Working Papers

FEDERAL RESERVE BANK OF BOSTON™

No. 17-13

Mortgage-Default Research and the Recent Foreclosure Crisis

Christopher L. Foote and Paul S. Willen

Abstract:
This paper reviews recent research on mortgage default, focusing on the relationship of this research to the recent foreclosure crisis. Research on defaults was advanced both theoretically and empirically by the time the crisis began, but economists have moved the frontier further by improving data sources, building dynamic optimizing models of default, and explicitly addressing reverse causality between rising foreclosures and falling house prices. Mortgage defaults were also a key component of early research that pointed to subprime and other privately securitized mortgages as fundamental drivers of the housing boom, although this research has been criticized recently. Going forward, improvements to data and models will allow researchers to explore the central unsolved question in this area: why mortgage default is so rare, even for households with high levels of negative equity or financial distress.

JEL Classifications: D12, G21, R21, R31
Keywords: Mortgage default, foreclosures, housing boom and bust, financial crisis

---

Christopher L. Foote and Paul S. Willen are senior economists and policy advisors in the Research Department of the Federal Reserve Bank of Boston. Their email addresses are Chris.Foote@bos.frb.org and Paul.Willen@bos.frb.org, respectively

This paper, which may be revised, is available on the website of the Federal Reserve Bank of Boston at http://www.bostonfed.org/economic/wp/index.htm.

This paper was prepared for the 2018 Annual Review of Financial Economics.

We thank Kris Gerardi, Lauren Lambie-Hanson, and Lara Loewenstein for many helpful comments and suggestions and Lawrence Bean, Lauren Lambie-Hanson, and Suzanne Lorant for editorial assistance.

This version: October 2017

1 Introduction

The foreclosure crisis that plagued the United States in the decade after 2005 put mortgage default on the research agendas of the world policymakers. But in the decades before the crisis, real estate economists and mortgage-finance professionals had been studying mortgage default, using advanced theoretical approaches and frontier-level econometric models. To measure their progress, we can go back to 1970, when the National Bureau of Economic Research published a slim monograph titled *Home Mortgage Delinquency and Foreclosure* (Herzog and Earley 1970), a volume that at the time was justly considered to be the definitive work in this area. The authors did not use a formal model of default; the Black-Scholes model, on which modern default theory is based, was still a few years away. But they did have a snapshot of nearly 13,000 mortgage loans on which they ran no fewer than 13 multiple regressions, an impressive feat given the technological limitations of the time. The authors concluded, among other things, that the presence of second liens and higher loan-to-value (LTV) ratios increased the likelihood of delinquency and foreclosure, whereas high debt-to-income ratios did not predict defaults very well.

In the 35-odd years between that publication and the start of the foreclosure crisis, a great deal about mortgage default was learned. On the theoretical side, arbitrage-based models shed light on the link between default and the time-series behavior of interest rates and house prices. Empirical researchers gained access to tens of millions of complete loan histories, on which they applied survival analysis to deal with problems of censored data and the competing risks of prepayment and default. By the beginning of the millennium, scholars viewed the use of arbitrage-based models and loan-level data not as the frontier but rather the standard for serious research.[1]

The foreclosure crisis caused interest in mortgage default to spread from real estate and urban economists to the economics profession at large, with the growth in research activity roughly keeping pace with the expanding stock of delinquent loans. This paper reviews how previous and contemporary default research informed policy discussions during the foreclosure crisis—or, in some cases, how this research *should have* informed those discussions. It evaluates default behavior during the crisis in light of previous findings, and it discusses the new research questions and opportunities that the crisis has opened up. Indeed, although research on default was advanced when the crisis began, we still have much to learn.

Like Herzog and Earley (1970), we focus on a "default function" that assigns a probability of default to a particular borrower with a particular house and loan. Investors can combine

[1] Vandell (1995) appeared somewhat before the millennium but represents an excellent recapitulation of the progress made since Herzog and Earley (1970). The *Econometrica* paper by Deng, Quigley, and Order (2000) is a good example of the degree of theoretical and technical sophistication present in empirical mortgage modelling at the start of the current century.

a default function with a house-price forecast to predict credit losses on a pool of loans. A default function also helps policymakers assess the effectiveness of a potential anti-foreclosure initiatives. Figure 1 shows the relationship between a borrower's equity and her probability of default implied by a prototypical default function. A researcher in 2005 would have found the downward slope in this function reasonable, because both theory and data predict that rising home equity makes a homeowner less likely to default. The researcher might have added that the curve shifts up for borrowers with low credit scores, those more likely to be unemployed, and those who took out their loans more recently.

The paper begins with a discussion of the formal theory of default that underpins functions such as that in Figure 1. In all the variants of this theory, negative equity is a necessary condition for default, for the simple reason that a borrower with positive equity can profitably sell her house to avoid default in the event she can no longer afford her monthly payments. Yet theory is less concrete about what she should do once equity becomes negative. On one hand, arbitrage-based models grounded in formal theories of household optimization imply that default depends only on aggregate factors such as house prices and interest rates, not on the individual characteristics or circumstances of the borrower. Such theories imply that, for example, credit scores should not shift the default curve in Figure 1. At the opposite extreme are so-called double-trigger models, which allow adverse life events such as job loss and illness to precipitate defaults.[2] Double-trigger models generate more-realistic empirical predictions, but they are less grounded in formal household optimization. As we will discuss, the foreclosure crisis has encouraged economists to blend these two extremes into a third alternative that provides an optimizing foundation for the default function and also formalizes the previously ad hoc approach of double-trigger models.

After laying out the theory in Section 2, we illustrate in Section 3 how lessons from theory were applied—or should have been applied—during the recent foreclosure crisis. Among other things, the section explains that once house prices fell, the foreclosure crisis proceeded pretty much as empirical mortgage researchers expected. The failure to predict the crisis stemmed from analysts' inability to foresee the massive fall in U.S. house prices in the late 2000s, which shifted the distribution of equity in Figure 1 to the left. In other words, the failure to anticipate the scope of the crisis rests largely on the science of asset pricing, not on estimates of the shape of the default function.[3]

Sections 4 and 5 examine the impact of the foreclosure crisis on current research. One

[2] The double-trigger models are so named because they contend that default arises from two triggers: negative equity and adverse life events. Without negative equity, a borrower suffering an adverse life event will simply sell the house if he can no longer afford the monthly payment. Yet without an adverse life event, borrowers with negative equity in double-trigger models are generally predicted to keep making their payments.

[3] For a discussion on how "reasonable people" disagreed about the rationality of house prices during the housing boom, see Gerardi, Foote, and Willen (2011).

silver lining of the crisis, as least as far as researchers are concerned, is that the large loan-level datasets traditionally used in default research are now more widely available, and new datasets have been developed. Default models have also been used to investigate new questions, such as the source of the steep increase in U.S. mortgage debt in the early 2000s (Demyanyk and Van Hemert 2011, Keys et al. 2010, Mian and Sufi 2009), although some of this work has been criticized recently. A pattern that consistently emerges in much empirical work is that default is surprisingly rare; people with very deep negative equity do not walk away from their homes as often as theory would predict, and even financially stressed borrowers display a strong aversion to default. Many explanations for abnormally low default rates have been suggested, and in the conclusion we speculate that a richer model of house-price expectations could be added to this list.

2 The Modern Theory of Mortgage Default

2.1 The Frictionless Option Model

Any discussion of default theory starts with the frictionless option model (FOM). This model is frictionless in the true sense of any standard asset pricing model, as there are no penalties for default (other than the loss of the house, of course) and no transactions costs of any kind. Additionally, as is typical in no-arbitrage models, a critical assumption is that households can take out unlimited unsecured loans at the market interest rate. The classic reference for this model is Epperson et al. (1985), and two equivalent formulations of the problem are linked by a put-call parity relationship.[4]

We start with the call-option formulation in which the mortgage contract involves the borrower's selling the house worth H to the lender and getting an option to repurchase the home by paying the lender an agreed upon sum of money M, the mortgage balance.[5] In the language of option theory, a borrower with negative equity is "out of the money," meaning that the borrower would lose money if he exercised the call option. No rational person would ever exercise an out-of-the-money option but as all finance textbooks point out, this does not mean that such an option is worthless. If there is any chance the option will someday be in the money, then the option has positive value. This situation will arise if future house prices rise by enough to restore positive equity. The default decision emerges because each month,

[4] Epperson et al. (1985) considers the narrow case of a borrower who can default but not prepay. Kau et al. (1992) extends the analysis to the more general situation where a borrower can also prepay, and Kau, Keenan, and Kim (1994) and Kau and Keenan (1999) conduct numerical simulations of the model.

[5] The put-option formulation may seem more realistic, but the the call-option formulation is actually very close to the original common law contract on which U.S. mortgages are based, and in most states, a mortgage contract still transfers legal title from the borrower to the lender. See the discussions in Gerardi, Lambie-Hanson, and Willen (2013); Kimball and Willen (2012); and Knoll (2002).

the borrower must decide whether to renew the option by making the monthly mortgage payment. The question for the borrower is whether the value of preserving the call option exceeds the monthly payment, in which case the borrower continues paying. Conversely, if the value of the call option falls short of the required payment, the borrower defaults.

In the alternative put-option formulation, the borrower has three assets: a house, a loan, and a put option that allows the borrower to sell the house to the lender at a strike price equal to the outstanding balance on the mortgage. Default consists of exercising the put option to sell the house and using the proceeds to pay off the loan. For illustration, consider a negative-equity borrower with a house worth H and a mortgage balance of $M > H$. Because the put option P gives the borrower the right to sell a house worth H for price M, it must be worth at least $P = M - H$, but it is typically worth more than $M - H$, because the distribution of future valuations of the option is asymmetric. If H falls, the option becomes more valuable, but if H rises, the value of the option can never fall below zero.

Table 1 shows the borrower's balance sheet before and after default. It surprises many that homeownership never makes a negative contribution to the borrower's balance sheet, even when equity is negative. Put another way, default can never *increase* a borrower's wealth. Does this mean that default is never optimal? No, because we must remember the monthly payment. By exercising the option to default, the borrower evades the monthly payment currently due on the loan. Consequently, default is optimal if the (positive) contribution of the house to the borrower's balance sheet is smaller than the (positive) monthly payment. This condition is exactly equivalent to the call-option default condition discussed above, as put-call parity ensures that the balance sheet benefit of homeownership equals the value of the call option.

It may seem counterintuitive that a home with negative equity makes a positive contribution to the household balance sheet, but this implication follows directly from the first principles of option theory. Basically, the property is an asymmetric bet. If house prices fall further, the increase in the value of the put option offsets part of the reduction in the value of the house. Yet house prices can also rise, and if they rise enough, positive equity will be restored and the borrower will be able to sell the house at a profit. In other words, in no state of the world will the homeowner have to pay money, and in some states of the world he may receive money.

Figure 2 shows how the value of the default option and the borrower's net position change with the price of the house, thereby illustrating why even borrowers with some negative equity will continue making payments. To construct this figure, we used a simplified mortgage contract set out in Kau and Keenan (1999), in which the mortgage balance never declines (that is, there is no amortization) and borrowers cannot prepay their mortgages. Each period, the borrower makes a payment of a percent of the constant unpaid principal balance M. The borrower also receives a per-period service flow of s percent of the current value of the house H; this flow can be interpreted as the per-period rents generated by the house.[6] The interest rate is fixed at r, and house prices are assumed to follow a standard stochastic process with variance σ^2. For the figure, we assume that M equals \$100,000, $a = s = r = 0.04$, and $\sigma^2 = .02$.

