toll on subprime borrowers and the lenders that serviced them.

At this point, it is important to be very clear in our interpretation of this variable. In our empirical model we interpret the estimates corresponding to the subprime indicator as correlations, and do not assign a causal interpretation. It simply tells us the relative frequency of foreclosure for a borrower who used the subprime lending channel compared with the frequency for a borrower who used the prime lending channel. This variable likely includes omitted selection effects, as borrowers who use the subprime market have poor credit histories relative to other borrowers. It also likely measures omitted treatment effects, since interest rates in the subprime market are often significantly higher than rates in the prime market, which results in higher payments, which, in turn, increases the likelihood of default.

Table 3 displays the size of the Massachusetts subprime market, beginning in 1993, in terms of the

²⁶It is important to stress that most, but not all, of a subprime mortgage lender’s business involves servicing subprime borrowers. Yet many subprime lenders also service prime borrowers, meaning borrowers with a sound financial credit history.

²⁷HUD defines a “high-cost” loan to be a mortgage with an initial interest rate that is at least 300 basis points larger than the yield of a treasury bill with a comparable maturity period. The reasoning behind this method is that lenders charge higher interest rates to subprime borrowers to compensate for the elevated credit risk.

²⁸It is the same method used by many researchers in the industry, including, for example, the Mortgage Banker’s Association.

percentage of purchases financed with subprime mortgages each year (column 1), as well as the subprime market’s contribution to the total number of foreclosures in Massachusetts. In terms of purchase mortgages, the peak of subprime lending occurred in the 2004–2006 period, when between 10 and 15 percent of the state’s purchase mortgage originations were made by subprime lenders. The last three columns in Table 3 show the disproportionate impact that the subprime mortgage market has had on the recent rise in foreclosures. In 2006 and 2007, approximately 30 percent of Massachusetts foreclosures were on ownerships initially financed with subprime mortgages, up from only 10 percent in 2003 and 2004. While these percentages are high, we see an increase when looking at the percentage of recent foreclosures on borrowers whose last mortgage was issued by a subprime lender. In 2006 and 2007, this percentage was approximately 45 percent. If we broaden the criteria even further and consider the contribution from borrowers who obtained a subprime mortgage at any point in their ownership experience, the foreclosure percentage increases to approximately 59 percent.

This distinction emphasizes the importance of differentiating between different segments of the subprime market, and specifically between purchase loans and refinance loans. For ownership experiences that begin with mortgages obtained from a prime lender, subprime refinances are often a signal of financial distress, especially for borrowers that extract equity with a subprime refinance. It is likely that in the absence of a subprime market, many of those borrowers that ended up defaulting would have defaulted on their previous prime mortgages.²⁹ In our empirical analysis we do not consider this group of borrowers, but instead focus on ownerships that initially financed their home purchase with a subprime mortgage.

4.2. Comparison of Massachusetts to the National Mortgage Market

As we discussed above, a big advantage of the Warren dataset comes from its coverage of the entire mortgage market. It includes all homeownerships financed with a mortgage, and thus covers the subprime, alt-a, and jumbo non-agency market segments as well as the prime conforming segment. However, a disadvantage is that it only covers a single state. While our analysis is only performed with data from Massachusetts, we believe that many of the results can be generalized to the national mortgage market. In this section we will present some basic facts regarding the time series of subprime shares, delinquency rates, and foreclosure rates in Massachusetts, in comparison to the nation as a whole, as well as to California, which is the state with the largest subprime mortgage market. The purpose of this presentation is to give a sense of just how representative the Massachusetts prime and subprime mortgage markets are of the national market.

Figure 3 contains three plots that summarize the state of the MA, CA, and U.S. mortgage markets over the 2004 – 2008 period. The first plot displays the share of outstanding mortgages that are prime in each market, the second plot displays the share of past-due mortgages that are prime in each market, and the third plot displays the share of foreclosure starts that have occurred on prime mortgages. The data is taken from the Mortgage Bankers Association (MBA), and to obtain the corresponding subprime shares from the figures, one only has to subtract the prime shares from 100. The figure shows a very high degree of similarity between Massachusetts, California, and the entire U.S. in terms of the patterns of the share of outstanding mortgages that are subprime (constant at about 10 percent over the period) and the share of past-due mortgages that are subprime. In addition, the share of foreclosure starts that have occurred on subprime mortgages in Massachusetts and California is very similar over the four-year period (the pattern for the entire U.S. is much more volatile).

²⁹This point is especially important in the context of the current public policy debate regarding whether or not subprime borrowers should even be allowed to obtain mortgages and purchase homes, since borrowers that refinanced from a prime to a subprime lender should not be included in the discussion.

4.3. The Decision to Sell

While our main focus in this paper is specifically modeling the default decision that leads to foreclosure, the competing risks nature of our estimation means that we must also address a household’s decision to sell. The mobility literature emphasizes the importance of life-cycle factors and factors that affect housing demand in the decision to sell and change residence.³⁰ Variables such as household size, age, marital status, income, wealth, and education have been found to impact a household’s decision to move. Unfortunately, our data do not contain such information at the household level. However, theory tells us that equity should also have an important, non-linear impact on the decision to sell. If a homeowner in a position of negative equity desires to move, he must come up with the difference between the sale amount and the remaining nominal mortgage balance in order to meet his mortgage obligation. For this reason it is financially beneficial for negative-equity homeowners to instead exercise their default option, and give their home to the mortgage lender in exchange for the outstanding balance of their mortgage. In contrast, for homeowners with positive equity, it is always better to sell and extract the positive equity, rather than to exercise the default option. Thus, we would expect different effects of equity on the probability of sale for homeowners in positions of negative equity compared to homeowners with positive equity.³¹ For this reason, we use the same equity specification (discussed above) in estimating the sale hazard as we use to estimate the foreclosure hazard.

5. Results

5.1. Non-parametric hazards

To gain insight about foreclosure and sale probabilities from our data, we calculate the non-parametric Kaplan-Meier estimates by length of the ownership experience.³² We use a quarterly interval length in our subsequent analysis, as this is the finest partition with which we are able to estimate the house price indexes.

Figure 4 contains the Kaplan-Meier hazard functions for foreclosure and sale, respectively, using our entire sample of data.³³ The foreclosure hazard exhibits positive duration dependence in the first 4 years of the ownership experience, peaking at a quarterly rate of almost 0.15 percent, and then negative duration dependence for the remainder of the ownership period. This hump-shaped pattern is consistent with findings from the mortgage default literature.³⁴ The sale hazard has a similar pattern as it also exhibits positive duration dependence early in the ownership experience (approximately 3 years), and negative duration dependence later, although it does not decline as much as the foreclosure hazard.

³⁰See, for example, Henderson and Ioannides (1989), Henly (1998), and Chan (1996).

³¹A study of U.K. homeowners (Henly, 1998), finds that households with negative equity have less mobility than those with positive equity. This result suggests that house price appreciation may have a positive impact on the decision to sell. Henly also finds weak evidence of a negative impact on mobility from regional unemployment rates.

³²The Kaplan-Meier estimates are calculated as follows: We assume that hazards occur at discrete times t_j where $t_j = t_0 + j$, $j = 1, 2, \dots, J$. If we define the number of loans that have reached time t_j without being terminated or censored as n_j , and the number of terminations due to risk k at t_j as d_{kj} , then the Kaplan-Meier estimates of the hazard and survival function is

$$\lambda_k(t_j) = \frac{d_{kj}}{n_j}, S(t_j) = \prod_{t_j \leq t} \left(1 - \frac{d_j}{n_j}\right).$$

The cumulative incidence function for cause k is

$$I_k(t_j) = \sum_{i=1}^j \lambda_k(t_i) S(t_i).$$

³³In estimating the foreclosure (sale) hazard, we treat sales (foreclosures) as “censored” data, based on the implicit assumption that they are independent of each other. However, in our multi-variate analysis below, we relax this assumption.

³⁴For example, von Furstenberg (1969) found that mortgage risk increased with the age of the mortgage, up to 3 or 4 years after origination, after which it declined.

Figure 5 displays Kaplan-Meier conditional hazard rates for foreclosure, broken down categorically for some of the explanatory variables discussed in Section 4.1. This exercise should provide some initial insight as to whether these variables have any potential to explain variations in foreclosure rates. In the upper left panel of Figure 5 we see that hazard rates are much higher for ownership experiences that begin with extremely high combined LTV ratios (1 or greater). This effect seems particularly strong early on in the ownership period. For example, conditional on surviving for two years, an ownership experience that begins with an initial LTV ratio of one or greater is more than six times more likely to default relative to an experience that begins with an LTV between 0.9 and 1. Initial combined LTV ratios between 0.9 and 1.0 have slightly lower hazard rates than those between 0.8 and 0.9, and initial LTV ratios of 0.8 or lower have the lowest hazard rates.

To construct the upper right panel of Figure 5, we use our house price indices to compute homeowner’s equity after purchase. The differences between the top two panels, which differ in that the left one divides the sample by equity at origination and the right one divides the sample by current equity, are instructive. Origination LTV has always presented a problem for analysis of mortgage performance. Borrowers with low equity may be more likely to default because they have less money to begin with, what we would call a selection effect, or because exercising the default option is more profitable, a treatment effect. The figures show that the effect of origination LTV, which has more to do with selection is initially strong but wears off, whereas the effect of updated equity, which has more to do with treatment, remains strong indefinitely.

The lower left panel of Figure 5 displays differences in hazard rates for various unemployment rates. From the plot, it appears that borrowers living in areas with very high unemployment rates (greater than 7 percent) are much more likely to default.

The lower right panel of Figure 5 displays the differences in hazard rates for borrowers that financed their house purchase using a mortgage from a subprime lender versus those that obtained a mortgage from a prime lender. The difference is substantial. The hazard rate for borrowers using prime mortgages never reaches 0.08 percent, while rates peak at approximately 0.7 percent after two years for borrowers using subprime mortgages. These non-parametric hazard rates imply that the cumulative survival (success) rate after 12 years for borrowers who finance their home purchase with a mortgage from a subprime lender is approximately 87 percent, while the rate for those using a prime lender is more than 98 percent.

5.2. Competing Risks Model

In order to ease the computational burden of the maximum likelihood estimation, we were forced to make a few assumptions. Our full data set includes all Massachusetts ownership experiences financed by a mortgage over the past 20 years. This accounts for approximately 1.6 million ownership experiences. After removing observations for which any of our explanatory variables were missing, this translates into a data set of approximately 1 million ownership experiences. We chose to truncate ownership experiences at 32 quarters (8 years), since we do not see many ownerships endure longer than 8 years in our data, and those we do see mostly come from the early period of our dataset.³⁵ This assumption proved to significantly speed up convergence in our maximum likelihood routine. To further facilitate estimation, we chose to use a 10 percent random sample of our data.³⁶

³⁵The estimation results are robust to changing this cutoff to 48 quarters (12 years)

³⁶In order to estimate the model, we must put the data into “long-form,” which entails expanding the data set to include observations for each quarter that each ownership experience is active. For example, if we have an ownership experience that began in the first quarter of 2000 and ended in the first quarter of 2007, we would need to create 29 quarterly observations. Thus, our data set of 1 million ownerships becomes almost 20 million observations, and estimation becomes computationally burdensome.

Table 4 displays the parameter estimates. The first column displays estimates corresponding to the foreclosure parameters, while the second column displays estimates corresponding to the sale parameters. The direction of the signs of the coefficient estimates for foreclosure are consistent with our expectations, and with evidence from prior studies. The estimated effects of short-term interest rates (LIBOR) and the unemployment rate on the incidence of foreclosure are positive. Borrowers who finance their house purchase with a mortgage from a subprime lender are more likely to experience a foreclosure than those who used a prime lender. Borrowers who purchase a condominium or a multi-family property are more likely to default than borrowers who purchased a single-family home.

While the signs of the coefficient estimates seem very reasonable, we must also verify that the magnitudes are sensible. The quantitative implications of the parameter estimates are displayed in Table 4. The top panel of this table displays the “standardized elasticities,” which correspond to the percentage changes in the conditional hazard due to a one-standard deviation change in each continuous explanatory variable, and a change from zero to one for each dichotomous explanatory variable. The bottom panel of Table 4 displays the effects of these changes on the level of the conditional default hazard for ownerships that have aged 5 years. For example, the panel shows that a homeowner who purchased with a subprime mortgage is 9 times more likely to default, all else equal, than someone who purchased with a prime mortgage and 1.1 times more likely to default if the unemployment rate is one standard deviation above average. Since this is a proportional hazard model, the effect of different changes affect the hazard multiplicatively, so the combined effect of subprime ownership and high unemployment is $9 \times 1.1 = 9.9$, not $9 + 1.1 = 10.1$.