As made clear in Epperson et al. (1985), we can use this type of model to define the *value of the mortgage* V as the present discounted value of all promised payments, including the current monthly payment.[7] The mortgage value V recognizes that in some states of the world, the borrower will default and exercise the option, and default occurs when the value of the mortgage V meets or exceeds the value of the house H. The default condition $V \geq H$, however, is exactly equivalent to the condition that the monthly payment exceeds the value of the call option and the condition that the monthly payment exceeds the value of the house plus the put option less the debt. In other words, the $V > H$ formulation is simply another formulation of the FOM. Indeed, the figure shows that we can read off the value of the put option P as the difference between the mortgage balance M (dashed line) and the value of the mortgage V (solid blue line).

This option value of default is small for large values of positive equity $H > M$, because the option to default has little value when there is little chance it will be exercised. As the value of of H declines, however, the value of the default option grows. At the boundary of negative and positive equity ($H = M$), the *only* contribution that the house makes to the borrower's net wealth $H - M + P$ is this option. The value of the option continues to grow as prices continue to decline, but eventually, as we move leftward along the horizontal axis, the house-price decline "catches up" with the growing value of the option P and the declining value of the mortgage V, and the house no longer makes a positive contribution to the borrower's wealth. At this default threshold, the house price has fallen by enough so that the restoration of positive equity through rising prices is sufficiently unlikely, given the volatility of prices σ^2, so $V = H$ and the borrower defaults.

Several predictions of the FOM are relevant for the discussion that follows. First, the FOM specifies a small number of factors on which default depends. Figure 3 shows the sensitivity of the default threshold to changes in the model's fundamental parameters. The baseline values for these parameters (denoted by the dots) are identical to the values that generate the previous figure with one exception: the monthly payment fraction a is raised by 1 percentage point, to 0.05, so that it exceeds r and thus makes the baseline default

[6] If the borrower is living in the house, then the service flow is the rent that the borrower pays himself, or, equivalently, the rent he does not have to pay someone else.

[7] By assuming that $a = r$ for the figure, we are also assuming that the present discounted value of the perpetual payment stream $\frac{aM}{r}$ equals M. This is an expositional simplification and will be relaxed below. The particular model we use to generate the figure is found in equations 22 and 23 of Kau and Keenan (1999).

threshold more realistic.[8] The top left panel shows that the default threshold declines with higher levels of house-price volatility σ^2. As with any option, the default option becomes more valuable given higher volatility in reference prices, so the default threshold declines as volatility increases.

The next two panels show how service flows (such as rents) and payment sizes affect the default threshold. Intuitively, higher service flows raise the value of the house to the homeowner, but in this formulation, higher rents are modeled equivalently as reducing the value of the mortgage V. Regardless of the way they are modeled, higher rents make continued homeownership more attractive, so they reduce the equity threshold at which borrowers default. In a similar vein, the third panel shows that higher monthly payments raise the default threshold, because they make continued ownership less attractive.

Returning to our earlier discussion of the shape of the default function, what does the FOM imply about this shape? Strictly speaking, the FOM implies a unique threshold level of equity such that everyone defaults, as shown by the line labeled "FOM" in Figure 4. The stark prediction in Figure 4 would seem to be easy to test (and easy to reject), but valuing individual homes is difficult. As a result, uncertainty about house values smooths out the measured default function, even if the true function looks like the one in Figure 4. Aragon et al. (2010) argue that the standard deviation of pricing errors using standard repeat-sales indices is more than 20 percent, and Figure 5 shows that pricing errors of this magnitude can smooth out the default function significantly. Even a 10 percent pricing error has a big effect on the shape of the curve.

2.2 Empirical Problems with the FOM

A problem for the FOM that is larger than its stark prediction for the shape of the default function is that it allows only a small set of factors to shift that function up or down. The FOM is a classic no-arbitrage model in the tradition of Black-Scholes, and one of the virtues of such models is that they do not require knowledge of individual characteristics of investors, who are in our case the mortgage borrowers. Whether the borrower is rich or poor, whether she has just lost her job or is newly divorced, or whether she has recently started a new job two time zones away has no impact on her default decision. In the FOM, any link between individual circumstances and default decisions is severed by the assumption that borrowers can take out unlimited unsecured loans. Consider a borrower who has few liquid resources because of, say, a recent job loss. If the FOM predicts that he should keep paying his mortgage, then he can simply borrow the money at the market rate to do so.

[8] Note that the baseline default threshold implied in the panels of Figure 3 is larger than the default threshold in the previous figure. This is intuitive and confirmed by the bottom left panel of Figure 3: borrowers reach the default threshold sooner if they must pay a higher payment each month.

In the real world, however, borrowers typically attribute their defaults to idiosyncratic shocks. Table 2, taken from Cutts and Merrill (2008), shows that more than 40 percent of a recent sample of delinquent Freddie Mac borrowers cited loss of income as a reason for delinquency. Other hardship-inducing shocks include illness and marital difficulties. And it's not just what borrowers say—empirical default equations reveal that individual characteristics help explain borrower behavior as well. Figure 6 shows the results of a standard competing-risks hazard model of the performance of a sample of Alt-A mortgages from Fuster and Willen (2015).[9] Some of these results mesh well with the FOM. For example, the top panel shows that there is a negative relationship between default and equity, and as noted earlier, mismeasurement of the value of the homes can explain why the curve does not jump from zero to 100 percent at a specific point. The top panel also shows that the default hazard decreases with declines in payment size, which are determined in this study by contract interest rates. This is again what the FOM predicts.

The problems with the FOM all appear in the bottom panel of Figure 6, because the FOM implies that most of the variables listed there should not matter. For example, the coefficient on credit score should be zero, but the panel implies that a 100-point increase in FICO score cuts the probability of default in half. Of course, the FICO score could proxy for some other variable that does matter in the FOM, but such a variable would have to stand in for something about the loan or the property, not something about the individual borrower.[10] One possibility is that low FICO scores are correlated with errors in our measures of the value of the property, or that borrowers with higher FICO scores have higher monthly payments.

In general, however, it is hard to see how any channel could cause the FICO score to have such a dramatic effect on the probability of default. For example, differences in house-price expectations may lie behind some of the individual-specific coefficients in the default regressions. Perhaps negative-equity borrowers disagree over the probability that house prices will recover enough to restore positive equity. But like all no-arbitrage models, the FOM places tight restrictions on how the stochastic house-price process relates to other fundamental variables in the model. For example, expected price growth is essentially pinned down by the interest rate and the rental service flow. If house prices were expected to rise sharply and interest rates were low, then expected returns would be abnormally high for people who bought houses with borrowed money, in violation of the no-arbitrage condition.[11]

[9] During the recent housing boom, Alt-A mortgages were typically taken out by borrowers with good credit scores who did not want to document their incomes. See Adelson (2003) for details.

[10] For example, borrowers with low FICO scores might be more likely to be unemployed, but that would not affect the probability of default in the FOM.

[11] In the Kau and Keenan (1999) model the relevant no-arbitrage condition is $\frac{\mu_H - \lambda_H \sigma}{H} + s = r$, where μ_H is expected price growth, and λ_H is a risk-adjustment term. As before, H and s are the values and service flows from the house, respectively, and σ^2 is house-price volatility. This condition implies that the risk-adjusted

Even more generally, the fundamental theorem of asset pricing implies that the risk-adjusted path of expected house prices should be the same for all houses, so one cannot explain any of the coefficients in the lower panel with unobserved differences in expected house-price appreciation. A more promising argument would be that the variables in question are correlated with unobserved differences in the *variance* of prices, which we denoted σ^2 above.[12] The problem here is that the sign goes the wrong way. One would imagine that house price volatility would be higher in neighborhoods populated by people more likely to be unemployed or to have low credit scores, yet higher volatility would cause the FOM to predict a *lower* probability of default for those borrowers.

Finally, households may differ in their default behavior because of differences in their transactions costs of default. Moving incurs both time and financial costs, and a foreclosure can adversely affect future chances of home purchase and employment. There may also be a psychological stigma attached to default that some households would prefer not to incur. Foster and Van Order (1984) suggested transactions costs as a way to improve the empirical performance of one of the earliest option-based models of default, and Downing, Stanton, and Wallace (2005) show how heterogeneous transactions costs can smooth out the FOM's default function, as in Figure 5. A stigma cost that is constant over time for an individual borrower can explain the strong effect of credit scores on default, because individuals with few defaults on their previous debts would be expected to default less often on their mortgages, too. But time-invariant costs do not explain why so many defaults are caused by high-frequency shocks, such as unemployment, illness, and divorce.

2.3 The Double-Trigger Model

To address the empirical failings of the FOM, researchers have turned to the so-called double-trigger model, which attributes default to the combination of negative equity and an idiosyncratic adverse shock.[13] The basic logic is that borrowers suffer adverse life events all the time, but normally, borrowers have positive equity, so these shocks translate into sales and refinances, not defaults. Only when the second trigger, negative equity, is also pulled do shocks lead to default. The top panel of Figure 7 depicts the two basic features of default in the double-trigger model. First, the figure explains why we see virtually no defaults when house prices have been rising so that most borrowers have positive equity. Second, the model provides a channel through which credit scores and other idiosyncratic shocks can matter for

return to owning a home (that is, price appreciation plus the service flow) is equal to the return to investing in a bond r. This condition also replicates the user-cost equation in standard housing models.

[12] As we saw in Figure 3, the value of the call option is increasing in this variance because of the higher probability that the call will end up in the money.

[13] The use of "double-trigger" in the academic literature goes back at least as far as Goldberg and Capone, Jr. (1998).

default. If borrowers with low credit scores are more likely to, say, lose their jobs or suffer other setbacks, then the double-trigger model predicts that these borrowers will also be more likely to default.

Yet in its purest form, the double-trigger model has a significant limitation. The top panel of Figure 7 implies that a borrower who loses a job and owes 5 percent more than the value of his house is just as likely to default as an unemployed borrower who owes 50 percent more. Similarly, a negative-equity owner who suffers a 20 percent decline in income is just as likely to default as an owner whose income stream completely dries up. The bottom panel of Figure 7 depicts a modified double-trigger model that places life events and negative equity on continua. Positive equity still inoculates borrowers against default, but default among negative-equity owners is now driven by an interaction between the extent of negative equity and the severity of the adverse life event. Borrowers facing larger income shocks default at lower levels of negative equity, and the threshold income shock is more severe for a borrower with more equity. Yet even this modified model is an ad hoc formulation that lacks foundations in optimizing behavior. If a negative-equity owner suffering an adverse life event believes that house prices could recover, then why doesn't he simply borrow the money to make his monthly payment?

2.4 Integrating FOM and Double-Trigger: Two-Period Models

Some intuition for how adverse life events matter for default is provided by two-period models that limit homeowners' unsecured borrowing opportunities. The first such model, in Gerardi, Shapiro, and Willen (2008), relaxes the FOM assumption that households can borrow unlimited amounts at the riskless rate. Instead, borrowers must pay a higher rate on borrowing than they receive from saving. Consistent with the FOM, the model interprets the monthly payment as an investment in an asset that pays off if the option is in the money in the future. The authors then show that one can redefine the optimal default condition in terms of the asset's expected payoff r_M: the borrower pays if r_M exceeds the borrower's opportunity cost of funds and defaults otherwise. Households with abundant cash-on-hand can essentially borrow from themselves, and their low implicit cost of funds makes them less likely to default than liquidity-constrained households. A similar approach is taken in Foote, Gerardi, and Willen (2008), although the two-period model in that paper is based on the formulation of the FOM that compares the value of the house to the value of the mortgage. Consistent with the FOM, the option to default reduces the value of the mortgage, and negative-equity borrowers stay in their homes if the possibility of future positive equity is sufficiently high. But the model also shows that because liquidity-constrained borrowers have higher discount rates, these borrowers value future house-price gains less than unconstrained borrowers do. As a result, future price gains are less of a carrot for the constrained borrowers to keep making their payments, so they default more often. Both of these two-period approaches illustrate why adverse life events that wipe out cash-on-hand can lead to default, so both can reproduce the qualitative predictions of the modified double-trigger model. Yet because these models are highly stylized and have only two periods, they cannot generate quantitative predictions.