The second column in Table 4 reports the estimates for the sale hazard when we parameterize the baseline hazard as a third-order polynomial in the length of the ownership, while the second column in Table 4 reports the percentage changes in the conditional hazard due to a one-standard deviation change in each continuous explanatory variable, and a change from zero to one for each dichotomous explanatory variable. Higher initial LTV ratios are estimated to increase the probability of sale. Subprime ownerships are approximately 45 percent more likely to sell at any point in the ownership than prime ownerships. Condominiums are also estimated to be more likely to be sold compared with single- and multi-family homes. This result is not surprising to the extent that condominiums tend to be occupied by households that are in the early or late stages of their life-cycle. The other notable result for sales is the significant, negative effect from an increase in unemployment. A one-standard deviation increase in the unemployment rate is estimated to decrease the probability of sale by approximately 12 percent.

The estimated effect of cumulative house price appreciation through our proxy of housing equity on foreclosure is, of course, at the heart of this paper. Because housing equity E_{it} is estimated with a linear spline, we do not display the estimated parameters for that variable in Table 4. Rather, we graph the predicted foreclosure and sale hazard as a function of equity in Figure 6. The other covariates have been set to their sample averages. Overall, the figure shows that increases in E_{it} have a large and negative effect on foreclosures, and a large and positive effect on sales, for the range of equity values between -100 and 25% of the purchase mortgage. For homeowners with equity values above 25% , further increases in equity have a much smaller effect on the foreclosure and sale hazards. This is not surprising. As discussed above, homeowners with positive equity are unlikely to default, because they are better off selling their homes instead. Thus, if a homeowner already has positive equity, additional positive equity is likely to matter little in the default decision. However, when one takes into account various transactions costs involved in selling a home, such as real estate broker commissions (often 6% of the sale price) and moving expenses, the equity threshold at which borrowers will default may be slightly greater than zero. Thus, the apparent kink in the predicted default hazard at 25% equity is not necessarily inconsistent with these theoretical arguments.

6. Identification

The FOM, as well as our theoretical model, concludes that negative equity is a necessary (but not a sufficient) condition for default. This finding, combined with the assumption that initial mortgage balances are not larger than house values, means that negative house price appreciation is a necessary condition for default. The estimation results from our empirical model confirm that house price appreciation has a strong, negative effect on the probability of foreclosure. That is, a decrease in equity resulting from cumulative house price depreciation from the time of purchase significantly increases the probability of foreclosure for a given ownership. One potential concern with this empirical finding is the direction of causality. The empirical duration model assumes that the causality runs from house prices to foreclosures, and thus, that house prices are exogenous in the model. However, if the direction of causality is the opposite, then the estimation would suffer from an identification problem, and our interpretation of the estimated parameters would be incorrect. For example, a large concentration of foreclosures in a given neighborhood could adversely affect the values of surrounding homes in that neighborhood. It is not unreasonable to argue that this may be an issue, since some evidence of negative, social externalities of residential foreclosure has been found in the literature.³⁷ Another related issue is the concern that house price movements could reflect some other macroeconomic influence that is the true cause of foreclosures, but which is left out of the empirical model. Perhaps the variable that best fits this description is unemployment. The intuition is that a negative unemployment shock in a given location could directly lead to cash-flow problems for households, thus causing a rise in foreclosures. At the same time, high levels of unemployment could result in migration out of the area, which could in turn result in a decrease in house prices. In both of these cases our causal interpretation would not be valid. Finally, a very popular explanation of the foreclosure crisis, both in the literature, as well as in the media, is that a relaxation of underwriting standards characterized by non-traditional mortgage products marketed to financially unsophisticated borrowers resulted in the foreclosure boom, which also resulted in a house price crash.³⁸ In this section we address these issues of identification with a battery of robustness checks.

We begin the discussion by making the observation that the variation over time in our equity measure comes exclusively from variation over time in *cumulative* house price appreciation since the time of house purchase for a given borrower. The variation in this variable has three sources: house price variation over time; house price variation across geographic locations (towns/cities); and finally, variation across homeownerships in the same town, resulting from different purchase times. In other words, there is significant household-level variation in this variable across ownerships within the same town, as a result of different purchase times, and thus, different housing market conditions at the time of purchase. While the time-series variation and the variation across geographic locations in equity may both suffer from reverse causality or simultaneity issues, the household-level variation does not. The intuition for this assertion is simple. If we look at two households in the same town in the same time period, who purchased their homes five years apart, for example, the difference in their equity positions comes from two sources: differences in initial downpayments (i.e. initial LTV ratios, which we control for separately in the estimation), and differences in purchase times that result in different house-price experiences. Since both households are prone to the exact same contemporaneous movements in house prices (recall that we measure house prices at the town level), this source of variation does not contribute to differences in equity positions between the households (i.e. it

³⁷See Immergluck and Smith (2006), Dubin (2008), Rogers and Winters (2008), and Leonard and Murdoch (2008) for some examples from the literature.

³⁸See Pavlov and Wachter (2006), Coleman IV, Lacour-Little, and Vandell (2008), Wheaton and Lee (2008), Wheaton and Nechayev (2008) and Sanders, Chomsisengphet, Agarwal, and Ambrose (2008).

acts like a fixed effect that is common to all households in the same town in the same time period).³⁹

In fact, we find that the majority of the variation in our equity measure (approximately 70 percent) comes from this household component. While this should substantially mitigate any potential bias induced by a negative effect of foreclosures on house price growth, the time-series and geographic variation could still introduce some bias in our maximum likelihood estimates. Ideally, we would like to isolate the household-level variation in our equity measure in the estimation of its effect on the probability of foreclosure. In a linear model, this would be trivial, as we would simply specify a fixed-effects specification, in which we would include a full set of time and geographic indicator variables (and interactions). Unfortunately, this strategy is not possible in the context of maximum likelihood estimation. First, there is a well known incidental parameters problem, in which inconsistent estimates of the time and geographic effects could also introduce bias into the estimates of the variables of interest. Second, such a strategy would involve the inclusion of thousands of right-hand-side variables, which coupled with the large size of the dataset, would prove computationally infeasible.

6.1. A simple linear probability model

To test the plausibility of our identification assumption, we explore an alternative to the non-linear, duration framework used above, and estimate a linear probability model, in which we are able to include a full set of time and geographic indicator variables. While a duration model is necessary to analyze many of the questions that we pose in this paper, such as characterizing the joint distribution of conditional foreclosure and sale probabilities, it is not necessary for estimating the relationship between equity and foreclosure. As we explained above, this relationship is non-linear, but this can easily be captured in the framework of a linear probability model, by specifying equity in terms of a spline or polynomial.⁴⁰ Our strategy is to estimate the linear probability model and isolate different components of the variation in our equity measure. If reverse causality or simultaneity is an issue, estimates of the different specifications will differ. For example, if foreclosure has a nontrivial effect on the path of house prices, then the estimate of the relationship between equity and foreclosure using only time-series variation in house prices will be biased, and thus different from the unbiased estimate of the relationship when only individual-level variation in purchase times is used. The specification that we use is the following:

$$\text{Prob}(\text{Foreclosure}_{ikt} = 1|E, X) = \alpha + \beta E_{ikt} + \gamma X_{ikt} + \delta_t + \eta_k + \xi_{kt} + \varepsilon_{ikt} \quad (16)$$

where the probability of foreclosure for individual i living in town k at time t is given by Foreclosure_{ikt} . Equity, E_{ikt} , is specified as a linear spline, in the same manner as it was specified in the duration model of section 3. The vector of other controls, X_{ikt} , is also identical to the corresponding vector included in the duration model. Depending on the specification, a full set of time effects, δ_t , town effects η_k , and their respective interactions, ξ_{kt} , may be included.⁴¹

Figure 7 displays the estimation results for four different specifications of the linear probability model. Each panel displays predicted foreclosure probabilities as a function of equity.⁴² The panels also include

³⁹Of course there is always the issue that the decision to purchase the home is endogenous, and it is possible that households may try to time the market with their purchase. But, compared to the identification issues discussed above, this would seem to be of second-order importance.

⁴⁰Of course one well known issue with linear probability models is the fact that predicted probabilities are not constrained to lie in the $[0, 1]$ interval.

⁴¹Since we are including multiple observations for each ownership in the data (specifically, one observation for each quarter that an ownership is active), we cluster the standard errors by ownership experience.

⁴²For the sake of brevity, we do not display all of the estimated parameters of the model. These are available upon request from the authors.

standard error bars, which are calculated using the Delta Method. The first panel does not control for time or town effects; the second panel includes a set of time effects; the third panel includes a set of town effects; and the fourth panel includes both time and town effects (and the respective interactions). The crucial observation to take away from the figure is the striking similarities between all of the panels. The specification with neither time or town effects (panel 1) produced virtually identical estimates to the specification with a full set of time and town effects (panel 4). This is substantial evidence in support of our identification assumption: That causality runs from house price appreciation to foreclosure, and not the other way around. In addition, the relationship between equity and foreclosure estimated in the context of the linear probability model is very similar to the relationship estimated in the competing-risks, duration model in Figure 6, which to some extent validates our choice of using a linear probability model in this section.

6.2. Further tests

In addition to estimating the linear probability models above, which we believe provides the most convincing support for our identification assumption, we also perform further robustness checks. First, we specify a formal Granger causality test on aggregate house price growth and foreclosure rates. Finally, we discuss a few empirical observations that provide further support for our assumption of the direction of causality.

Our econometric test for Granger causality follows directly from the autoregressive specification in Hamilton (1994).⁴³ We perform the estimation using quarterly data over the period 1975-2007, and include four lags of aggregate house price appreciation, the aggregate foreclosure rate, the unemployment rate, price inflation (CPI), wage-inflation, the Federal Funds Rate, and the output gap. We find a p-value of 0.60 from the F-test associated with the null hypothesis that foreclosures do not Granger-cause house price appreciation, and a p-value of 0.05 from the F-test associated with the null hypothesis that house price appreciation does not Granger-cause foreclosure rates. These results support our assumption that the direction of causality runs from house price appreciation to foreclosures. We view the inclusion of the unemployment rate in the Granger-causality test as especially important for validating the exogeneity assumption of cumulative house price appreciation given our discussion about the potential simultaneity issues above. Specifically, the inclusion of the unemployment rate in the Granger-causality test ensures that movements in house prices are not merely reflections of variation in unemployment.⁴⁴

Further support for our causality assumption is found in Figure 1. The figure shows that delinquencies, defined as a missed mortgage payment, are highly correlated with the Massachusetts business cycle, but that foreclosure rates are not as correlated with the business cycle. In recessionary periods, such as the early 1990s, and the early 2000s, we see large increases in delinquency rates. However, the key point is that we only see increases in foreclosure rates in periods of house price depreciation. For example, in the recessionary period of 2001, delinquency rates increased significantly, but foreclosure rates actually decreased. In contrast, during the recessionary period of 1991, both delinquency and foreclosure rates increased significantly. One of the big differences between these periods is the behavior of house price appreciation. In the 2001 recessionary period, Massachusetts house price growth was positive and increasing, while in the 1991 recessionary period, house price growth was negative and decreasing. This makes sense, since net equity is likely positive in periods of positive house price appreciation, and thus a household experiencing cash-flow problems can either sell their house or extract equity from their home by refinancing. In periods of negative house price growth, households suffering cash-flow problems and with negative equity in the home will be unable to sell

⁴³Hamilton (1994), section 11.2

⁴⁴In addition, unemployment rates at the town-level are included as a control variable in the duration model.

or refinance, and will thus be forced to default. Further evidence supporting this claim can be seen in the behavior of delinquency rates and foreclosure rates during the mid-1990s. In this period we see delinquency rates rise, but again foreclosure rates are falling as house prices are rising. These empirical observations are completely consistent with the theoretical prediction that negative equity is a necessary condition for default. Thus, based upon the implications from our theoretical model and the evidence presented above, we believe that the causality assumption in the empirical model is reasonable.

6.2.1. Measuring the Fit of the Model

To get an idea of how well the model fits the foreclosure data, we use the predicted probabilities from the model to calculate the expected number of foreclosures at each quarter of the ownership experience. The predicted probability of foreclosure for the i th ownership at quarter j can be calculated using the estimates from the competing risk model and the expression for the foreclosure hazard. Given the predicted probabilities for each ownership, λ_{ij} , we can calculate a predicted hazard rate at each quarter of the ownership:

$$\widehat{\lambda}_j = \frac{\sum_{i=1}^N \widehat{\lambda}_{ij}}{n_j} = \frac{\widehat{f}_j}{n_j}, \quad (17)$$

where \widehat{f}_j is the expected number of foreclosures at age j , and n_j is the number of borrowers “at-risk.” Figure 8 plots the predicted hazard rate $\widehat{\lambda}_j$ against the non-parametric, Kaplan-Meier hazard rate. The model seems to fit the data fairly well. The largest difference between the model and the data seems to be in the last four quarters of the ownership experience, when the model over-predicts default probabilities.