2.5 Integrating FOM and Double-Trigger: Life-Cycle Models

Other combinations of the FOM with the double-trigger model take a life-style perspective and are thus more suitable for specific analysis. Campbell and Cocco (2015), Laufer (2017), Corradin (2014), and Schelkle (2014) each follow a common formula: a dynamic life-cycle model in which borrowers earn uncertain labor income and face limits on unsecured borrowing. Campbell and Cocco (2015), Corradin (2014), and Schelkle (2014) model the decision to default on an individual mortgage. Laufer (2017), by contrast, ignores the individual characteristics of each mortgage and instead focuses on the entire homeownership experience from purchase to default or sale, paying special attention to the additions to debt arising from second mortgages and refinances. All four models can generate the basic contours of the modified double-trigger model. Negative equity is a necessary condition for default, but the incidence of default is increasing in the intensity of the borrower's income shock and decreasing in her level of wealth.

Laufer (2017) includes three different idiosyncratic shocks as "second triggers" that lead to default. Borrowers face a permanent income shock, an unemployment shock that reduces their income only temporarily, and a "preference shock." The last of these is a random increase in wealth that the borrower gets from moving and that induces sale or default depending on the borrower's equity position. Laufer finds that the preference shocks are the most important second trigger, although it is unclear why or, in a sense, which of the theoretical shocks maps most closely to the real shock of income loss and employment disruption. Corradin (2014) emphasizes how the initial leverage position of borrowers is affected by the volatility of house prices and the likelihood that a borrower will undergo a lengthy spell of unemployment. Imperfect financial markets prevent borrowers from smoothing over income disruptions, so lenders seek to insure themselves against default with down payments that rise when borrower-income risk or house-price volatility increases.

Schelkle (2014) sets up his discussion of his dynamic optimizing model by explicitly considering both the frictionless model of Section 2 and the simple double-trigger model from the top panel of Figure 7. Specifically, he uses generalized method of moments (GMM) to estimate a critical negative equity threshold in the frictionless model and a probability of life event in the double-trigger model. He then argues that the estimated double-trigger model performs far better than the frictionless model, although neither performs as well as the fully optimizing model. The paper does a good job of showing the differences between the two basic models described above, as well as how these models can be combined into a richer optimizing framework.

Indeed, one shortcoming of some dynamic models is that they fail to draw a clear link to the FOM benchmark. On the page it is hard to see the link between, say, Campbell and Cocco (2015) and Epperson et al. (1985), but both papers solve the problem of the optimal exercise of an option, although the former imposes important limitations on the investor's potential portfolio strategies. Without those limitations, the solution of any optimizing model should converge to the FOM.

3 Mortgage-Default Theory in the Crisis

3.1 Was This Crisis Different?

Many of the basic results in the previous section were developed in the 1980s and 1990s by real estate economists and mortgage-finance professionals, so much was already known about default before the recent foreclosure crisis began. In this section, we discuss how this understanding was (or was not) applied to specific questions of interest during the crisis. Perhaps the most basic question early on was whether the huge extent of the crisis would cause its effects to be different from anything economists had seen before. Many predicted it would; in a recent column on the evaluation of macroeconomic risk, Jeffrey Frankel writes that it was common to view the national housing crash as an unprecedented event:

> When the housing market did crash, it was regarded as a surprise. The crash lay outside any standard probability distribution that could have been estimated from past data, analysts declared, and was therefore a black swan event, or a case of "Knightian uncertainty," radical uncertainty, or unknown unknowns. After all, the analysts argued, [national] housing prices had never fallen in nominal terms before. But, while nominal housing prices had not fallen in the United States in the previous 70 years, they had fallen in Japan in the 1990s and in the United States in the 1930s (Frankel 2017; insertion added).

Frankel might have added that prices had also fallen even more recently in specific regions of the United States. In the late 1980s and early 1990s, house prices in coastal states such as California, New York, and Massachusetts underwent a boom-bust cycle that left many owners in those states with negative equity only a decade or so before the national housing downturn started. A similar cycle had occurred in the mid-1980s in oil-patch states such as Louisiana and Oklahoma. These price declines provided mortgage researchers with the variation they needed to estimate the effect of negative equity on default, and as the national crisis began, these researchers predicted that defaults would rise. But the researchers also knew that the large majority of homeowners with negative equity would not default—and not simply because these owners were being nice to their lenders. As explained in the previous section, even the most hard-headed homeowners with negative equity often find it optimal to keep making payments because of the potential for house prices to recover—which in the case of the regional crises they did within a few years.[14] Figure 8 depicts house prices for two coastal states, California and Massachusetts. Both states experienced declines in house prices during the early and mid-1990s. But both states had also recovered their previous peaks by the turn of the century.

Figure 8 also shows that California and Massachusetts experienced deeper price declines during the national housing cycle of the 2000s, with prices taking somewhat longer to recover as well.[15] Yet for Massachusetts at least, the behavior of negative-equity owners was similar across the two cycles, in that few of them defaulted. Table 9 updates some estimates from Foote, Gerardi, and Willen (2008), who pointed out the low default rate in Massachusetts during the 1990s. As seen in Figure 8, prices peaked in the Bay State in early 2006, fell through 2012, and recovered thereafter. Accordingly, the top panel of Figure 9 shows a significant number of Massachusetts homeowners with mild amounts of negative equity in 2007. As prices continue falling thereafter, the number of Bay State homeowners with more severe negative equity grows. However, the bottom panel shows that for the most part, fewer than 10 percent of homeowners with any amount of negative equity in any recent year wound up losing their homes. These low default rates mimic those among Massachusetts homeowners during the 1990s.

One feature of the housing landscape that had changed between the two housing cycles was the higher prevalence of subprime mortgages, which defaulted in large numbers at the start of the housing bust and helped precipitate the subsequent financial crisis. Even here, however, past performance was instructive. In the 1990s, subprime mortgages were generally "hard money" loans that were extended to current homeowners with low credit scores and backed by existing home equity. Over time, as house prices rose in the late 1990s and early 2000s, subprime loans were increasingly extended for home purchases. Using default equations, lenders could model how many defaults could be expected from subprime borrowers, conditional on a particular scenario for house prices. These predictions turned out to be strikingly accurate when house prices fell. Gerardi et al. (2008) point to an August 2005 analysis by Lehman Brothers that forecast if house prices fell by 5 percent per year

[14] As we discuss below, homeowners with deep negative equity tend to default even less often than models like the FOM predict. While this fact is hard to explain theoretically, it is not hard to see in the data.

[15] Indeed, California had yet to recover its mid-2000s peak completely by 2017. California is one of the "sand states" that experienced the most significant housing cycles during the 2000s; below we present price data for the other three sand states (Arizona, Florida, and Nevada) as well as for the nation as a whole.

for the next three years—what the analysis labeled the "meltdown" scenario for the U.S. housing market—then cumulative losses on pools of securitized subprime mortgages would be a little more than 17 percent. This is a large number. Using a standard recovery rate of 50 percent for each defaulted mortgage, a loss of 17 percent on a pool of loans implies that about one-third of the mortgages in the pool will default.

One index of subprime prices, the 2006-1 ABX, traces the performance of subprime loan pools originated at the end of 2005, so it allows us to evaluate the Lehman forecast. Table 3 shows predictions of cumulative losses on deals in this index, according to J.P. Morgan. The earliest forecast, from July 2008, is a little more than 17 percent, making the Lehman analysts appear preternaturally accurate. However, the decline in actual house prices turned out to be even larger than the 5 percent decline imagined in Lehman's meltdown scenario. As prices kept falling after 2008, J.P. Morgan kept revising its loss forecast, until the forecast stabilized in 2010 at about 23 percent, where it has remained ever since.[16] The bottom line of this analysis is that subprime lenders were under no delusion that these loans would continue to perform if house prices fell. Past experience had taught them that loans extended to borrowers with low credit scores were highly sensitive to house prices.

3.2 Unemployment and Default at the Aggregate Level

Another set of issues in the recent crisis concerned the relationship between unemployment and default. Defaults started rising in 2006, but aggregate unemployment did rise much until 2008, and this pattern led some analysis to contend that unemployment was not an important driver of default at the individual level.[17] The tenuous relationship between unemployment and default ostensibly indicated that defaults were best prevented by permanent principal reduction, as opposed to temporary relief targeted to unemployed borrowers.

To understand how individual-level unemployment can generate default even when the aggregate unemployment rate is constant, consider again the double-trigger model in the top panel of Figure 7. In 2005, house prices had been rising rapidly for years, so most borrowers had positive equity and were thus located in the right-hand column of the figure. Consequently, even borrowers who did suffer income shocks could avoid default by selling or refinancing their homes. And to be clear, even though the unemployment rate was stable before 2008, millions of people suffered income shocks during this period. Figure 10 shows that around 1.75 million people flowed from employment into unemployment *each month* throughout 2005, 2006, and 2007, according to the Current Population Survey. In 2006,

[16]The 23 percent estimate is guaranteed to be close to the final loss amount, because few loans covered by the 2006-1 ABX remain active (that is, most loans have either prepaid or defaulted).

[17]Goodman et al. (2010, p. 67) state that "default transition rates picked up long before unemployment picked up—thus unemployment did not 'cause' defaults." Mian and Sufi (2010) and Mian (2010) make similar arguments.

house prices started falling, so anyone suffering a subsequent income shock was likely to have been pushed into negative-equity territory, where income shocks make default likely. Figure 10 also shows that foreclosure starts are small relative to employment-to-unemployment flows. Consequently, even a small change in the share of job losers who default on their mortgages would cause foreclosure numbers to skyrocket.

City-level data provides some more direct evidence that unemployment shocks are important drivers of default. Table 4, taken from Goodman et al. (2010), depicts default transition rates for borrowers binned by combined loan-to-value (CLTV) and unemployment rates at the level of the metropolitan statistical area (MSA). Borrowers in the right-most column of the table have abundant positive equity, so they correspond to the area to the right of the zero-equity line in the bottom panel of Figure 7. Consistent with the modified double-trigger model, the city-level unemployment rate has little effect on default for these borrowers; the default rate for borrowers in the lowest jobless bin (0.24 percent) is essentially identical to the rate of those in the highest bin (0.23 percent). Occupant owners with CLTVs above 100 have negative equity, so they are located on the left side of the modified double-trigger map. Data matches theory for these borrowers as well. Among borrowers with CLTV ratios above 120, the default rate for borrowers in the highest unemployment bin is 2.21 percent, while the rate from the lowest bin is only 0.86 percent. The model predicts income shocks would matter less for non-occupant owners, and a similar comparison among the cells of the lower panel of Table 4 bears out this prediction as well.

Although the match between Table 4 and the modified double-trigger model is strong, the need to use area-level unemployment data for both this table and for empirical default regressions means that the jury is still out on just how strongly unemployment is linked to default at the individual level. In a geographic area, some unknown third factor might cause both a high default rate and a high unemployment rate, but the people who default may not be the people who are unemployed. Consider a city with a lot of new construction and a great deal of housing speculation and house-price appreciation. A subsequent housing bust would simultaneously lead to large numbers of mortgage defaults (because of falling house prices and rising negative equity) and high levels of unemployment (because of the decline in construction). An even more basic problem with using aggregate unemployment rates in default regressions is that doing so leads to massive attenuation bias in the resulting coefficient. The simulation evidence in Gyourko and Tracy (2014) shows that using the area-level unemployment rates as a proxy for an individual resident's unemployment experience understates the true effect of unemployment on default by a factor of more than 100.

To make progress in understanding the role of unemployment in default, researchers will need to develop datasets that match labor market experiences and default behavior at the individual level. One of the few datasets that currently does so is the Panel Study of Income

Dynamics (PSID), and we discuss some results that use the PSID below. Unfortunately, the size of the PSID, its two-year frequency of observation, and the limited housing information it contains make it less than ideal for studying unemployment and default. One could improve upon the PSID by combining the Census's Longitudinal Employer-Household Dynamics dataset (LEHD) with either public property-deed records or with a loan-level dataset of the type that has long been used to estimate default regressions. The resulting dataset could then track the timing of both job loss and delinquency at the individual level.

3.3 Policies to Prevent Foreclosures

As the number of foreclosures began to rise, policymakers naturally looked for a way to reduce them. Devising such a policy would be easy if all borrowers ignored the option value of default and stopped making payments the moment equity became negative. In this scenario, a mass principal reduction that restored each borrower to positive equity would reduce defaults dramatically.[18] Indeed, if all borrowers followed this naive default strategy, the size of the required monthly payment would have no effect on the default decision, so trying to reduce foreclosures by reducing payments would be pointless.

Because both theory and data refute this naive characterization of borrower behavior, crafting an effective anti-foreclosure policy is complicated. As discussed above, during previous housing busts most negative-equity borrowers did not default, so after 2007 most negative-equity borrowers were expected to continue making payments with or without any mortgage relief. Of course, these owners might still *ask* for relief, which would transfer wealth from lenders' pockets to theirs. Conversely, borrowers who did need help had typically suffered adverse life events that are difficult for lenders to verify. Is a borrower who has lost his job really unable to find new employment, or is he failing to search in hopes of getting a break on his loan? Lenders and mortgage servicers were well aware of these incentive issues, and they did not hesitate to inform policymakers of their typical response to negative-equity borrowers who threatened to default. Philip Swagel, who served as Assistant Treasury Secretary for economic policy in the early part of the crisis, later recounted that

> As a practical matter, servicers told us, reputational considerations meant that they did not write down principal on a loan when the borrower had the resources to pay—never. They would rather take the loss in foreclosure when an underwater borrower walked away than set a precedent for writing down principal, and then have to take multiple losses when entire neighborhoods of homeowners asked for similar writedowns (Swagel 2009, p. 19).

[18]Of course, this policy would require the knowledge of each house's price, so a second condition needed for such a principal-reduction policy work would be for the prices of houses to be known without error, which is not possible in practice (recall Figure 5).

Related to the question of whether to extend mortgage relief was what form that relief should take. All of the models discussed so far—FOM, double-trigger and the combinations thereof—imply that sufficient principal reduction can prevent foreclosures, but these models also predict that a sufficiently large reduction in the monthly payment will do so as well. Determining the optimal form and size of mortgage relief requires knowledge of parameters that are unobservable, which makes modifications less attractive to lenders.[19]

Some basic guidance for constructing modifications emerges from the double-trigger model, which predicts that borrowers who default are likely to be liquidity constrained. These constraints cause at-risk borrowers to discount future gains highly, so large reductions in principal—that is, large reductions in payments they are scheduled to make *in the future*—generally provide less help than reductions in the payment that is due today. This logic is spelled out most clearly in Eberly and Krishnamurthy (2014), who compare various anti-foreclosure policies using a model that is rich enough to capture many of the relevant theoretical complexities and institutional details. The authors find that when troubled borrowers are liquidity constrained,

> transfers to households during the crisis period weakly dominate transfers at later dates and hence are a more effective use of government resources. These initial transfers could include temporary payment reductions, such as interest rate reductions, payment deferrals, or term extensions. This result is robust to . . . various forms of deadweight costs of default, debt overhang, and the easing of credit constraints through principal reduction. Generally, any policy that transfers resources later can be replicated by an initial transfer of resources, although the converse is not true (Eberly and Krishnamurthy 2014, p. 77).

Some empirical support for the effectiveness of payment reduction comes from Fuster and Willen (2015), who use a sample of loans for which required payments were reduced automatically due to a decline in overall interest rates.[20] The typical loans in the sample are Alt-A 5/1 ARMs originated in 2005 and 2006, which reset from their initial interest rates of more than 6 percent to a rate of 3 percent in 2010 and 2011.[21] The data provide a control group in the form of 7/1 ARMs originated at the same time but not due to reset until 2012 and 2013. The authors find that the size of the monthly payment has an economically large effect on the probability of repayment. The top panel of Figure 6, reprinted from their paper, shows that reducing the rate by 3 percentage points was equivalent to reducing a borrower's

[19] See the discussions in Foote, Gerardi, and Willen (2008) and Adelino, Gerardi, and Willen (2013) for a formal treatment of this problem.

[20] Tracy and Wright (2012) report similar results in a study of the effects of refinancing programs on defaults.

[21] Because the loans reset down and not up, there was no wave of prepayments shortly before the reset, eliminating a main selection issue discussed below.

CLTV from 135 to 95. The theories discussed in Section 2 lead to two possible explanations for this result. On one hand, payment reductions in the FOM induce repayment because they reduce the price of the call option on the house. On the other hand, these reductions also make it possible for households to better withstand other negative triggers such as job loss. Fuster and Willen (2015) argue for the latter explanation, because the improvement in repayment behavior appears to coincide with the reset and not anticipate it, as it would if the borrowers were living in a frictionless world.

3.4 Rate Resets, DTI Ratios, and Default-Prevention Policy

Policymakers have also instituted new mortgage-lending rules intended to prevent future problems. For the most part, the new rules reflect the view that the crisis was caused by the origination of mortgages that were unaffordable because of significant upward interest-rate resets from low initial "teaser" rates, or because of debt-to-income (DTI) ratios that were too high from the start. Accordingly, Section 1411 of the Dodd-Frank Act of 2010 requires lenders to make "a reasonable and good faith determination based on verified and documented information that, at the time the loan is consummated, the consumer has a reasonable ability to repay the loan." One way lenders can satisfy the ability-to-repay requirement is to originate a so-called qualified mortgage (QM). These mortgages are free of certain characteristics, such as interest-only periods and negative amortization. Additionally, the new Consumer Finance Protection Bureau (CFPB) has set the maximum possible DTI ratio for a qualified mortgage at 43 percent, although there are some important exceptions to this rule.[22] Income problems and liquidity constraints are crucial components of the modified double-trigger model, so it would seem that this model calls for ability-to-repay rules as effective ways to reduce future housing problems. Unfortunately, the reality is more complicated.

A first point is that even though concern over interest-rate resets figured prominently in policies to combat the crisis, especially early on, these resets were not a serious problem. Table 5 shows that the vast majority of borrowers who defaulted on their mortgages did so while being asked for monthly payments that were the same size or smaller than their payments at origination.[23] The most important reason that interest-rate resets caused so few foreclosures is that only a small fraction of borrowers actually saw their interest rates rise. Would resets have caused problems if they had been more common? This question is difficult

[22]DTIs can be larger for mortgages eligible for purchase by Fannie Mae and Freddie Mac. For more on qualified mortgages and the ability-to-pay rule, see two websites maintained by the CFPB: https://www.consumerfinance.gov/ask-cfpb/what-is-a-qualified-mortgage-en-1789/ and https://www.consumerfinance.gov/ask-cfpb/what-is-the-ability-to-repay-rule-why-is-it-important-to-me-en-1787/.

[23]The table comes from Foote, Gerardi, and Willen (2012); for confirmatory evidence, see Mayer, Pence, and Sherlund (2009).

to answer empirically, because creditworthy borrowers tend to prepay their mortgages if and when their mortgages reset. As a result, borrowers experiencing resets tend to be less creditworthy and therefore unrepresentative of the general borrower population.[24]

The use of DTI limits to prevent foreclosures presents another set of complications. Recall that one of the findings in the original default study by Herzog and Earley (1970) was that DTIs at origination were not significantly related to defaults. The authors explained this result by noting that DTIs lacked sufficient variation to identify their effect, because very few lenders allowed front-end DTIs greater than 25 percent. Yet the relative unimportance of origination DTI has emerged repeatedly in subsequent default studies, even as higher DTIs have been permitted.[25] The double-trigger theory can explain why, because that theory links defaults to adverse income *shocks*, not to low income *levels*. That is, default occurs when, say, one of two wage earners in a family loses a job, so that the household's DTI of 33 percent suddenly becomes a DTI of 67 percent. Defaults are not so much a problem for borrowers who take out mortgages with DTIs of 40 percent as compared to borrowers with DTIs of 35 percent, which is the comparison identifying the DTI coefficient in a standard default regression.

Of course, setting a low DTI limit reduces the probability that an income shock of a given size will cause a loan to become truly unaffordable. The problem is that income variances in the U.S. labor market are large, so trying to prevent defaults in this way allows only very low DTIs to be permitted. Assume, for example, that unaffordablity is defined as a DTI of 50 percent or more, but a 38 percent limit for DTIs at origination is imposed. Under these assumptions, the empirical volatility of a standard individual-level income process implies that just over a third of mortgages will still become unaffordable within the first three years.[26] A better way for lenders to prevent default would be to base lending decisions on variables correlated with the future *variance* of income, and the natural candidate for such a variable is the credit score. If income volatility contributed to default on previous loans, and if a borrower's income volatility is persistent over time, then credit scores will be much better than DTIs in predicting future defaults—which in practice the scores turn out to be.

All told, the double-trigger model implies that the affordability of mortgages (and not just negative equity) is critical for borrowers considering default. But the importance of unforecastable income shocks means that for any individual borrower, defaults are hard to predict and thus hard to prevent with origination-income restrictions. Lenders and policy makers can reduce defaults by limiting mortgage loans to people with stellar credit scores

[24]The Fuster and Willen (2015) paper cited above exploits plausibly exogenous variation in *downward* resets, so prepayments before the resets are not a problem. For more on the relationship between resets, prepayment, and default, see Foote et al. (2009).

[25]See, for example, Avery et al. (1996), Foote et al. (2009), and Lam, Dunsky, and Kelly (2013).

[26]The income process used here is estimated from the Panel Survey of Income Dynamics; see Foote et al. (2009) for details.

and sizable down payments, but restrictions of this type are likely to shut out large numbers of U.S. residents from home ownership. At the end of the day, policy makers must balance the benefits of dispersed home ownership with the increased number of defaults implied by a more liberal lending policy.

4 The Legacy of the Crisis: New Data Sources

The remainder of this paper discusses the new research opportunities and data sources that the recent housing cycle has brought about. We start with the data, but before doing so we stress an important point: compared to data used in virtually any other area of household finance, the data available for mortgage-default research has always been exceptionally good. There was still room for improvement, however, and the crisis brought about at least three major changes in this regard.