Figure 9 displays this same exercise for some of the specific cohorts of ownership experiences in our data that ranges from 1989 to 2008. While the model does not fit the data perfectly, it appears to do a very good job at getting the magnitudes and general shape of the default hazards correct.⁴⁵ Borrowers who bought their houses in 1989 and 2005 experienced the highest default hazards. This risk peaked at almost 0.7 percent after four years for the 2005 cohort, and approximately 0.5 percent after four years for the 1989 cohort. Both of these cohorts experienced significant price declines, as Massachusetts house prices (levels) reached a trough in 1993 and were still declining at the end of our data in December 2008.

Figure 10 displays the model’s in-sample goodness of fit with respect to subprime purchase ownership defaults. The figure displays the actual number of subprime foreclosures in each year in our sample of subprime borrowers versus the expected number of foreclosures \widehat{f}_j predicted by the model of this group of borrowers (a comparison of all ownerships is also displayed in the figure). The figure shows that the model is not perfect, as it slightly over-predicts foreclosures from 2007–2008. However, it does a remarkably good job of replicating the general pattern of foreclosures between 1993 and 2007, and it does predict the dramatic rise in foreclosures beginning in 2005 that we see in the data.⁴⁶

7. Counter-factual Exercise

We motivated this paper by claiming that we could address the viability of the role of bad underwriting and falling house prices in the foreclosure crisis that emerged in Massachusetts in 2007. To do this, we use the model estimated in Section 3 to calculate the predicted number of foreclosures under different house price scenarios, which gives us variation in HPA, and for different vintages of home buyers, which yields variation in underwriting standards. We know that vintages with bad underwriting also had bad house-price

⁴⁵It is important to note that time effects are not included in the estimation.

⁴⁶As a robustness check, we have also estimated the model separately for the subprime purchase ownerships and find very similar results.

outcomes and large numbers of foreclosures. We also know that vintages with good underwriting had good HPA outcomes and small numbers of foreclosures. To address the motivating question of the paper, we can look at vintages with good HPA and do the counter-factual exercise of giving them bad underwriting; if we find that outcomes are still extremely bad, that means that underwriting was the driver. In contrast, we can look at vintages with good underwriting and do the counter-factual exercise of giving them bad HPA; if we still find good outcomes, that would be further evidence that underwriting was the driver.

We focus on two vintages of home buyers, 2002 and 2005, because they perfectly illustrate our concept. Subprime played almost no role in the 2002 vintage, accounting for fewer than 3 percent of the purchases in the state, and house prices soared a cumulative X percent statewide between 2002 and 2005. And in the first three years of the vintage, there were almost no foreclosures; by the end of 2005 less than Y percent of the 2002 buyers had lost their homes. The contrast with 2005 could not be more marked. Subprime exploded, accounting for 15 percent of new home purchases; and house prices peaked sometime in the summer of 2005; by the end of 2008, they had fallen by Z percent. And after three years, Q percent of the 2005 buyers had lost their homes to foreclosure, an order of magnitude more than the 2005 vintage.

Our first counter-factual exercise is to take the successful 2002 vintage and give them the price experience of the 2005 buyers. Everything else about the vintage, all the purchase characteristics including LTV and the subprime purchase dummy as well as the dynamic variables like unemployment and interest rates remain the same. Comparing the lines marked 2002 with 2002 prices and 2005 with 2002 prices in Figure 11 illustrates our claim that deteriorating underwriting alone cannot explain the foreclosure crisis. Foreclosures increase but they remain small relative to what actually happened with the 2005 vintage. This is the evidence for our claim that bad underwriting alone could not have caused the foreclosure crisis. Because the proportional hazard specification is non-linear, one cannot simply interpret the gap between the counter-factual line for 2005 with 2002 prices and the 2005 line as the effect of prices, a point which we return to shortly.

Our second counter-factual exercise is to take the disastrous 2005 vintage and replace the buyers with the buyers from 2002. Again, we let the household experiences remain exactly the same except for the prices. Foreclosures fall by roughly half, but are still enormous relative to what actually happened to the 2002 buyers – the ones who got the good prices. We interpret this as meaning that prices alone would have resulted in a level of foreclosures that would have constituted a “crisis,” though, as we already said, not nearly of the magnitude of the one we have experienced.⁴⁷

Subprime loans do end up playing a large role in the crisis, because it is the difference between 15 percent subprime in the 2005 vintage and 3 percent in the 2002 vintage that accounts for the difference between the 2002 vintage with 2005 prices and the actual results for the 2005 vintage with 2005 prices. Subprime’s contribution to the crisis, therefore, was as an accelerant but not as the root cause. Fifteen percent of buyers in 2005 had some cocktail of problematic credit history, high payments and high LTV’s which would lead to a high likelihood of cash-flow problems. However, the logic of the lenders that made the loans was that those borrowers could either sell or refinance in that event, an outcome that became problematic when house prices began to fall in 2005.

8. Conclusion

Using micro-data that encompasses two housing cycles in the state of Massachusetts, and a novel identification strategy, we distinguish between two competing theories of the cause of the foreclosures crisis in the United States in 2007-2009. One theory holds that underwriting standards were substantially relaxed over

⁴⁷In fact, one of the authors participated in a “Foreclosure Crisis” panel in March of 2006 at which time the number of foreclosures was still a small fraction of what would obtain in 2008.

the course of the past decade, as evidenced by the subprime mortgage boom. The subsequent expansion of credit to lower income and minority borrowers involved the origination of a substantial number of high-cost, onerous loans that borrowers had trouble repaying, and which led directly to the foreclosure boom, and in turn, the significant decline in house prices. The other theory holds that a decline in house prices that began in late 2005 or early 2006 directly resulted in the foreclosure crisis, and that the growth in subprime mortgage credit exacerbated the situation. Our results lend support to both explanations, but place much more emphasis on the latter. We find that had house prices not fallen, the foreclosure crisis would not have occurred, regardless of whether lenders had lowered underwriting standards. By contrast, even if lenders maintained the same underwriting standards of the early 2000s, the observed fall in house prices would have generated a slightly smaller, but still significant increase in foreclosures.

To arrive at these conclusions, we employ a novel strategy to identify the effects of changes in house prices on changes in foreclosure rates. Our identification strategy involves exploiting within-town, within-time variation in household-level measures of housing equity, which has the advantage of not suffering from potential reverse causality or simultaneity issues. We employ this strategy on panel data from Massachusetts deed records spanning the past 20 years. The data is novel and especially useful in addressing these issues for several reasons. First, it contains the universe of all residential property and mortgage transactions in Massachusetts, allowing us to follow borrowers for the entire length of stay in their respective homes, regardless of the number of times they refinance. Second, it contains a long time series that spans two substantial house price downturns. Finally, the data allow us to observe and match all liens on a given property, providing us with accurate measures of the cumulative loan-to-value (CLTV) ratios that are essential for calculating a borrower's equity position. Our estimates of a proportional hazard, competing risks model of homeownership termination show that the relationship between equity and foreclosure is highly non-linear, with the foreclosure hazard essentially zero for equity levels greater than 25 percent, but then rising sharply for equity levels below 25 percent.

To motivate our empirical model, we construct a simple, two-period model of homeownership termination that allows for a non-trivial role for household characteristics in the default decision. It does so by embedding the traditional, option-theoretic default model in a portfolio choice framework. In the model, the household's financial situation affects the stochastic discount factors used to discount the costs and benefits of continuing to make mortgage payments. Through this channel, the model provides a theoretically robust argument for the heuristic "double-trigger" model of default that is often referenced (and assumed) in the literature, but which to our knowledge has not yet been formally derived. This model posits that foreclosure occurs when a borrower faces negative equity and cash-flow problems, but a crucial point that differentiates our model from others in the literature is that it does not distinguish between "strategic" and "non-strategic" default. In the model, borrowers have access to unsecured credit, implying that a mortgage payment is never, strictly speaking, unaffordable. In such an environment, a borrower in financial distress with negative equity will choose not to default as long as they believe that house prices will rise rapidly in the future. But, it is also possible in the model that a borrower in good financial health may default if he or she believes that house prices are going to continue to fall.

References

- [1] Black, F., and M. Scholes, 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81(3):637–654.
- [2] Campbell, T., J. Dietrich, 1983. "The Determinants of Default on Conventional Residential Mortgages." *Journal of Finance* 38(5): 1569–1581.
- [3] Campbell, J. Y., Giglio, S., and P. Pathak. 2009. "Forced Sales and House Prices." *NBER Working Paper* No. w14866.
- [4] Case, K., and R.J. Shiller, 1987. "Prices of Single Family Homes Since 1970: New Indexes For Four Cities." *NBER Working Paper* No. 2393.
- [5] Chan, S., 1996. "Residential Mobility and Mortgages." *Regional Science and Urban Economics* 26:287–311.
- [6] Coleman IV, M. D., M. Lacour-Little, and K. D. Vandell (2008). Subprime Lending and the Housing Bubble: Tail Wags Dog? SSRN eLibrary.
- [7] Cox, J., J. Ingersoll, and S. Ross, 1985. "An Intertemporal General Equilibrium Model of Asset Prices." *Econometrica* 53(2): 363–384.
- [8] Deng, Y., and J. Quigley, 2004. "Woodhead Behavior and the Pricing of Residential Mortgages." *Program on Housing and Urban Policy Working Paper Series* No. W00-004.
- [9] Deng, Y., J. Quigley, and R. Van Order, 2000. "Mortgage Terminations, Heterogeneity and the Exercise of Mortgage Options." *Econometrica* 68(2): 275–307.
- [10] Dubin, Robin. "Foreclosures in Cleveland." Paper presented at the Midyear Meeting of the American Real Estate and Urban Economics Association, Washington, D.C., May 28, 2008.
- [11] Epperson, J., Kau J., D. Keenan, and W. Muller, 1995. "The Valuation at Origination of Fixed-Rate Mortgages with Default and Prepayment." *Journal of Real Estate Finance and Economics* 11: 5–36.
- [12] Hamilton, J., 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.
- [13] Han, A., and J. Hausman, 1990. "Flexible Parametric Estimation of Duration and Competing Risk Models." *Journal of Applied Econometrics* 5(1): 1–28.
- [14] He, H., and N. Pearson, 1991. "Consumption and Portfolio Policies with Incomplete Markets and Short-Sale Constraints: The Finite-Dimensional Case." *Mathematical Finance* 1(3): 1–10.
- [15] Henderson, J., and Y. Ioannides, 1989. "Dynamic Aspects of Consumer Decisions in Housing Markets." *Journal of Urban Economics* 26: 212–230.
- [16] Henly, A., 1998. "Residential Mobility, Housing Equity and the Labour Market." *The Economic Journal* 108(447): 414–427.
- [17] Herzog, J., and J. Early, 1970. *Home Mortgage Delinquency and Foreclosure*. New York: National Bureau of Economic Research 2006-042A.
- [18] Immergluck, D., and G. Smith 2006. "The Impact of Single-Family Mortgage Foreclosures on Neighborhood Crime." *Housing Studies* 21(6): 851–866.
- [19] Kau J., and D. Keenan, 1995. "An Overview of the Option-Theoretic Pricing of Mortgages." *Journal of Housing Research* 6(2): 217–244.
- [20] Kau J., D. Keenan, and T. Kim, 1994. "Default Probabilities for Mortgages." *Journal of Urban Economics* 35: 278–296.
- [21] Leonard, T., and Jim Murdoch. "The Neighborhood Effects of Foreclosure." Paper presented at the Midyear Meeting of the American Real Estate and Urban Economics Association, Washington, D.C., May 28, 2008.
- [22] McCall, B., 1996. "Unemployment Insurance Rules, Joblessness, and Part-Time Work." *Econometrica* 64(3):647-682.
- [23] Meyer, B., 1995. "Semiparametric Estimation of Hazard Models." MIT working paper.
- [24] Mian, A. R. and A. Sufi, 2009. The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis. *Quarterly Journal of Economics*, forthcoming.
- [25] Munnell, A., Tootell, G., Browne, L., and J. McEneaney. 1996. "Mortgage Lending in Boston: Interpreting HMDA Data." *American Economic Review* 86(1): 25-53.
- [26] Pavlov, A. D. and S. M. Wachter (2006). Underpriced Lending and Real Estate Markets. SSRN eLibrary.
- [27] Pennington-Cross, A., and G. Ho, 2006. "The Termination of Subprime Hybrid and Fixed Rate Mortgages." *Federal Reserve Bank of St. Louis Working Paper Series* 2006-042A.
- [28] Quercia, R., and M. Stegman, 1992. "Residential Mortgage Default: A Review of the Literature." *Journal of Housing Research* 3(2): 341–379.
- [29] Rogers, William H., and William Winter. "The Impact of Foreclosures on Neighborhood Housing Sales." Paper presented at the Midyear Meeting of the American Real Estate and Urban Economics Association, Washington, D.C., May 28, 2008.
- [30] Sanders, A., S. Chomsisengphet, S. Agarwal, and B. Ambrose (2008). Housing prices and alternative mortgage concentrations. Working Paper.
- [31] Wheaton, W. C. and N. J. Lee, 2008. Do Housing Sales Drive Housing Prices or the Converse? SSRN eLibrary.