The first is that in 2007, data suppliers began to provide loan-level information on a broader array of mortgage types. In years prior, researchers were generally limited to loan-level data on mortgages held in private-label securities, which are bonds not guaranteed by government-sponsored enterprises (GSEs) such as Fannie Mae or Freddie Mac. The main private-label dataset, created by the LoanPerformance company and now owned by CoreLogic, includes loan-level data on the subprime, Alt-A, and prime jumbo loans that make up most private-label bonds. The long-standing availability of these data for most private-label securities contradicts the common claim that private-label issuers tried to hide information about the loans in their securities. In reality, the opposite is true.[27] In fact, at the start of the crisis, virtually nothing was known about the performance of individual mortgages *outside* private-label securities, including the portfolio loans held on bank balance sheets and the loans packaged into securities and backed by the GSEs. Starting in 2007, researchers began working with a dataset constructed by loan servicers who oversaw the payment collection on portfolio, GSE, and private-label loans. This servicer dataset, now called the McDash data, allowed researchers to compare the performance of different types of loans along a number of dimensions.[28] For example, Adelino, Gerardi, and Willen (2013) provided evidence against the widely held view that institutional problems in the private-label market prevented loan modifications there. In fact, modifications of loans in private-label securities were about as frequent as modifications of portfolio and GSE loans.

A second dimension of data improvement is that researchers are matching different datasets together at the individual level. One example of this matching has occurred be-

[27] See the discussion in Foote, Gerardi, and Willen (2012).

[28] The McDash dataset was formerly called the LPS dataset, after the name of the company that maintained it (Lender Processing Services). The McDash dataset is now owned by Black Knight Financial Services.

tween lender-supplied data and information about properties and mortgages from public deed-registry offices. The lender data, which come from trustees of securities or from servicers, include various characteristics of the borrower (for example, the credit score) as well as the loan's month-to-month payment history. But these data typically lack information about the property's other liens, so a combined loan-to-value (CLTV) ratio cannot be calculated. By contrast, public-records data from deed-registry offices lack detailed information about borrowers but do allow researchers to track all outstanding liens on the property, so a CLTV can be constructed. Of course, it is the CLTV that should matter in the borrower's default decision, and Figure 11 shows an example of how using the correct CLTV ratio can make a significant difference when estimating the default function.

A third data improvement is the use of credit-bureau data to explore mortgage default. One example is Elul and Hunt (2010), who use a matched sample of credit-bureau data from Equifax and loan-level data from McDash to investigate borrower debt positions around the time borrowers stop making their mortgage payments. In the next section, we discuss the use of credit-bureau data to identify so-called strategic default, which occurs when borrowers have the financial wherewithal to make their payments but decide not to because of deep negative equity. Credit-bureau data open up new possibilities to study default. But by themselves, these data have yet to yield ironclad insights, because the credit history of the borrower is endogenous with respect to the default decision. Having a lot of additional debt may cause a borrower to default on his mortgage, but the additional debt may also reflect a deeper problem, such as bad health, that leads to both the additional debt and the mortgage default.

Armed with these new data sources, researchers are now investigating a variety of questions related to recent macroeconomic history, general equilibrium in housing markets, and fundamental aspects of consumer behavior. We discuss three such questions in the next section.

5 The Legacy of the Crisis: New Questions

5.1 What Caused the Housing Boom?

Identifying the underlying causes of the early 2000s housing boom is an important project for economists today. Why did so many borrowers take out mortgages they did not repay, and why did their lenders give them the money? Early research on the boom tended to blame lax lending standards for the large number of subsequent defaults, and some of the most cited papers in this literature are based fundamentally on patterns of mortgage defaults.

One of the most direct uses of defaults to infer something about the boom is the one by Demyanyk and Van Hemert (2011), who run a series of default regressions on subprime mortgages that were securitized into private-label bonds. It is well known that subprime mortgages originated early in the boom tended to default less often than mortgages originated at the peak. A natural explanation for this pattern is house prices, which were generally rising for the early mortgages but falling for the later cohorts, so that the later mortgages were more likely to experience negative equity. The key contribution of Demyanyk and Van Hemert (2011) is to show that the later cohorts of securitized subprime loans defaulted more often *even after* controlling for such variables as the subsequent behavior of house prices. The implication is that lending standards had deteriorated significantly by the end of the boom. Figure 12 depicts the Demyanyk and Van Hemert (2011) result graphically, with the line in this panel showing the yearly cohort dummies from their regressions. The monotonic increase in this line indicates that the explanatory variables in the default regressions (including house-price appreciation) do not explain all of the increase in defaults among the later vintages of securitized subprime mortgages.

The panel also provides some context for the Demyanyk and Van Hemert (2011) results by relating their private-label cohort coefficients to those for other types of loans. The blue and red bars in the figure depict cohort coefficients from Lam, Dunsky, and Kelly (2013), who run default regressions on loans securitized by the GSEs (Fannie Mae and Freddie Mac) and on loans insured by the Federal Housing Administration (FHA). Two features about these other cohort effects are noteworthy. First, the effects for GSE and FHA mortgages also rise during the housing boom.[29] This increase in GSE and FHA coefficients indicates that if the private-label results point to problems in that market, such as bad incentives for mortgage originators, then these problems were relevant in the GSE and FHA markets as well.

A second point is that if GSE cohort effects reflect the strictness of lending standards, then they display some strange behavior from 1996 through 2000. The blue bars in the figure decline from 1996 through 1998 and then rise rapidly, so that the 2000 cohort effect is almost as large as the effects during the height of the housing boom. There is little outside information to corroborate increasing strictness in GSE standards in the late 1990s; in fact, some have criticized the GSEs for an unwise loosening of their standards during this period.[30] To be sure, the cohort effects reflect the *unobserved* component of underwriting standards, so if the GSEs were allowing, say, higher DTI ratios and smaller down payments in the late 1990s, then default regressions that treat those characteristics as explanatory variables

[29]Because the panel graphs *coefficients* from default regressions rather than *conditional probabilities* of default, comparisons of the levels of coefficients across different loan types is not informative. The information in the figure is limited to the time-series patterns for the set of coefficients estimated using a single type of loans.

[30]See, for example, Wallison and Pinto (2012).

would not necessarily generate cohort effects that rise over time. But the early GSE pattern does suggest that cohort effects could also be influenced by features of the macroeconomic landscape that are difficult to model precisely in default regressions. For example, the good performance of the late-1990s loans could relate to expectations about future house prices (which are unobservable), while the bad performance of the 2000 mortgages could stem from high unemployment after that year (which, as we have seen, is difficult to proxy at the individual level with area-level unemployment rates). Because a cohort effect sweeps up any misspecification in a default regression for a given year, interpreting this effect solely as the unobserved component of underwriting standards may be too broad.

A second paper that uses default patterns to infer something about boom-era lending standards is Keys et al. (2010). This paper also uses private-label mortgage data, but unlike Demyanyk and Van Hemert (2011), it focuses not on coefficients from a default regression but on a straightforward regression discontinuity. Because borrowers with higher credit scores tend to default less often, a scatter plot of average default rates and credit scores will have a negative slope. Keys et al. (2010) show that in plots of this type, the default rate among private-label loans appears to jump *up* at the credit score of 620. That is, borrowers with credit scores of 621 have somewhat higher default rates than borrowers with scores of 619, even though the continuous nature of credit scores would imply that the 621 borrowers should default less often. Additionally, the authors show that the number of loans in their private-label dataset also rises discretely at the 620 cutoff. They argue that this increase in quantity results from long-standing guidelines promoted by the GSEs that made it easier to securitize loans with scores at or above 620. Putting these two facts together, the authors argue that securitization led to lax lending standards. Because loan originators could offload the mortgages of the 621 borrowers into the securitized market, these originators did not screen the 621 borrowers carefully, and defaults among these borrowers were higher as a result.

Bubb and Kaufman (2014) raise questions about this analysis, and their paper highlights the value of the data sources that are now available outside the private-label market. Using the more comprehensive McDash data, the authors repeat the Keys et al. (2010) discontinuity analysis for portfolio rather than securitized loans. Surprisingly, they find the same pattern that Keys et al. did: defaults are higher for borrowers just above the 620 cutoff and lower just below it. (A similar pattern exists at the score of 660.) This pattern cannot result from securitization, because none of the loans in this particular exercise was securitized. Second, the authors calculate the securitization rate at the 620 and 660 cutoffs—a rate that Keys et al. could not calculate because all of the loans in their private-label dataset were securitized. In another unexpected result, Bubb and Kauffman find no large differences in the securitization rate at either the 620 or 660 cutoffs, even though a large difference in securitization probabilities at 620 underpinned Keys et al.'s interpretation of their findings.

What model can explain all of these results? Bubb and Kauffman point out that underwriting mortgages is costly, so lenders will naturally try to focus their attention in a cost-effective way. Borrowers with very good credit scores are unlikely to default, so loan officers do not need to spend much time with them. Lenders should instead expend time and energy on marginal borrowers, who have low credit scores, so that they can determine whether any special circumstances should offset those poor scores in their lending decisions. Under some general conditions, Bubb and Kaufman show that a lender is likely to adopt a *cutoff rule*, so that all borrowers above the cutoff get less attention, and those below the cutoff get more. In fact, Bubb and Kauffman note that a cutoff rule of this type was the source of the ostensible securitization limit cited in Keys et al. (2010). The GSEs never said that they would not securitize loans below 620, only that loans below that cutoff needed extra attention by loan officers before they could be sold to the GSEs.

In a market of this type, Bubb and Kaufman contend that we would expect fewer loans to be made to borrowers below the 620 cutoff, because most of those borrowers would not have the extenuating circumstances needed to convince loan officers to decide in their favor. But the sub-620 borrowers *who did get mortgages* had to have been special in an important sense, because they convinced their lenders that extenuating circumstances made their credit scores inaccurate indicators of their true creditworthiness. The special nature of the sub-620 borrowers implies that we should not be surprised that they defaulted less often than borrowers just above the cutoff, who received no special scrutiny, even though their credit records were almost as bad. And because the benefit of using a cutoff rule is a general feature of *lending*, not an institutional feature of *securitization*, it is also unsurprising that the lending and default patterns found in the original Keys et al. paper show up in portfolio loans as well. In later work, Keys et al. have responded to this critique by contending that the analysis must be performed on finely disaggregated loan types, for example, non-agency loans with low documentation. But Bubb and Kauffman's basic point is that cutoff rules are more relevant for lending in general, not securitization in particular. This observation seems to introduce serious complications for researchers trying to discern the ultimate effects of securitization on the housing boom with a regression discontinuity approach.

A third paper linking default rates to the underlying causes of the boom is Mian and Sufi (2009), which investigates the patterns of default across different ZIP codes. The paper begins by noting that "a salient feature of the mortgage default crisis is that it is *concentrated* in subprime ZIP codes throughout the entire country" (p. 1449, emphasis added). Specifically, the authors show that ZIP codes where people tended to have low credit scores (and, typically, low incomes) were also the ZIP codes that saw the biggest increases in mortgage defaults. That fact and others in Mian and Sufi (2009) lead the authors to interpret the recent boom and bust as fundamentally related to subprime lending. Adelino, Schoar, and Severino (2016) take issue with this concentration claim. The authors use McDash data to show that in *absolute* terms, the Mian-Sufi statement is correct: increases in defaults were indeed larger in low-income ZIP codes. But on a *proportional* basis, the increase in defaults in high-income ZIP codes were as great or greater than the percentage increase in low-income areas. These two facts are consistent, because in any given year, defaults tend to be higher in poorer ZIP codes. Consequently, a large absolute increase in defaults in a low-income ZIP code could still result in a small percentage increase due to a high initial level of defaults. In a sense, Adelino, Schoar, and Severino (2016) turn the concentration finding on its head, arguing that the large percentage increase in defaults in richer areas indicates that the boom-bust cycle was not just a subprime event, but one that affected middle-class borrowers as well.