- [32] Wheaton, W. C. and G. Nechayev, 2008. The 1998-2005 Housing Bubble and the Current Correction: Whats Different this Time? *Journal of Real Estate Research*, Vol. 30, No. 1, 2008.
- [33] Sueyoshi, G., 1992. "Semiparametric Proportional Hazards Estimation of Competing Risks Models with Time-Varying Covariates." *Journal of Econometrics* 51: 25-58.
- [34] Vandell K., 1995. "How Ruthless is Mortgage Default? A Review and Synthesis of the Evidence." *Journal of Housing Research* 6(2): 245-264.
- [35] von Furstenberg, G. 1969. "Default Risk on FHA-Insured Home Mortgages as a Function of the Term of Financing: A Quantitative Analysis." *Journal of Finance* 24: 459-477.
- [36] Williams, A., W. Beranek, and J. Kenkel, 1974. "Default Risk in Urban Mortgages: A Pittsburgh Prototype Analysis." *AREUEA Journal* 2(2): 101-102.
- [37] Yu, X., 2006. "Competing Risk Analysis of Japan's Small Financial Institutions." *IMES Discussion Paper Series* No. 2006-E-7.

Appendix

A. Proof of proposition

Proof:

We start the proof by dividing the problem into two optimization problems, one in which the borrower always keeps the house (K) and one in which the borrower always defaults (D). Let $\bar{c}_K(w)$ and $\bar{c}_D(w)$ be the constant level of consumption which yields the same utility as the optimal consumption profile for keepers and defaulters respectively. We now characterize the critical value of wealth, w^* such that $w > w^* \Leftrightarrow \bar{c}_K(w) > \bar{c}_D(w)$. Let w_L^i be the value of wealth such that if $w > w_L^i$ the consumer saves, and if $w < w_B^i$ the consumer borrows. It follows that for each problem the two critical values $w_L^i > w_B^i$, where $i = K, D$, have the following properties:

$$w > w_L^i \Rightarrow \theta_L > 0, \theta_B = 0 \quad (18)$$

$$w_L^i \geq w \geq w_B^i \Rightarrow \theta_L = \theta_B = 0 \quad (19)$$

$$w < w_B^i \Rightarrow \theta_L = 0, \theta_B > 0, \quad (20)$$

where θ_L is dollars saved and θ_B is dollars borrowed. Since we assume that it is more expensive to own the house than to rent the house, keepers have less current consumption than defaulters. Thus, under (7), keepers' marginal utility of consumption is high relative to that of defaulters. Keepers therefore wish to consume more than defaulters, which pushes their thresholds for both borrowing and saving higher:

$$w_L^K > w_L^D \text{ and } w_B^K > w_B^D. \quad (21)$$

To solve the general problem, we start with $w > w_L^K$ which by equations (18) and (21) implies that $\theta_L > 0$ for both keepers and defaulters. Since both consumers are at interior solutions, the permanent income hypothesis holds and consumption is proportional to lifetime income, Y , which is

$$Y_i = \begin{cases} y_0 + w + \frac{y_1}{R_L} + \frac{(H-M)^+}{R_L} - r_M M, & i = K \\ y_0 + w + \frac{y_1}{R_L} - \rho, & i = D \end{cases}$$

Our assumption that $R_K > R_L$ implies that $Y_K > Y_D$ since,

$$R_K = \frac{(H-M)^+}{r_M M - \rho} > R_L \quad (22)$$

$$\frac{(H-M)^+}{R_L} > r_M M - \rho \quad (23)$$

$$\frac{(H-M)^+}{R_L} - r_M M > -\rho \quad (24)$$

Thus, when $w > w_L^K$, $\bar{c}_K(w)$ exceeds $\bar{c}_D(w)$. Similar logic shows that when $w < w_B^D$, $\bar{c}_D(w)$ exceeds $\bar{c}_K(w)$.

To prove our basic claim, we simply need to show that the slope of $\bar{c}_K(w)$ always exceeds the slope of $\bar{c}_D(w)$ which implies that the two lines cross at a single point. To do this, define the ‘‘shadow riskless rate’’ as:

$$1 + r = \frac{c_1}{\beta c_0}. \quad (25)$$

Solving the Euler equations reveals that:

$$\partial \bar{c} / \partial w = \frac{1}{1 + \beta} \beta^{\frac{\beta}{1+\beta}} (1 + r)^{\frac{\beta}{1+\beta}} \quad (26)$$

We now show that the shadow riskless rate for keepers exceeds the shadow riskless rate for defaulters in the range $[w_L^K, w_B^D]$. We look at three regions within this range.

We start with the region $w_L^K \geq w \geq w_L^D$, in which the keeper does not save but the defaulter does, so the shadow riskless rate for the keeper must exceed R_L , otherwise $\theta_L > 0$. Thus, in this first region $(1 + r_K) > R_L = (1 + r_D)$ which implies that $\frac{\partial \bar{c}_K}{\partial w} > \frac{\partial \bar{c}_D}{\partial w}$.

In the second region, $w_L^D \geq w \geq w_B^K$, $\theta_B = \theta_L = 0$ and neither type saves. When both keepers and defaulters do not save, equations (1-2) and (5-6) yield:

$$\begin{aligned} \frac{c_1}{c_0} &= \frac{y_1 + (H - M)^+}{y_0 + w - r_M M} \text{ for } i = K \\ \frac{c_1}{c_0} &= \frac{y_1}{y_0 + w - \rho} \text{ for } i = D. \end{aligned}$$

Since we assume that $r_M M > \rho$ and $(H - M)^+ > 0$, equation 25 gives us the inequality:

$$1 + r_K = \frac{y_1 + (H - M)^+}{y_0 + w - r_M M} > \frac{y_1}{y_0 + w - \rho} = 1 + r_D$$

Which from (26) implies $\frac{\partial \bar{c}_K}{\partial w} > \frac{\partial \bar{c}_D}{\partial w}$.

In the region $w_B^K \geq w \geq w_B^D$, in which the defaulter holds no assets but the keepers borrows so the shadow riskless rate for the defaulter must fall short of R_B , otherwise $\theta_B > 0$. Increases in r_M , and M_{-1} and decreases in H shift the \bar{c}_K curve up without affecting the \bar{c}_D curve, leading to increases in w^* . Increases in ρ shift the \bar{c}_D curve down reducing w^* . To show the remaining results requires a straightforward application of the Implicit Function Theorem \square

B. Construction of Weighted, Repeat-Sale House Price Indexes

We follow Case and Shiller (1987) and Deng, Quigley, and Van Order (2000) in constructing our house price indexes. The index for each town is based on the assumption that the house price of household i at time t is given by the following process:

$$\ln P_{it} = \ln \bar{P}_t + \mu_{it} + \eta_{it}, \quad (27)$$

where \bar{P}_t is the house price level of the town, η_{it} is white noise, and μ_{it} is a Gaussian random walk with mean equal to zero, $E[\mu_{i,t+k} - \mu_{it}] = 0$, and variance proportional to the age of the loan, $E[\mu_{i,t+k} - \mu_{i,t}]^2 = k\sigma_1^2 + k^2\sigma_2^2$.

The index is calculated using a three-stage process on paired sales of one-family houses. High frequency sales—occurring six months or less apart—were dropped. Sales with appreciation greater than 50 percent within one year, greater than 100 percent within 1.5 years, and 150 percent within two years were dropped due to the high probability of renovation. Sales with depreciation larger than 25 percent within one year and 50 percent within two years were dropped because of the high probability of a subdivision or heavy deterioration reducing the property's previous value. House prices less than \$15,000 and greater than \$10 million were dropped.

House price indexes were developed for every Massachusetts city with more than 5,000 sale observations between January 1987 and August 2007. All other towns were excluded from the sample. In the first stage, the log price of the second sale minus the log price of the first sale is regressed on a set of time dummy variables, D_t :

$$\dot{p}_i = \sum_{t=1}^T \beta_t D_t + \omega_i, \quad (28)$$

where $\dot{p}_i = \ln P_{i,t+k}^{2nd} - \ln P_{i,t}^{1st}$, and the dummy variables have the value +1 for the time of the second sale, and the value -1 for the time of the first sale.

In the second stage, the squared residuals, ω_i^2 , are regressed on k and k^2 :

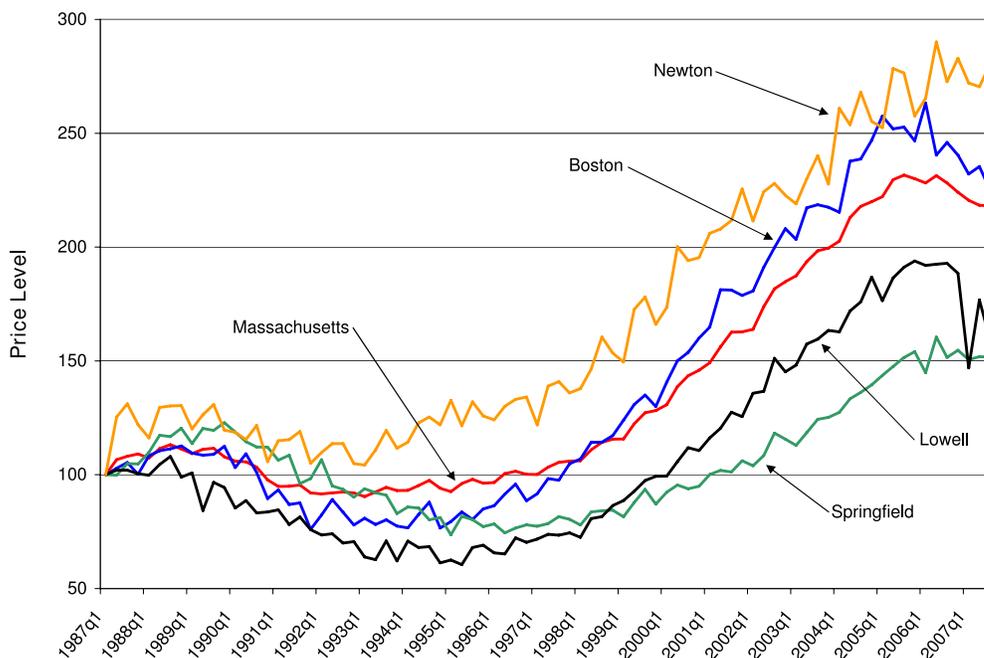
$$\omega_i^2 = A + Bk + Ck^2. \quad (29)$$

In the third stage, equation 28 is estimated by GLS, using $\hat{\omega}_i$ —the square roots of the predicted values of equation 29—as weights.

The house price index is then constructed from the estimates β_t :

$$P_t^{Index} = 100 * \exp(\beta_t). \quad (30)$$

In the figure below, we plot P_t^{Index} for the entire state of Massachusetts, as well as four cities—Boston, Springfield, Lowell, and Newton—which are indicative of the state's various geographic regions, as well as different ethnic and income pools.



C. Robustness Check for Subprime Lender Indicator

This appendix provides a detailed robustness check of our subprime purchase indicator variable. As we discussed above, this variable is based on HUD’s annual list of lenders who originate predominantly subprime loans. This list is certainly an imperfect proxy for the subprime market, as the mortgage lenders on the HUD list do not do business in the subprime market exclusively. Moreover, there are lenders who are not on the HUD list, but who do provide loans to subprime borrowers. Therefore, by using the HUD subprime list, we are introducing two potential mistakes into our analysis. First, we are falsely labeling some mortgages as subprime that were really made to prime borrowers, a misidentification that we call “false positives.” Second, we are falsely labeling some mortgages as prime that are really issued to subprime borrowers, a misidentification that we call “false negatives.” If our results are to be taken seriously, it is important to estimate the magnitude of both of these identification errors.