To sum up, initial research on the source of the housing boom made intensive use of defaults. These papers pointed to private-label mortgages in general, and subprime mortgages in particular, as the source of most problems. But these early interpretations have been challenged. New investigations of defaults during the bust and of debt increases during the boom have led some authors to call for a "new narrative" on the housing cycle that stresses its ubiquity across all income classes. Consensus has yet to be reached, but mortgage defaults will no doubt feature prominently in future research.[31]

5.2 Do Foreclosures Reduce House Prices?

Until now, causality in the default function has run from equity to default. In general equilibrium, however, large numbers of defaults could cause house prices to decline and negative equity to become more prevalent. Economists have suggested two plausible mechanisms for how defaults could drive prices lower. First, borrowers who know they are going to lose their homes may spend less on keeping them up, generating negative externalities for nearby properties. Second, a foreclosure could lead to an increased supply of residential property on the market, driving down prices through a supply effect.

Figure 13 illustrates how either channel could confound the estimate of the default function. It posits a world in which equity has no effect on default probabilities, so the true default curves in the figure are horizontal. But suppose that between, say, 2006 and 2007 some exogenous factor (perhaps a wave of interest-rate resets) shifts the horizontal default function upward. If defaults lead to lower house prices, then the distribution of equity will shift leftward, as seen in the lower panel of the figure. The black dots in the upper panel represent data generated by the confluence of lower equity and an upwardly shifted default

[31] Aside from those cited above, papers in the new-narrative literature include Kaplan, Mitman, and Violante (2017); Albanesi, De Giorgi, and Nosal (2017); and Foote, Loewenstein, and Willen (2016).

function, with the dots on the right and left sides of the upper panel corresponding to 2006 and 2007 data, respectively. An econometrician might estimate the downward sloping default function shown in the top panel if he ignores the general-equilibrium effect of defaults on house prices. Of course, researchers exploring these issues are not interested only in the default function, because any general equilibrium effect running from defaults to prices would be interesting in its own right.

Most researchers exploring these issues have looked for very local spillover effects, with a default on one house affecting the prices only of nearby houses. Consider the following spatial-externality regression:

$$\log(P_{it}) = \alpha + \beta X_{it} + \gamma NF_{it} + \varepsilon_{it},$$

where P_{it} is the sale price of property i in period t, X_{it} is a vector of controls, and NF_{it} is a measure of the number of properties that experience some type of foreclosure event within a certain distance of property i in some window around period t. The coefficient of interest is γ. The original example of this general type of spatial externality regression is Galster, Tatian, and Smith (1999), which looked at the effect of investment on the sale prices of nearby properties, and Immergluck and Smith (2006) adapted the methodology to look at foreclosures. Others followed with different variations on the same basic regression, including Schuetz, Been, and Ellen (2008); Rogers and Winter (2009); Harding, Rosenblatt, and Yao (2009); and Lin, Rosenblatt, and Yao (2009). All of these papers find statistically significant but relatively minor spillover effects, which were typically confined to a small geographic area over a short time period.

In recent years, researchers have advanced the state of the art significantly. Ironically, Campbell, Giglio, and Pathak (2011, henceforth CGP) innovate by raising an obvious issue with the spatial-externality equation above, which is that an unbiased estimate of γ requires the assumption that NF is exogenous to P. Yet all the theories discussed above propose a strong causal link running from prices to defaults, because lower prices raise the incidence of negative equity and thereby raise foreclosures. Consequently, if either the FOM or the double-trigger model is correct, then the equation above is misspecified. CGP address this problem by taking two sets of differences designed to measure the effect of an additional foreclosure within 1/10 of a mile of a forced sale. First, the authors include in their regressions all foreclosures occurring in the year prior to the sale as well as the foreclosures occurring in the year after. The logic here is that the coefficient on sales that occur after a foreclosure includes both the effect of P on NF and the effect of NF on P, whereas a foreclosure that occurs after the sale cannot affect the price of the property and therefore measures only the effect of P on NF. By subtracting the coefficient on foreclosures that occur after from the coefficient on foreclosures that occur before, the authors argue, they are essentially subtracting the effect of P on NF from the combined effect of P on NF *and* the effect of NF on P. They are thus left with the effect of NF on P, the quantity of interest. CGP also perform a second round of differencing by including, as a control, foreclosures before and after within 1/4 mile of the sale. The idea here is that doing so controls for a common shock that causes those foreclosures as well as price declines.

CGP find statistically significant effects for the entire sample of sales in Massachusetts, but these effects are small. When they restrict their sample to single-family properties, however, the effect is statistically and economically insignificant. The interpretation of these results is difficult. The authors do find small but significant effects for condominiums, but those effects may reflect issues with common ownership of property and common payments into maintenance accounts. Moreover, CGP implicitly assume that sales prices are affected neither by foreclosures that occur after a property sells, nor by those that occur more than 1/10 mile away. Both are strong assumptions. The timing assumption, for example, implies that properties that are poorly maintained prior to foreclosure cannot affect the sale price of neighboring properties.[32] Consequently, one must be cautious in concluding that foreclosures have no effect on the prices of single-family homes.

Gerardi et al. (2015) also attempt to address the problem that P should causally affect NF by combining a repeat-sales methodology with highly disaggregated geographic fixed effects. Essentially, the authors' regression exploits variation in the number of nearby foreclosures across two properties bought in the same year, sold in the same year, and in the same census block group. For the effect of P on NF to explain their results, one would have to posit that there were distinct submarkets within an area of just a few city blocks, and that those submarkets were hit by different economic shocks over the relevant period. The dataset used in the paper includes not just the location of contemporaneous foreclosures, but also the location of seriously delinquent properties. It also includes information about the condition of bank-owned foreclosed properties. Sales prices for homes within 1/10 mile radii appear to be affected by delinquent properties prior to foreclosure, by bank-owned properties, and by properties sold by the bank within the last year. The price effect for all of these types of distressed properties is roughly the same, 1 percent. Finally, the authors also find that the worse the condition of a bank-owned property, the lower the price of nearby sales. But a bank-owned property in above-average condition actually sells for a higher price than a property with no foreclosures nearby.

Hartley (2014) looks at foreclosures in Chicago and estimates different γs by structure type. He finds that while the γ associated with a single-family foreclosure on a single-family sale is significantly negative, the γ associated with a multi-family foreclosure on a single-family sale is not. Hartley interprets this as evidence that the supply of property, rather

[32] For more on the effect of poor maintenance on prices, see Lambie-Hanson (2015).

than investment externalities, drives the discount. However, while the multi-family γ is statistically insignificant, it is economically large and positive, raising questions about its interpretation. Further, Hartley finds effects only within 0.05 mile and no effect between 0.05 and 0.10 mile, contradicting most previous research, including CGP. Anenberg and Kung (2014) examine sales of single-family properties near San Francisco from 2007 to 2009, and augment CGP by including information about listings. The authors argue that there is no foreclosure externality, because their estimate of the γ for foreclosures that occur prior to a sale is zero. In this respect, their results contradict not only CGP but also the other research cited above.

Guren and McQuade (2016) build a search model to illustrate how sales of foreclosed properties can affect the prices of nondistressed homes. In their setup, banks foreclose on properties and then aggressively sell them, while foreclosed borrowers go into a separate rental market for a certain period of time. Foreclosures drive down prices for two reasons. The first is a supply effect, because the eagerness of banks to sell increases the bargaining power of buyers. The second is a demand effect, which stems from the relegation of the defaulted borrowers to a separate rental market, where they cannot bid on homes. The supply effect turns out to have minimal effects on house prices. The demand effect is very strong, but it depends on assumptions about the structure of housing markets that are probably unrealistic. Specifically, the authors assume that the increased demand of the defaulted borrowers in the rental market has no effect on prices in the housing market as a whole. But foreclosed properties are typically purchased by investors who intend to rent them out. The higher rental demand generated by recently defaulted homeowners is therefore likely to raise the amount that investors will be willing to bid for foreclosed properties. For this reason, casual empiricism argues against the authors' contention that the rental market is sealed off from the housing market as a whole. Interestingly, the small supply effect in Guren and McQuade (2016) argues against most popular thinking about foreclosures and prices. Many have argued that a flood of foreclosed properties will drive prices lower, but the results in the formal model are driven almost entirely by flows of potential buyers in and out of the housing market.

Indeed, when thinking about a potential feedback between defaults and prices, one must remember that at a basic level, the price of houses is determined by the supply of and demand for houses, not by the supply of and demand for houses *on the market*. All houses are in some sense available for sale at the right price, so the supply of houses already includes all foreclosed and unforeclosed properties. As Guren and McQuade (2016) effectively illustrate, foreclosures create new supply but, through the rental channel, foreclosures provide new demand for housing as well. Is is possible to construct a model in which foreclosures reduce prices? Yes, as Guren and McQuade (2016) show, this is possible. But different assumptions about the structure of housing markets might instead generate a positive effect of defaults on prices. For example, if foreclosures lead to a deterioration of the capital stock through neglect, then models might predict that the stock of well-maintained homes would become more valuable, not less. And the historical record shows that house prices can rise rapidly in places with large numbers of foreclosures and bank-owned properties. The rapid price appreciation observed a few years ago in Phoenix, Las Vegas, and other foreclosure hotspots is a case in point.

While most researchers interested in foreclosure externalities have attempted to measure the effect of defaults on the prices of nearby properties, other researchers have explored a conceptually distinct question: whether a foreclosure on one house makes foreclosure on nearby houses more likely through a contagion effect. One example is Goodstein et al. (2012), who include the foreclosure rate in nearby ZIP codes as an explanatory variable in a standard default regression. The authors argue that the positive and significant value of the resulting coefficient is evidence of contagion. Unfortunately, all such regressions suffer from the famous reflection problem of Manski (1993), which comes about when a researcher measuring the behavior in a population tries to infer whether the group's average behavior influences the actions of its individual members. The reflection problem is not easy to solve (Graham and Hahn 2005), and in our view, no one in the contagion literature has been able to do so.[33]

5.3 How Important Is Strategic Default?

As defaults mounted in 2007 and 2008, commentators expressed concern about a phenomenon they called "strategic default." Economists have provided an array of definitions for this concept. Guiso, Sapienza, and Zingales (2013) write that strategic default occurs when borrowers "choose to walk away from their houses even if they can afford to pay their mortgages." Tirupattur, Chang, and Egan (2010) define the concept as "the proclivity of borrowers to default on their mortgage payments when they have the ability to make them," and Experian and Oliver Wyman (2009) define strategic default as "borrowers defaulting on their mortgages *only* because the value of their home has declined *well below* their mortgage [emphases added]."

These definitions of strategic default are somewhat problematic. The first two suffer from the presence of the words "afford" and "ability to pay." A consumer can afford a bundle of

[33] One avenue that might be explored in future work is the possibility that high levels of contagion lead to multiple equilibria. If, for example, stigma against default prevents borrowers from defaulting, then we could envision small exogenous differences in stigma leading one community to a low-stigma equilibrium with many foreclosures, while another community settles on a high-stigma equilibrium with few foreclosures. Researchers have explored this possibility in other areas, such as crime (Glaeser, Sacerdote, and Scheinkman 1996), but to our knowledge no one has yet tried to do so for foreclosures.

goods or services if the cost falls short of the funds he has available. Consider a borrower whose wife becomes ill and can't work. Suppose he can make his mortgage payment if he takes two jobs, rents his bedroom to boarders, sleeps on the kitchen floor, and eats generic boxed macaroni and cheese for every meal. Yet he still chooses to default. Strictly speaking, the borrower can afford the mortgage payment, but would one label this a strategic default? Probably not, but according to the first two definitions, it would be. The definition of Experian and Oliver Wyman (2009) is in some ways the best, because it focuses attention on the absence of a trigger, but it ignores the monthly payment. Consider a negative-equity borrower with plenty of money in the bank who is notified that his monthly payment will soon increase. The borrower concludes that the monthly payment will exceed the value of the implied call option, so he defaults. Would we call this a strategic default? Probably yes, but according to Experian and Oliver Wyman (2009) we would not, because the payment increased.