Fortunately, we have a limited amount of information regarding mortgage interest rates in the Warren Group data that makes it possible to check the magnitude of these errors. We have detailed interest rate information on approximately 5 percent of mortgages originated between 2004 and 2006. This includes the initial interest rate, the date when the interest rate is first allowed to change (often referred to as the “hybrid” term), the index that the rate is tied to once it is allowed to fluctuate, and the difference between the mortgage rate and this index. This difference is commonly referred to as the “margin.” However, this information is available only for adjustable-rate mortgages.⁴⁸

With these variables we construct our own definition of a “high-cost” mortgage and use this to calculate subprime mortgage indicators to compare with the indicators taken from the HUD list. First, we took all of the adjustable-rate mortgages with initial interest rate information, and divided these into two categories: those with initial fixed-rate terms (hybrid terms) less than three years inclusive, and those with initial fixed-rate terms over three years. Then, we took the Federal Home Loan Mortgage Corporation (FHLMC), prime, 1-year adjustable interest rate series, and subtracted it from the initial interest rates for mortgages with initial fixed-rate terms less than three years. Similarly, we took the Freddie Mac, prime, 5-year adjustable interest rate series and subtracted it from the initial interest rates for mortgages with initial fixed-rate terms over three years. We labeled as subprime all mortgages for which this difference was greater than 200 basis points, *or* for which the margin was greater than 350 basis points. This method is a somewhat arbitrary way to define a subprime mortgage; however, the term “subprime” itself is somewhat arbitrary, and a universal definition of a subprime mortgage does not currently exist. Some analysts use a secondary market definition and define a subprime mortgage to mean a loan that is part of a pool of securitized mortgages that has been labeled “subprime.” However, in many cases subprime mortgage pools also contain a non-negligible number of mortgages held by borrowers that many analysts would consider to be credit-worthy prime borrowers.⁴⁹ Other analysts, including those who work with HMDA data, define a subprime mortgage as we do, based on the spread between the contracted interest rate and a market index; however, there is no general consensus on the magnitude of the spread. We believe that our subprime definition is conservative, as a 200 basis point spread between the initial interest rate and the corresponding FHLMC prime series is large, and a 350 basis point margin is also quite large. For example, a two-year hybrid, adjustable-rate mortgage originated in July 2007, with an initial interest rate of 8 percent, and a margin of 350 basis points with respect to the 6-month

⁴⁸The information is taken directly from the mortgage documents in the Registries of Deeds. Adjustable-rate mortgage documents contain “riders” that state all of the information necessary for a borrower to calculate her interest at each stage of the mortgage. Fixed-rate mortgage documents do not contain any interest rate information. Interest rate information for fixed-rate mortgages is located in the mortgage note, which is not publicly available.

⁴⁹For example, we found using data from LoanPerformance Inc. for Middlesex County, MA, that 20 percent of mortgages in subprime pools are made to borrowers with FICO scores above 700.

LIBOR, which would translate into an interest rate of approximately 9 percent (assuming that interest rates remain constant over the next two years), would barely qualify as a subprime mortgage using our definition.

The table below displays a cross-tabulation of the subprime variable constructed with the HUD list, versus the subprime variable constructed using the interest rate information. The most important statistics for our purposes are the two that are highlighted in bold font. The interpretation is that 7.85 percent of mortgages from a subprime lender on the HUD list are considered to be prime by our definition, using the spread between the contract rate and the appropriate FHLMC series. Similarly, 26.04 percent of mortgages from a non-subprime lender on the HUD list are considered to be subprime.

	HUD subprime indicator		Total
	<i>prime</i>	<i>subprime</i>	
<i>prime</i> (<i>< 200 basis point difference</i>)	43,485 96.97% 73.96%	1,361 3.03% 7.85%	44,846
<i>subprime</i> (<i>≥ 200 basis point difference</i>)	15,308 48.92% 26.04%	15,981 51.08% 92.15%	
Total	58,793	17,342	

With these two statistics we can develop a simple model to estimate the effect of these errors on foreclosure rates. If we let D_{sub}^* be the observed foreclosure rate for subprime ownerships, D_{prime}^* the observed rate for prime ownerships, and D_{sub} and D_{prime} the respective true rates, we can write the observed rates as linear functions of the true rates and the errors:

$$D_{sub}^* = D_{sub} * (1 - \alpha) + D_{prime} * \alpha \quad (31)$$

and

$$D_{prime}^* = D_{sub} * (1 - \beta) + D_{prime} * \beta. \quad (32)$$

In equation 31, α is the fraction of mortgages identified as subprime from the HUD list that are really prime, and β is the fraction of mortgages identified from the HUD list as prime that are really subprime. Since we have estimates of both α and β from the table above, the system of equations 31 and 32 contains two unknowns (the true, unbiased foreclosure rates), and thus it is easy to solve. If we assume a spread of 200 basis points and 350 basis points as our subprime definition criteria, our result that approximately 18 percent of subprime ownership experiences end in default after 12 years, then changes to approximately 19.5 percent of subprime ownerships ending in default. It is important to note that the unbiased foreclosure rates are sensitive to our estimates of α and β , which in turn are a function of the spread between the contracted mortgage rate and the rate upon which we based our subprime definition. Furthermore, the key assumption on which the above relationship depends is that the sample of subprime mortgages that we capture with the HUD list and the entire population of subprime mortgages are not systematically different.

Table 1: Some facts about the Massachusetts housing market, 1988-2008

	Fore-closures	Sales	Purchase Initial Loan-to-Value Ratios					
			All ownerships			Ownerships that Default		
			#	mean	median	#	mean	median
1988			68,982	0.76	0.80	5,761	0.83	0.80
1989			57,309	0.77	0.80	4,412	0.85	0.88
1990	1,730	60,878	46,407	0.79	0.80	2,551	0.88	0.90
1991	5,638	62,147	47,890	0.80	0.80	1,279	0.90	0.95
1992	9,236	73,532	56,398	0.80	0.80	955	0.91	0.95
1993	8,119	81,305	62,545	0.82	0.85	965	0.92	0.95
1994	7,049	88,603	68,668	0.82	0.87	1,015	0.92	0.95
1995	4,662	82,444	63,687	0.83	0.87	936	0.94	0.97
1996	4,202	92,240	72,493	0.83	0.87	945	0.94	0.97
1997	3,825	98,514	77,759	0.83	0.85	959	0.93	0.97
1998	2,764	109,294	88,004	0.83	0.85	859	0.92	0.95
1999	2,030	110,484	88,950	0.82	0.85	974	0.92	0.95
2000	1,461	103,414	83,452	0.81	0.82	1,007	0.91	0.95
2001	1,078	99,803	81,638	0.82	0.85	977	0.92	0.95
2002	952	107,010	87,281	0.81	0.82	1,172	0.92	0.95
2003	587	109,866	89,432	0.82	0.85	1,564	0.92	0.95
2004	643	121,337	98,057	0.82	0.87	2,765	0.93	0.98
2005	914	115,816	95,505	0.83	0.90	3,641	0.95	1.00
2006	2,805	96,985	81,016	0.84	0.90	2,598	0.96	1.00
2007	7,142	82,922	68,764	0.84	0.90	349	0.95	1.00
2008	9,362	50,036	39,776	0.82	0.85	3	0.83	0.97

Sale and foreclosure numbers come from data compiled by The Warren Group, and reflect only residential properties. Data from 2008 are available through September. Each year corresponds to a specific ownership cohort. Loan-to-value ratios are cumulative in the sense that they include all mortgages obtained at the time of purchase, including second and third mortgages where applicable.

Table 2: Subprime Lenders in Massachusetts

	# loans	% of subprime purchase mortgages	status
Option One Mortgage Corp.	11,204	18.7	shut down
New Century Financial Corp.	5,973	10.0	shut down
Freemont Investment & Loan	5,567	9.3	shut down
Argent Mortgage Company	3,602	6.0	shut down
Summit Mortgage Company	3,282	5.5	shut down
Mortgage Lender Net	2,825	4.7	shut down
Long Beach Mortgage Corp.	2,504	4.2	shut down
WMC Mortgage Corp.	2,307	3.9	shut down
Accredited Home Lenders	2,174	3.6	shut down
First Franklin Financial	1,871	3.1	shut down
Total	41,309	69.0	-

Note: This is a list of the top ten subprime lenders in terms of number of purchase mortgage originations in Massachusetts from 1993 to 2008. The status of each lender is updated through January 2009.

Table 3: Subprime Foreclosure Statistics in Massachusetts

	Subprime Ownerships Started	Foreclosures from ownerships with		
		subprime purchase	last mortgage subprime	at least one subprime
1993	0.09%	0.00%	0.00%	0.00%
1994	0.37%	0.00%	0.00%	0.00%
1995	0.42%	0.00%	0.15%	0.17%
1996	0.89%	0.05%	0.52%	0.67%
1997	1.91%	0.16%	1.59%	1.83%
1998	2.58%	0.90%	4.16%	4.78%
1999	2.46%	3.35%	8.92%	11.18%
2000	2.47%	4.38%	12.53%	15.81%
2001	2.90%	4.73%	14.10%	19.67%
2002	3.92%	7.35%	17.86%	24.47%
2003	6.91%	9.71%	22.83%	29.47%
2004	10.06%	10.89%	26.28%	35.46%
2005	14.81%	17.83%	37.53%	47.92%
2006	13.04%	28.31%	44.99%	58.68%
2007	3.12%	30.52%	43.00%	57.39%
2008	0.00%	30.21%	39.29%	55.61%

Note: Subprime foreclosure numbers are calculated as percentages of the total number of foreclosures in a given year.

Table 4: Estimation Results

(1) Descriptive Statistics

	Mean	Std. dev.	Min	Max
initial CLTV (fraction)	0.82	0.17	0.1	1.25
subprime purchase indicator	0.04	0.19	0	1
condo indicator	0.22	0.41	0	1
multi-family indicator	0.1	0.31	0	1
equity (% of original mortgage balance)	0.74	0.95	-0.78	23.2
LIBOR(6-month)	4.33	1.85	1.17	10.02
cumulative unemployment change (town-level)	0.06	0.5	-1	5
median family income (2000 zip code-level)	56517	18877	15648	153918
% of minority households (2000 zip code)	15.1	16.22	0	98.72

(2) Coefficient Estimates

	Default		Sale	
	Coef	Std. Err.	Coef	Std. Err.
initial CLTV	0.17	0.19	1.50	0.04
LIBOR	-0.03	0.01	-0.06	2.88e-03
cumulative unemployment	0.023	0.04	-0.08	9.83e-03
median income (2000 zip code)	-2.23e-05	1.54e-06	1.59e-06	2.63e-07
% minority (2000 zip code)	5.35e-03	9.19e-04	1.33e-03	3.06e-04
subprime purchase indicator	1.82	0.04	0.34	0.02
condo indicator	0.10	0.05	0.57	0.01
multi-family property indicator	0.52	0.04	6.26e-03	0.02
# unique borrowers	194,374		194,374	
Log Likelihood	-310,158		-310,158	

Note: Specification includes semi-parametric, baseline hazard, expressed in 4-quarter intervals.

(3) Standardized Elasticities

	(+/-) std. dev.	Default	Sale
		% change hazard	% change hazard
initial ltv (fraction)	0.17	1.03	1.29
LIBOR – 6-month (percentage points)	1.85	0.94	0.90
cumulative unemployment change (fraction)	0.5	1.12	0.96
% minority (2000 zip code)	19.58	1.09	1.02
median income (2000 zip code)	\$24,493	0.66	1.03
multi-family indicator	.	1.68	1.01
condo indicator	.	1.11	1.01
subprime purchase indicator	.	6.18	1.40

Note: Continuous explanatory variables are increased/decreased by one standard deviation while dichotomous explanatory variables are changed from zero to one.

Figure 1: Massachusetts House Price Growth, Foreclosures and Delinquencies, January 1989 to August 2007

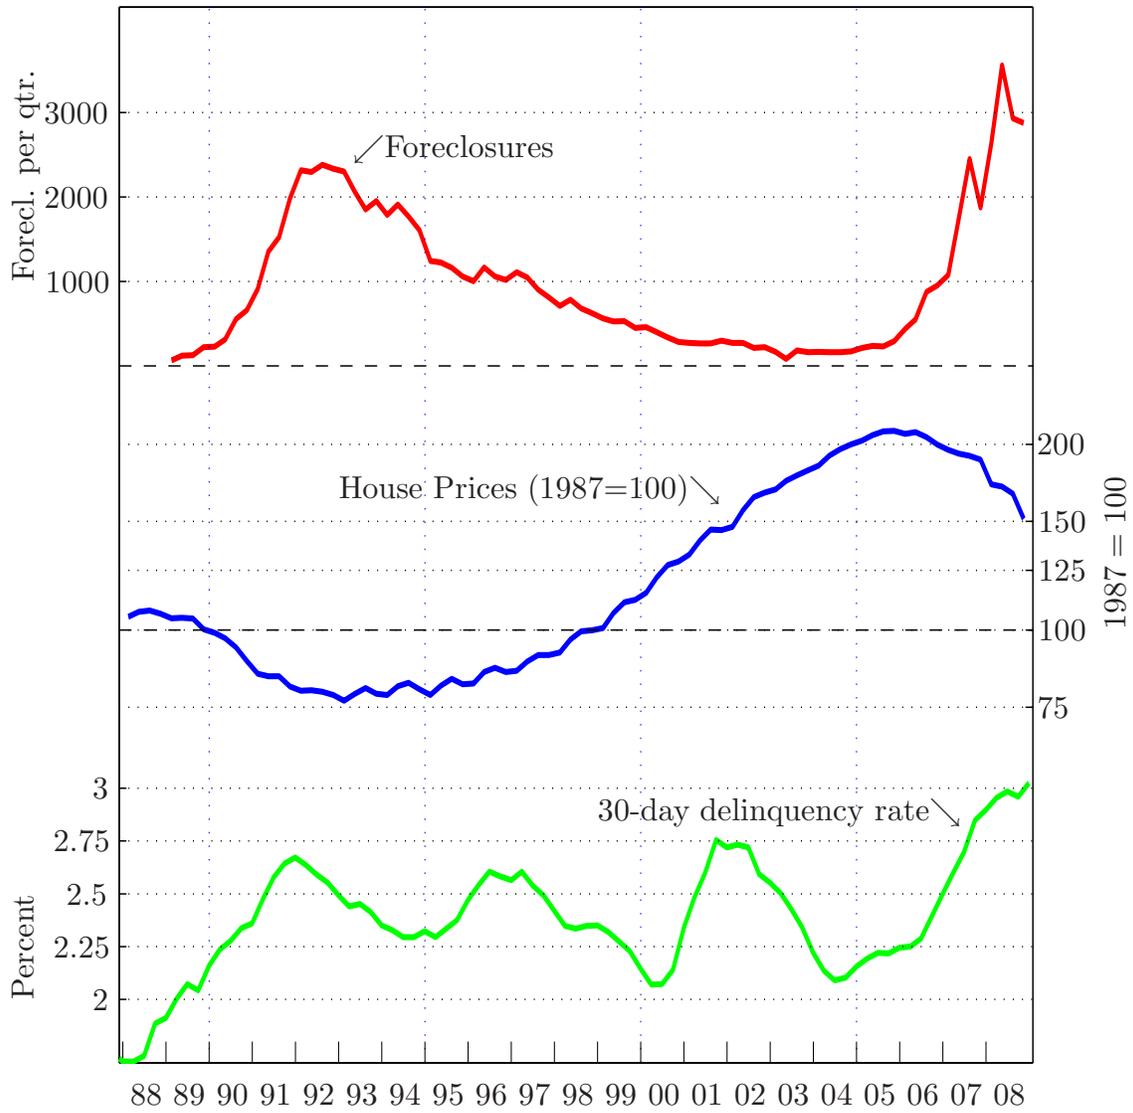
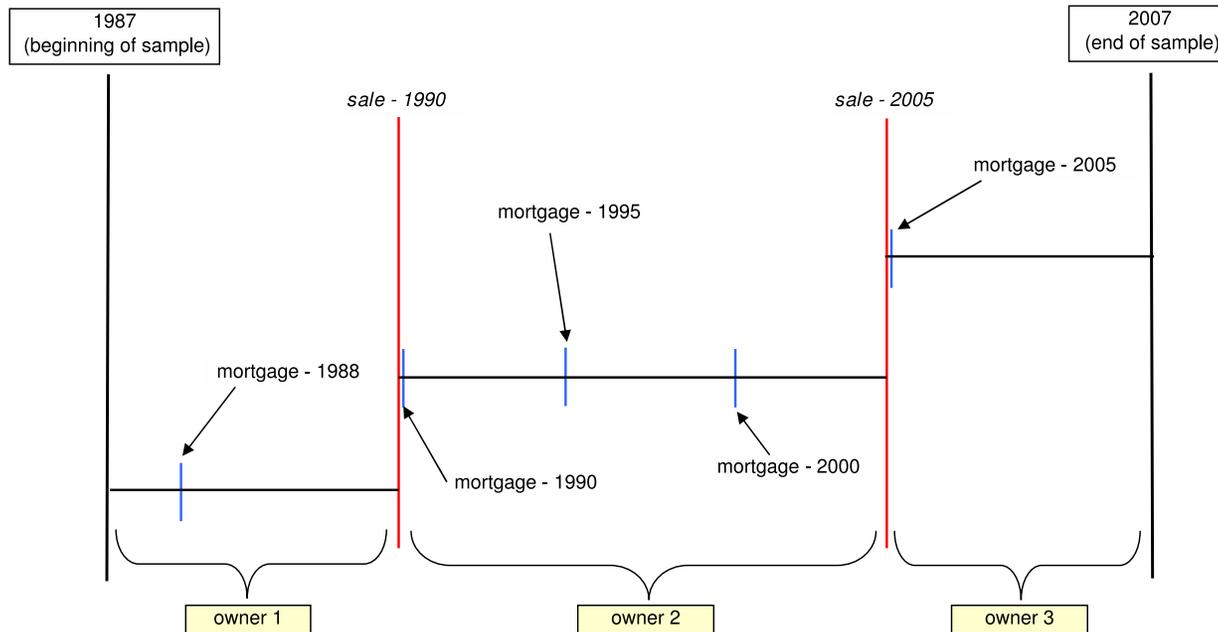


Figure 2: Example of the Ownership Experiences Associated with a Single Property



37

Notes: This example considers a property in which we see two sale transactions (red lines) and five mortgage transactions (blue lines). In this case we are able to identify three separate ownership experiences associated with this property. The first ownership experience comprises all transactions between the beginning of our data and the first sale of the property (1 mortgage). The second ownership experience includes all the transactions between the first and second sale of the property (3 mortgages), and finally, the third ownership experience includes all transactions between the second sale and the end of our data set (1 mortgage). The example in the figure illustrates that censoring is an issue with our data. In this particular example, the first ownership experience is left-censored, since we do not see the start of the experience, while the third ownership experience is right-censored, since we do not see the end of the experience. We address how we treat censored observations in our discussion of the estimation method in section 3.

Figure 3: Comparison of Massachusetts and the rest of the country. Source: MBA National Delinquency Survey.

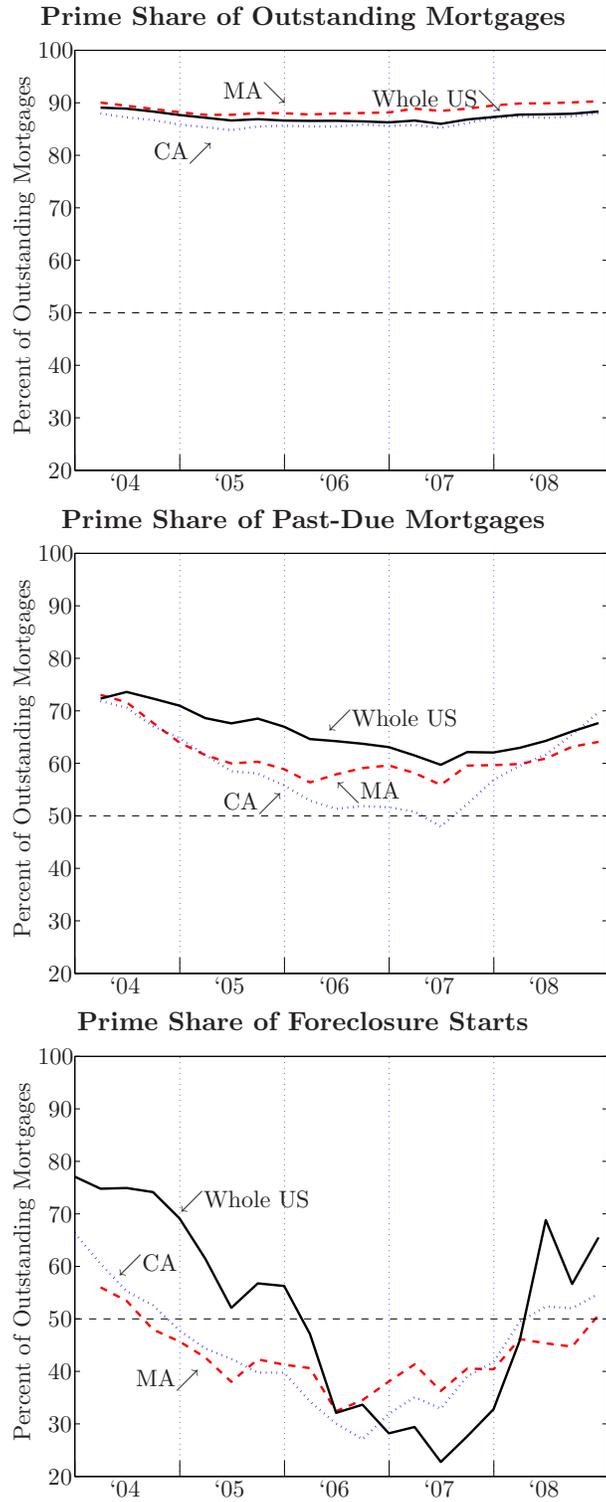


Figure 4: Kaplan-Meier Non-Parametric Sale and Foreclosure Hazard Rates for Massachusetts