Perhaps the most rigorous way to approach this question is to define non-strategic default as one that occurs because of a life event suffered by the borrower, so that strategic default can be defined as the complement. In the context of the models in Section 2, all defaults in the FOM are strategic, because life events play no role in that framework. Conversely, none of the defaults in the simple double-trigger model is strategic, because life events are necessary for default. One of the values of the new optimizing models that combine the FOM with double-trigger is that strategic and non-strategic defaults can co-exist in one model. Some wealthy borrowers will default despite suffering no shocks, while other borrowers will default explicitly because of shocks.

Empirically, researchers have tried to identify strategic default in several ways. Experian and Oliver Wyman (2009), Keys et al. (2012), and Tirupattur, Chang, and Egan (2010) use credit-bureau data to identify borrowers who default on their mortgages but not on other credit lines. The thinking here is that if a borrower is making payments on all other debts, then financial distress cannot be motivating the mortgage default. The credit-bureau approach implies that strategic default increased dramatically over the course of the crisis, and that highly creditworthy borrowers defaulted strategically more often than less creditworthy borrowers. At the time, many viewed this pattern as an ominous sign for the housing market. But the pattern is most likely an artifact of this particular method of identifying strategic default, because the credit-bureau approach suffers from what we might call Type I and Type II errors.

The Type II error is that the credit-bureau method identifies many borrowers as *not* strategically defaulting when in fact they are. Borrowers with low credit scores are routinely delinquent on their mortgages and obligations. Years ago, Herzog and Earley (1970) referred to 30 days past due as "casual delinquency," and it has long been known in the industry that casual delinquency is not a cause for concern when it occurs among borrowers with poor credit histories. In fact, Adelino, Gerardi, and Willen (2013) show that borrowers with low credit scores were *more* likely to cure from delinquency than borrowers with high credit scores. These results imply that the credit-bureau approach more or less precludes subprime borrowers from defaulting strategically, because they are much more likely to have cured from some delinquency in the months prior to default on the mortgage. The finding of Tirupattur, Chang, and Egan (2010) that strategic delinquency rates were low early in the crisis probably reflects the fact that early delinquency problems were concentrated in the subprime population.[34]

Type I errors in the credit-bureau approach identify borrowers as strategic defaulters, when in fact they are not. These errors will be more numerous among creditworthy borrowers, who manage their finances well and usually pay all of their bills on time. When confronted with a loss of income, borrowers with a good credit history are more likely to carefully analyze their finances. If they realize that paying their mortgage represents an unbearable sacrifice, they will default so that they can divert funds to pay other bills. The credit-bureau approach would label this a strategic default, so measurement errors of this type help explain why strategic defaults were found to be more common among creditworthy borrowers.

The second approach to measuring strategic default is to survey households. A novel example is Guiso, Sapienza, and Zingales (2013), who ask people how many people they know who have defaulted on their mortgages, and what fraction of those people could have afforded their payments. Survey data from March 2009 to September 2010 imply that the share of defaults that were strategic grew from 25 percent to 35 percent. Of course, it is unconventional for economists to ask people about their own motivations, let alone those of other people.[35] Survey responses could be influenced by media discussions of strategic default and by the respondent's own economic situation. Irresponsible borrowers were blamed by some for causing the crisis, so a borrower who was truly suffering during the crisis might be more likely to accuse his neighbor of defaulting strategically. At the same time, the economic impact of the Great Recession may have led respondents to have more sympathy for struggling families, so they would be less likely to label borrowers as strategic defaulters.

Guiso, Sapienza, and Zingales (2013) then ask borrowers to envision whether *they* would default at certain negative-equity thresholds, even if they could afford the payment at each threshold. Yet as discussed in Section 2, standard theory implies that the size of the monthly

[34]Consider a borrower with a 640 FICO score in 2005 who received a reduced documentation loan to speculate on a condominium in Las Vegas. When the market collapses, he realizes he is unlikely to make a profit on the transaction, and he defaults. Because he has other delinquent accounts, researchers using the credit-bureau approach would classify him as a nonstrategic defaulter.

[35]For example, economists have long studied whether unemployed workers actively seek work. To our knowledge, no one has ever asked survey respondents if the unemployed people they knew were really searching.

payment matters for default in a way that calls into question the validity of the resulting answers. For example, self-reported default propensities are higher among African-American borrowers, a finding that the authors attribute to social attitudes. But African-American borrowers also tend to have lower wealth and credit scores, so they must pay higher interest rates. The resulting higher monthly payments, not social attitudes, could then explain the reported default propensities. A deeper problem is that this evidence relies on the ability of households to imagine a decision in a hypothetical situation, something that behavioral economists have flagged as problematic (Loewenstein, O'Donoghue, and Rabin 2003).

Bhutta, Dokko, and Shan (2017) take yet a third approach to the problem that is similar to taking differences of default rates across various classes of borrowers. The easiest way to understand the differencing method is to apply it to the data in Table 4, which presents city-level default rates for borrowers in different equity and unemployment categories. The identifying assumption behind the Bhutta, Dokko, and Shan (2017) method is that borrowers with positive equity or modest negative equity must be exclusively double-trigger defaulters, because the FOM predicts payment in those cases. We can therefore use the 81–100 percent CLTV column in Table 4 as a baseline at which all defaults are double-trigger. As seen in the third row of this column, borrowers in cities with unemployment rates from 8.1 percent to 10 percent had a default rate of 0.52 percent. The default rate was 1.81 percent for borrowers two columns to the left, who had the same unemployment rates but a substantially higher level of negative equity: CLTV $>$ 120 percent. Thus the inferred strategic default rate for the CLTV $>$ 120 borrowers is $1.81 - 0.52 = 1.29$ percent, implying that roughly two-thirds of defaults for this group were strategic.

While creative, this approach is risky, because it assumes that all observed defaults at low levels of negative equity are double-trigger. Yet Figure 5 shows that measurement error in house prices—and thus in home equity—is a nontrivial problem. Even if all borrowers were FOM, and thus all borrowers defaulted only at high levels of *actual* negative equity, we might still see large numbers of defaults at mild levels of *measured* negative equity. A careful reading of Table 4 indicates another potential problem with the differencing method. The basic principle behind strategic default is that it should be independent of individual characteristics such as unemployment. Defaults at high levels of negative equity—which are more likely to be strategic—should therefore be less sensitive to unemployment. But Table 4 shows that defaults at very high CLTVs are more sensitive to unemployment, not less.

Finally, Gerardi et al. (2017) make use of the PSID, which is small in size but includes information on labor earnings and home values at the household level. The authors first construct household budget sets, which allow them to identify a "can pay" household as one that can make its required mortgage payment without reducing its consumption below some reference consumption level. A useful way to construct this limit is by referencing the level of consumption consistent with home affordability according to the Department of Veterans Affairs (VA), which administers a government-insured mortgage program for service members.

The main results of the exercise are summarized in Figure 6. Column 1 refers to households who definitely can make their mortgage payments, because their incomes after these payments ($y - m$) are larger than their reported consumption levels (c). At the other extreme are the "can't pay" households in column 3, each of which has income less mortgage payments that is below the VA-defined consumption level for a household with a similar demographic makeup: $y - m < c(VA)$. The middle column contains intermediate households, for whom making the mortgage payment requires a consumption level that is below reported consumption c but above the VA minimum $c(VA)$. The different panels of the table correspond to the entire population of homeowners (Panel A) and to those with either negative or positive equity (Panels B and C, respectively).

Several findings emerge. The first is that negative equity and income shocks are both important drivers of default. The default rate of can't-pay households with negative equity is more than 30 times higher than the rate for can-pay households with positive equity (19.7 percent versus 0.07 percent). Second, strategic default is a non-trivial issue, because many defaulters can afford their monthly payments. Of the 196 total defaults in the PSID (Panel A of column 4), 74 (37.7 percent) came from the can-pay households in column 1. The third finding is that despite the importance of strategic motives in overall default, almost no one who can afford the payment defaults. The 74 can-pay defaulters are a significant fraction of all defaulters (196), but they comprise less than 1.5 percent of all can-pay households (5,173). The seemingly paradoxical second and third facts are reconciled by a fourth finding: default is very rare across the entire population. In fact, the 19.7 percent default rate among the most troubled group—can't-pay with negative-equity—indicates that more than 80 percent of even these distressed households will keep their homes. Because overall defaults are so low, strategic defaults (that is, can-pay defaults) can be a large share of defaults, even if almost all can-pay households remain current on their mortgages.

6 Conclusions and Directions for Future Research

So where does mortgage-default research go from here? Although mortgage data are already quite good, there is still scope for improvement, and matching different individual-level datasets offers the largest potential payoffs. Combining loan-level datasets with employment records would yield a better understanding of the role of unemployment in default, while linking loans to other vital records (such as marriages, births, deaths, and divorces) would help identify other default triggers. On the theoretical front, work in the spirit of Campbell and Cocco (2015), Corradin (2014), Laufer (2017), and Schelkle (2014) will continue to combine the rigorous math of the FOM with the intuition and real-world liquidity constraints embedded in the double-trigger model.

As these new data sources and models are developed, we believe the central question that researchers will confront is why default is so rare. As the foreclosure crisis has waned in recent years, many economists report being surprised that the default rate did not turn out to be even higher than it was.[36] Some of these economists may not have been familiar with how the option value of waiting inhibits default at modest levels of negative equity. But at the end of the previous section, we noted that even financially distressed households tenaciously hang on to their homes, and other research shows that default is also rare among households with very deep negative equity. Bhutta, Dokko, and Shan (2017), for example, find default rates of less than 10 percent for borrowers with CLTV ratios of 200 percent. As these authors and others have pointed out, default rates this low are difficult for any model of the housing market to rationalize (White 2010).

A potential institutional explanation for low defaults rate is lender recourse. Ghent and Kudlyak (2011) find that states allowing lenders to enter deficiency judgments against defaulting borrowers experience lower defaults, especially on higher-priced homes. But these laws are limited geographically and difficult to enforce, so they do not explain why default rates are so low for the nation as a whole. Other potential explanations for low defaults include transactions costs of default, the non-pecuniary attachment of households to their homes, social stigma, and the increasingly negative consequences that poor credit has on other aspects of economic life, such as employment.

We also point to another potential explanation: after housing busts, nominal house prices tend to recover more quickly than standard models imply. Earlier, we saw rapid price recovery in California and Massachusetts after the coastal busts of the early 1990s. Figure 14 shows similar price recoveries after the recent bust for the nation as a whole and for the three other "sand states" (Arizona, Florida, and Nevada). As noted earlier, expected reversion of house prices is hard to embed in formal models of the housing market, because no-arbitrage conditions typically restrict the house-price process. But homeowners may know something that the models do not—price busts are temporary—so owners may rationally expect the restoration of positive equity within a few years. The correct model of mortgage default, yet to be devised, may have to account for this phenomenon, despite the difficulty of modelling mean reversion in asset prices.

[36] For evidence that economists were surprised at the low default rate, see the published general discussion of Eberly and Krishnamurthy (2014).