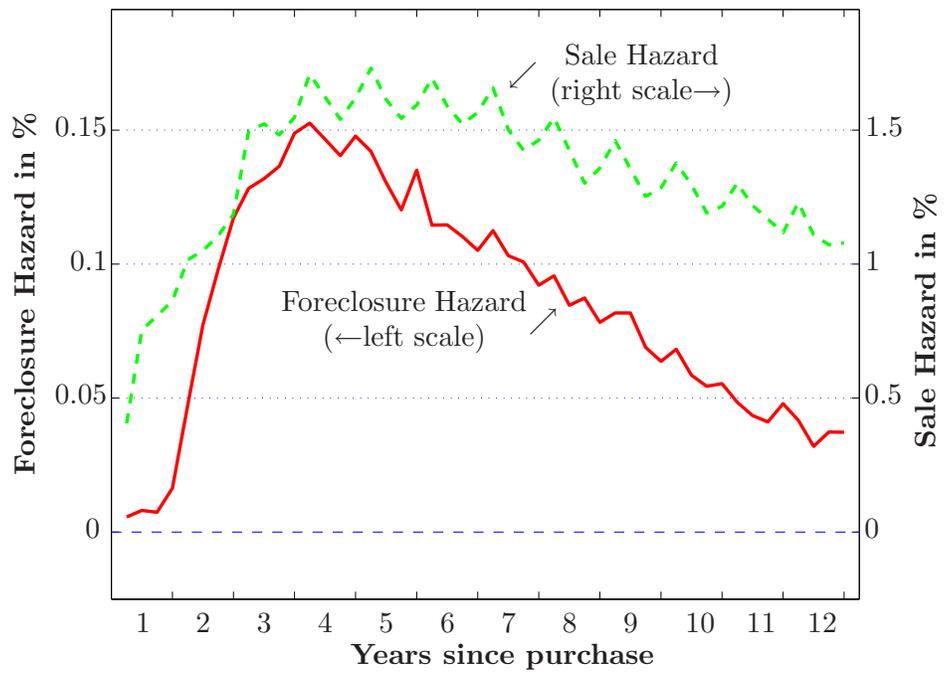


Figure 5: Kaplan-Meier Non-Parametric Default Hazard Rates for different samples

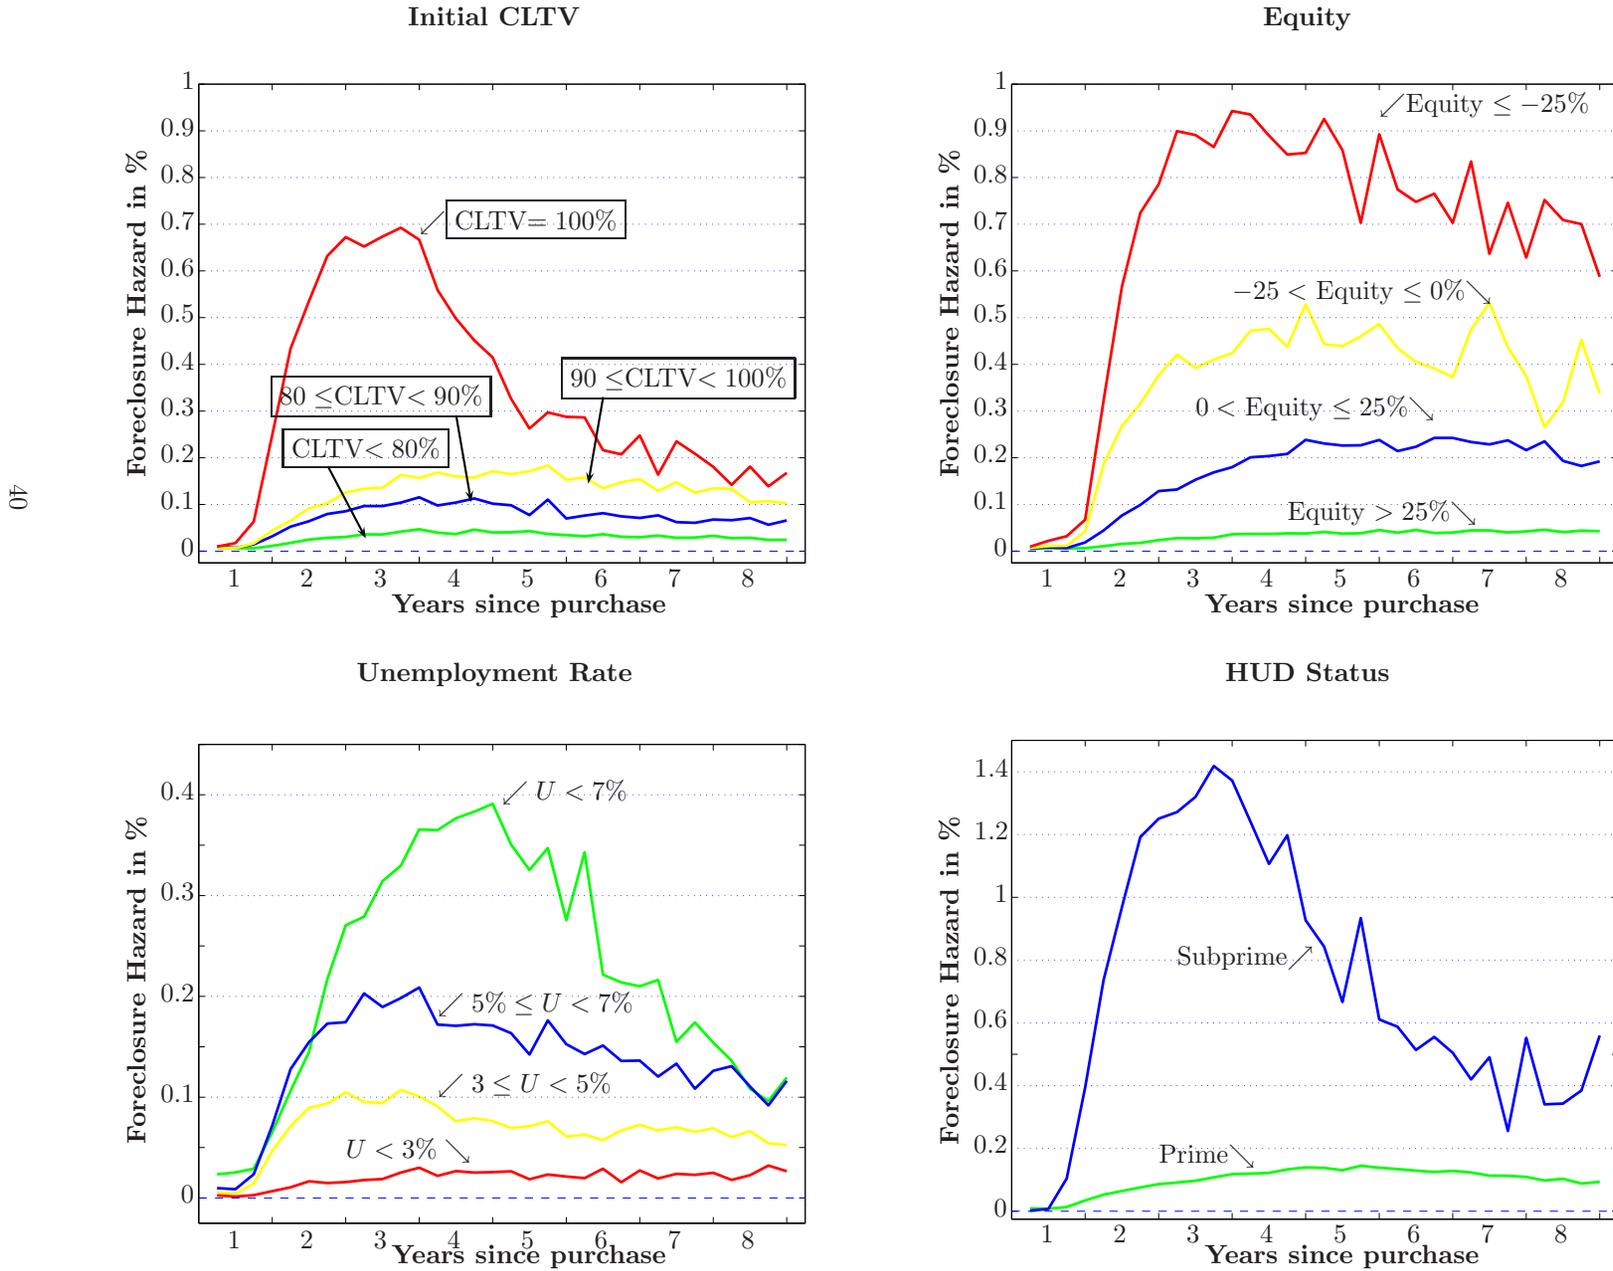


Figure 6: Estimated Effect of Equity on Foreclosure and Sale

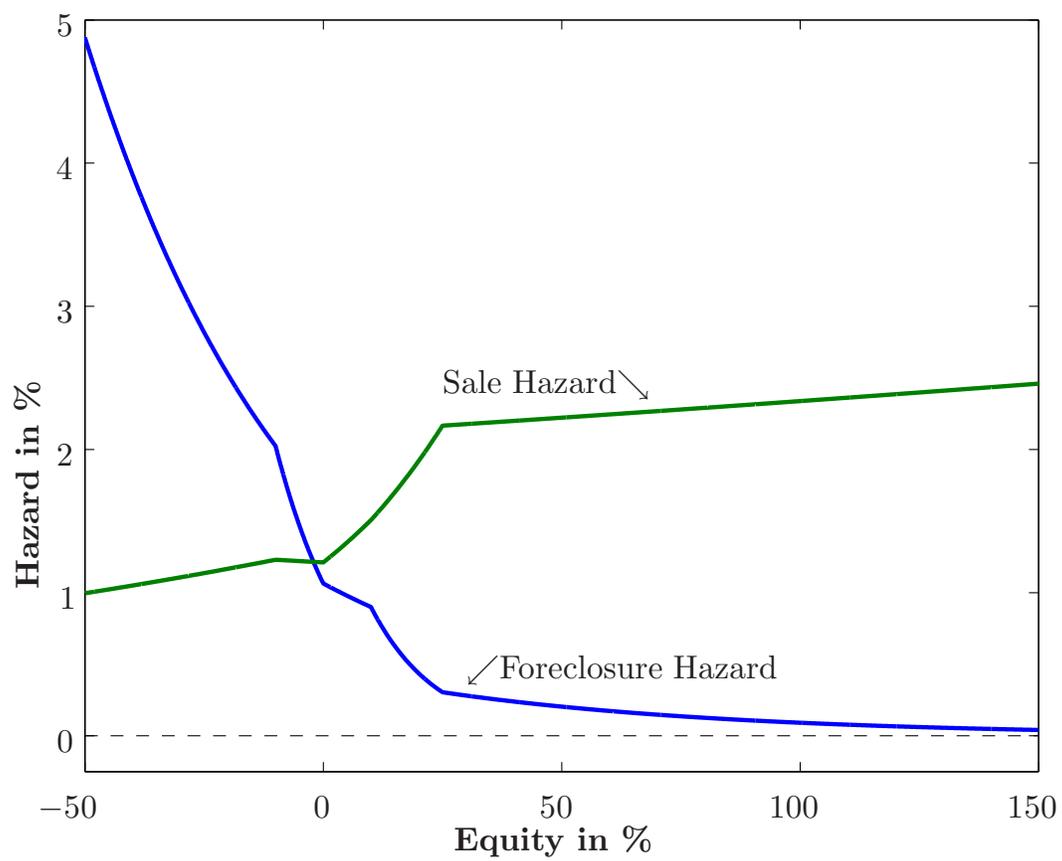


Figure 7: Estimated Effect of Equity on Foreclosure: Results from a Linear Probability Model

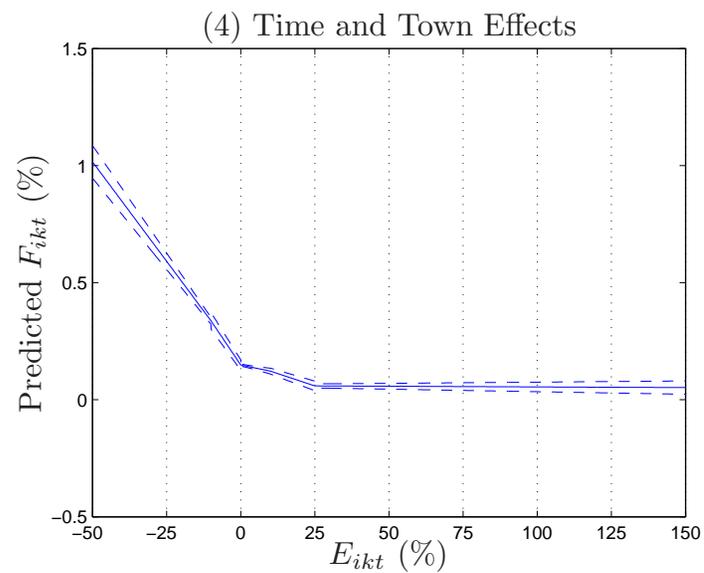
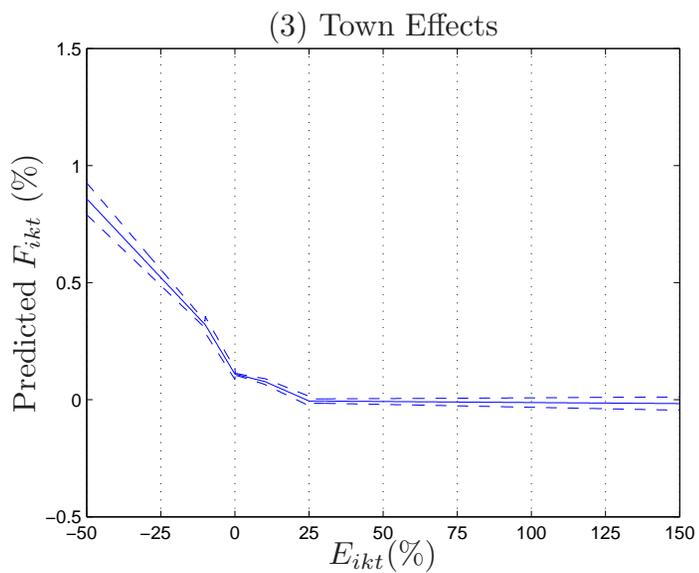
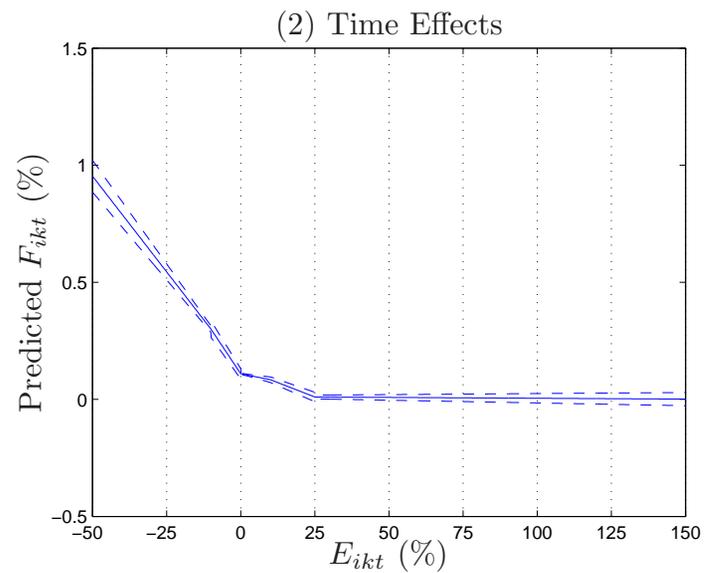
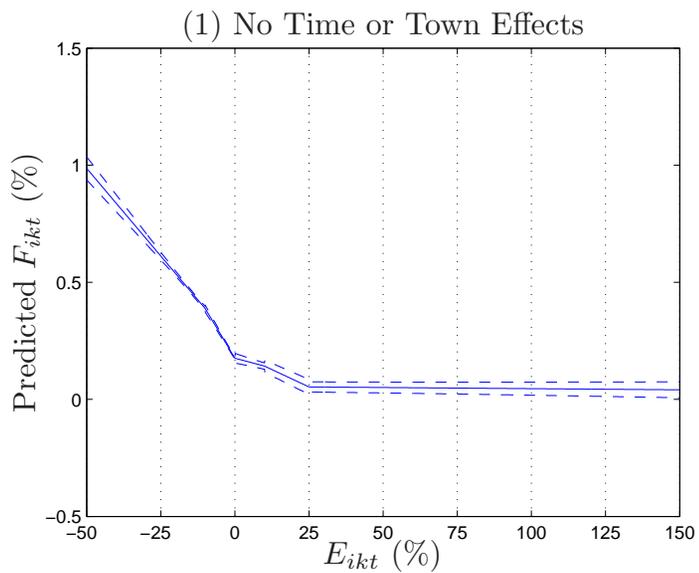


Figure 8: Predicted Hazard Rate versus Kaplan-Meier Hazard Rate (1987-2008)

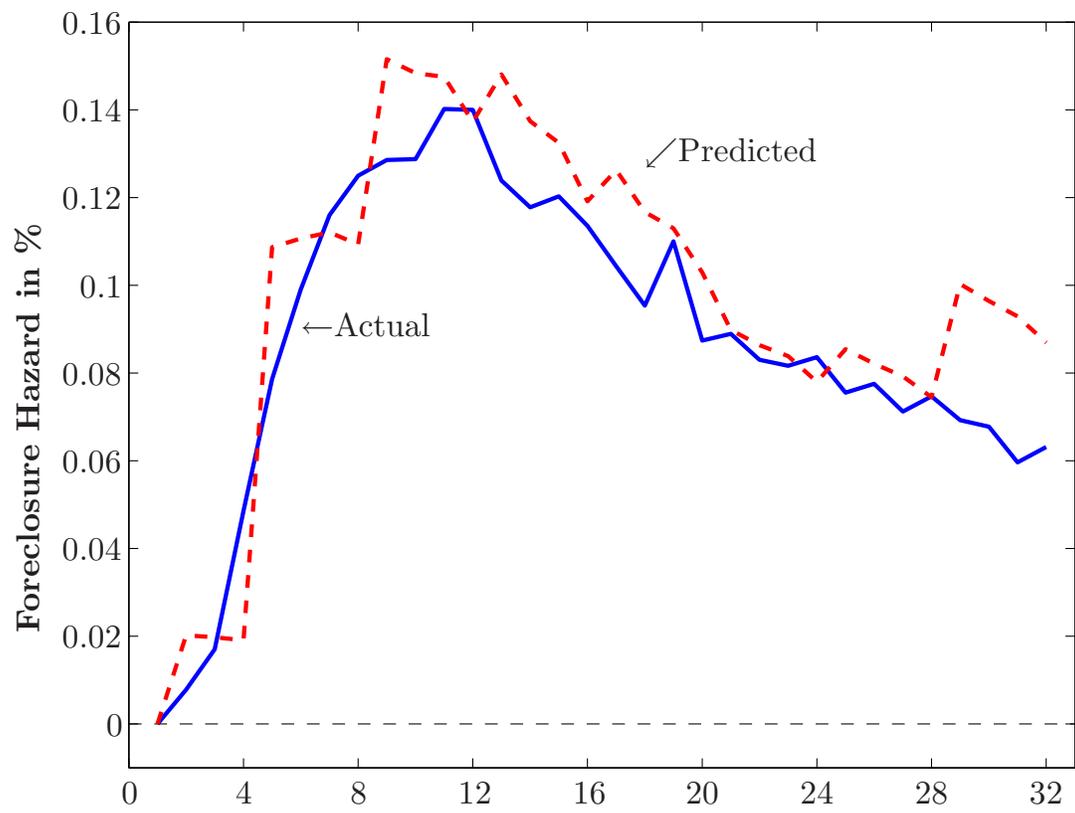


Figure 9: Predicted versus Actual Default Hazard Rates

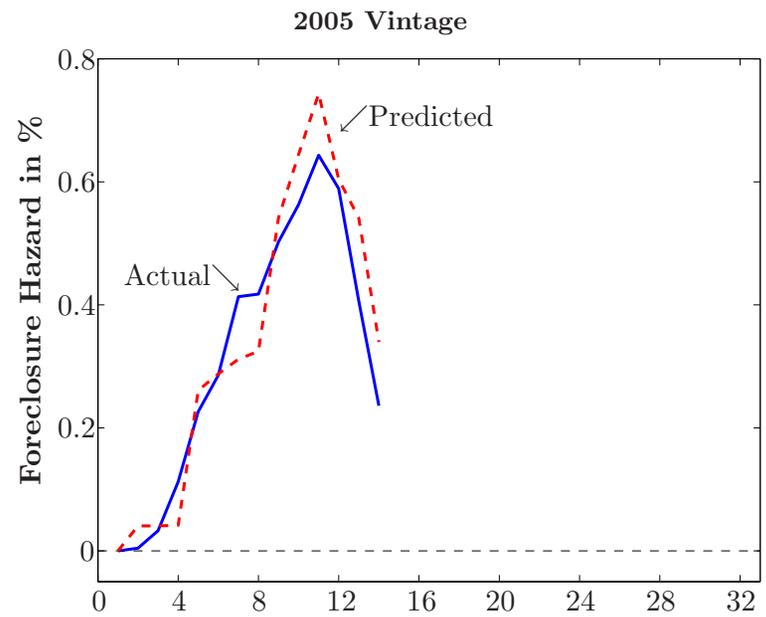
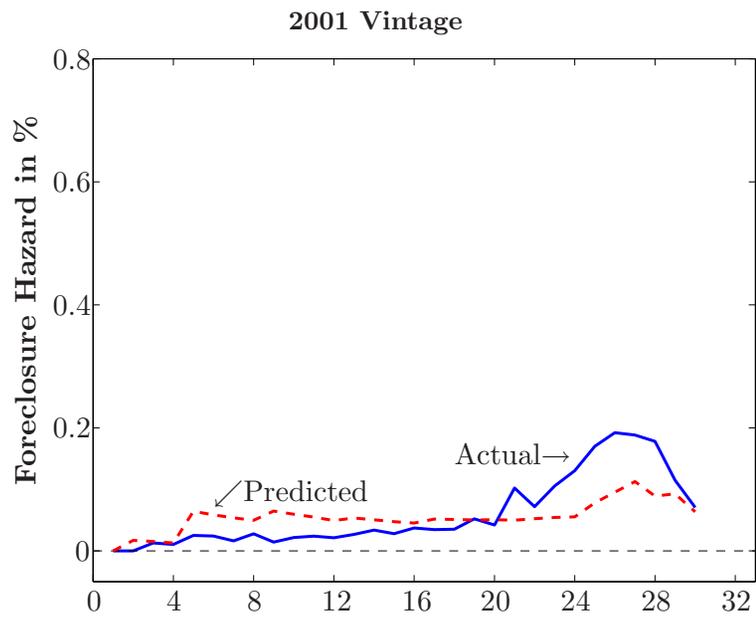
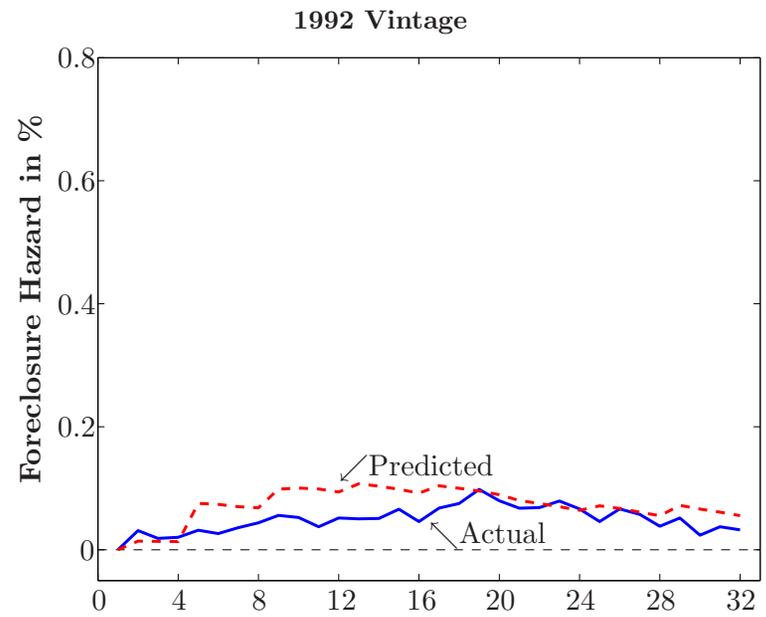
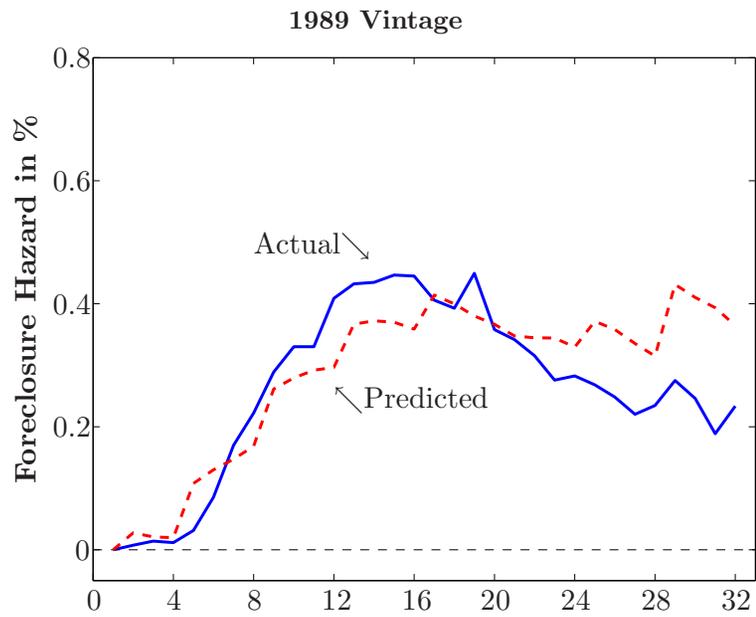


Figure 10: Predicted vs. Actual Massachusetts Defaults (1993–2008)

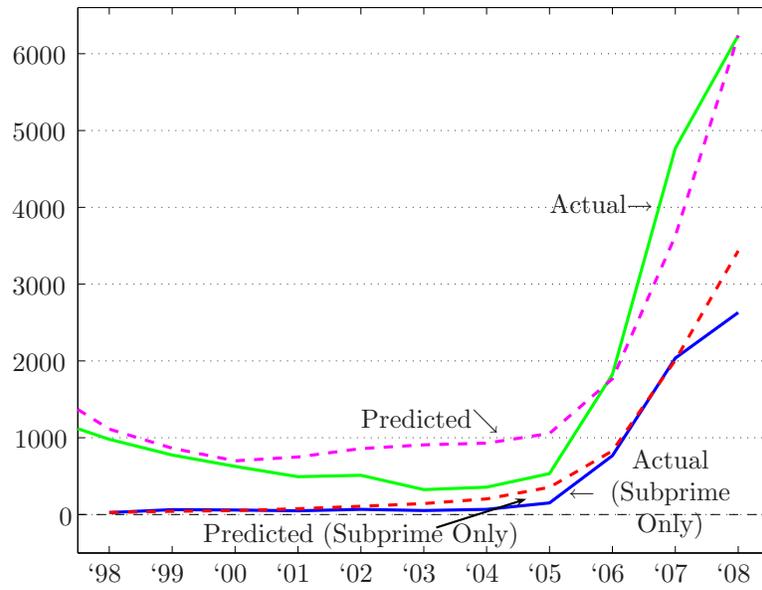


Figure 11: counter-factual Analysis