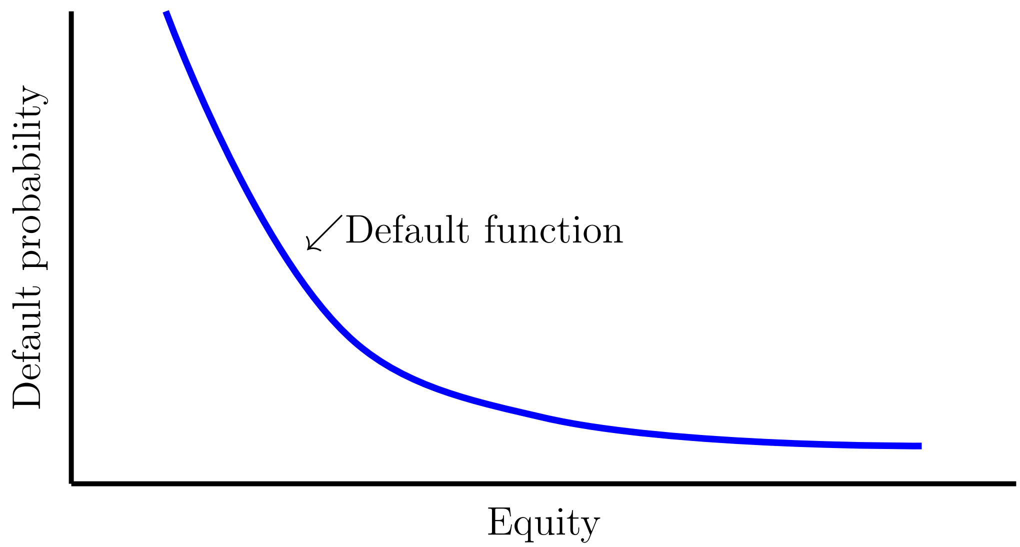

Figure 1. The Default Function.

	Assets	*Liabilities*	*Net Position*
Before default	House (H), Put Option (P)	Loan (M)	$H - M + P \geq 0$
After default	No House, No Put Option	No Loan	0

Table 1. The Borrower's Balance Sheet Before and After Default.

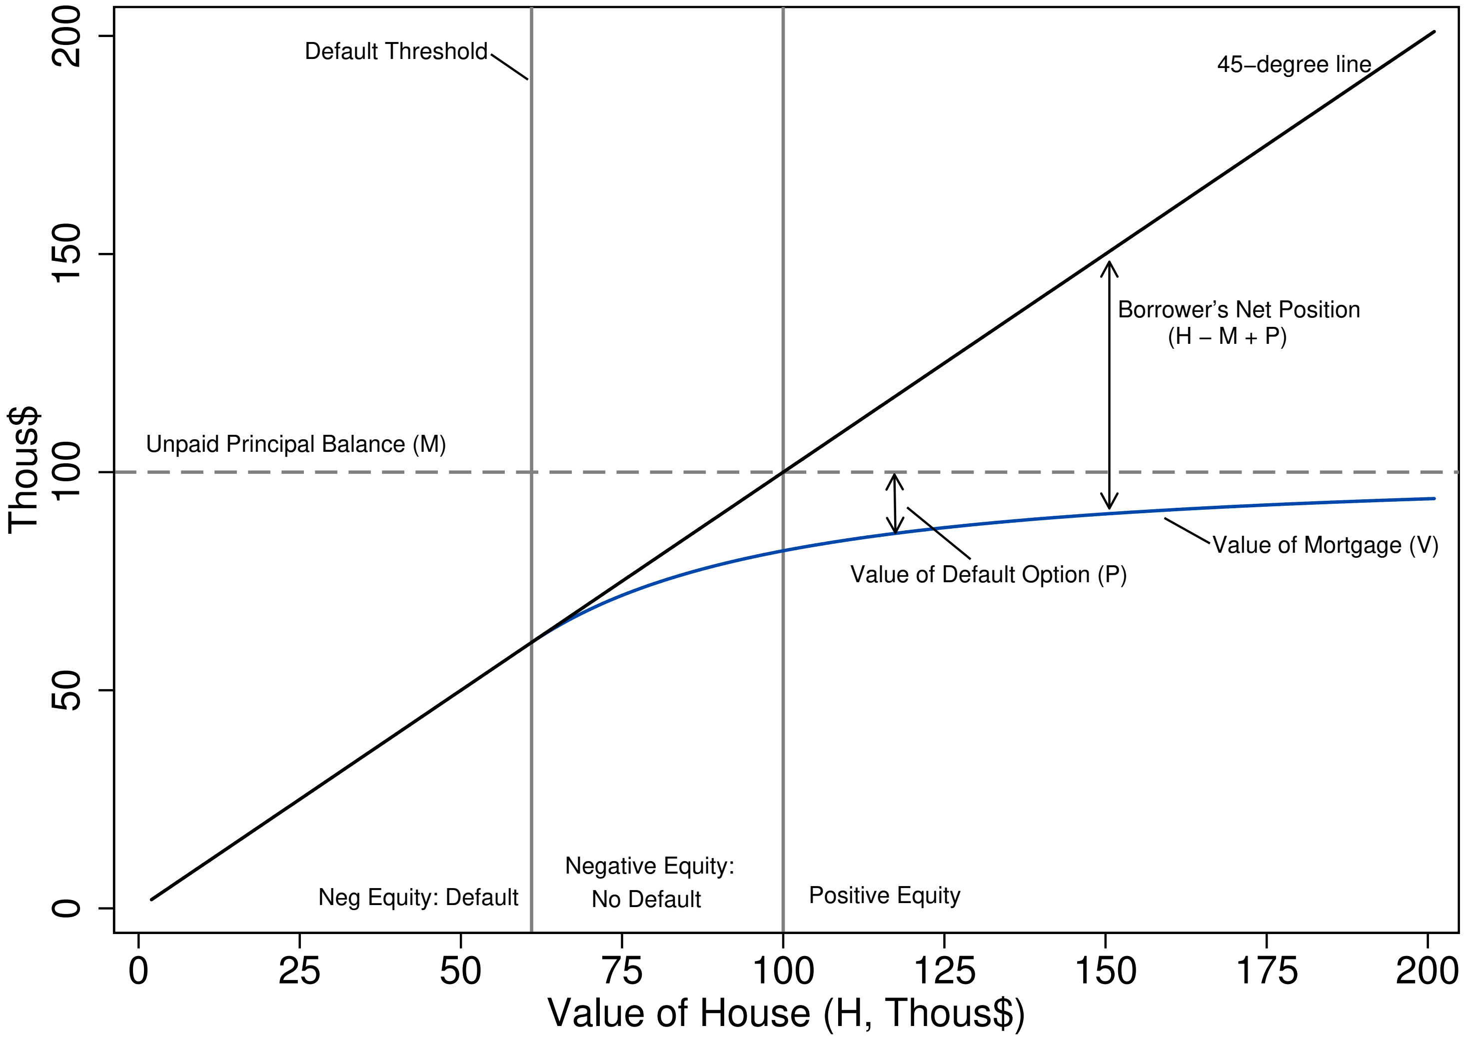

Figure 2. A Graphical Depiction of the Frictionless Option Model (FOM). *Note:* See Section 2 for a description of the model and for parameter values chosen.

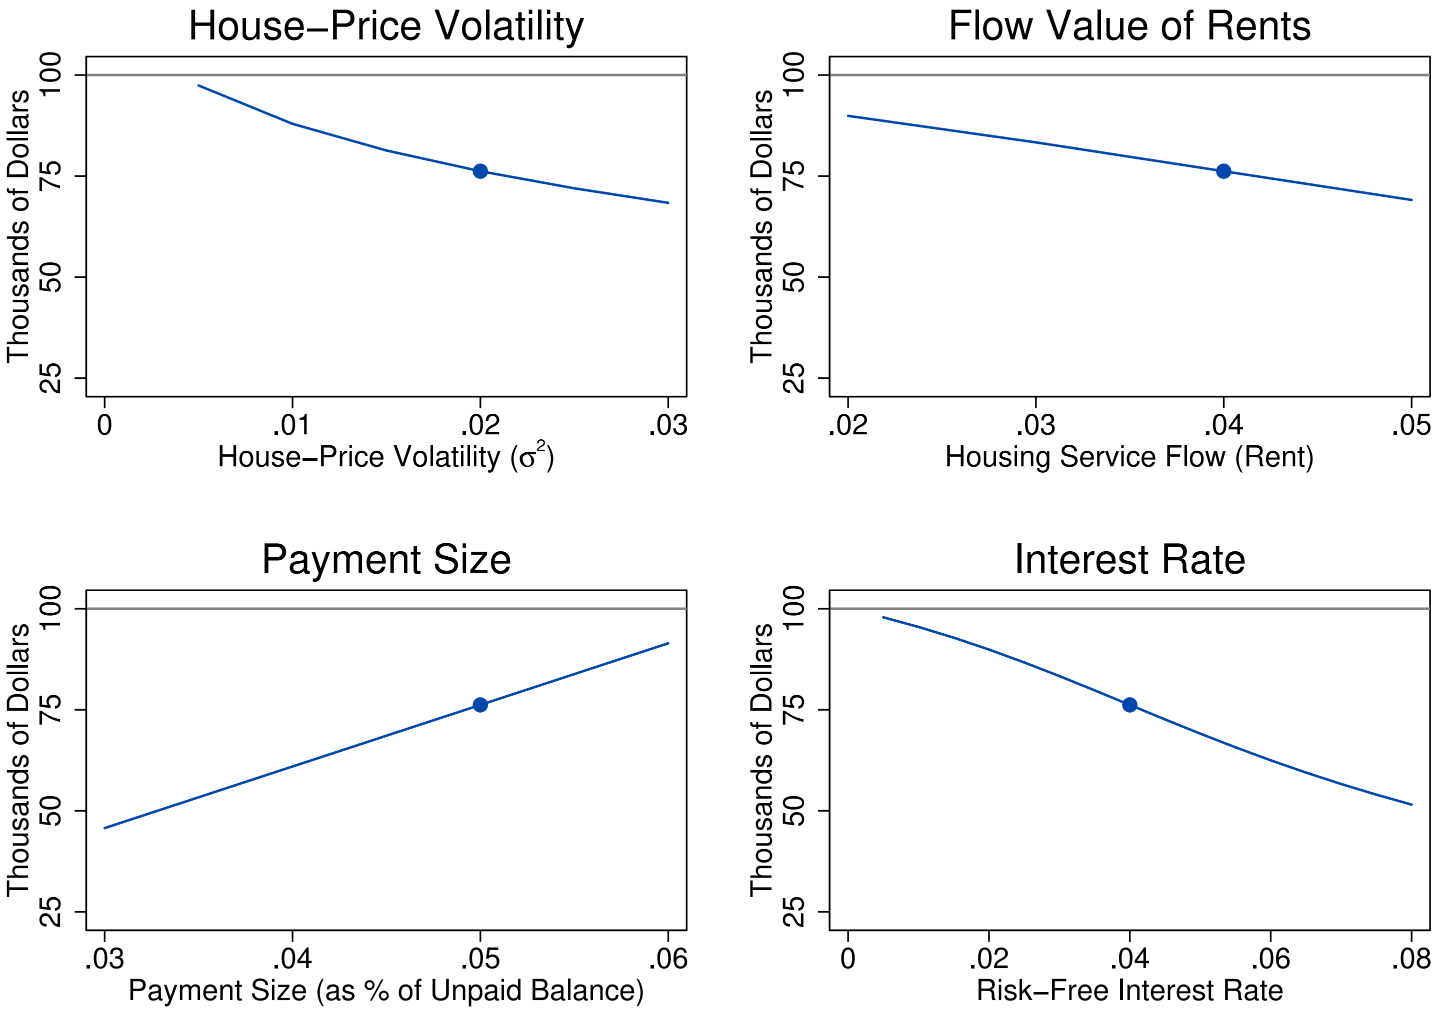

Figure 3. The Sensitivity of Default Threshold to Parameters of the FOM. *Note:* The dots represent a particular set of baseline values for the variables in the FOM. Each panel shows how the default threshold changes when a particular parameter value is changed while all of the other parameter values are held constant.

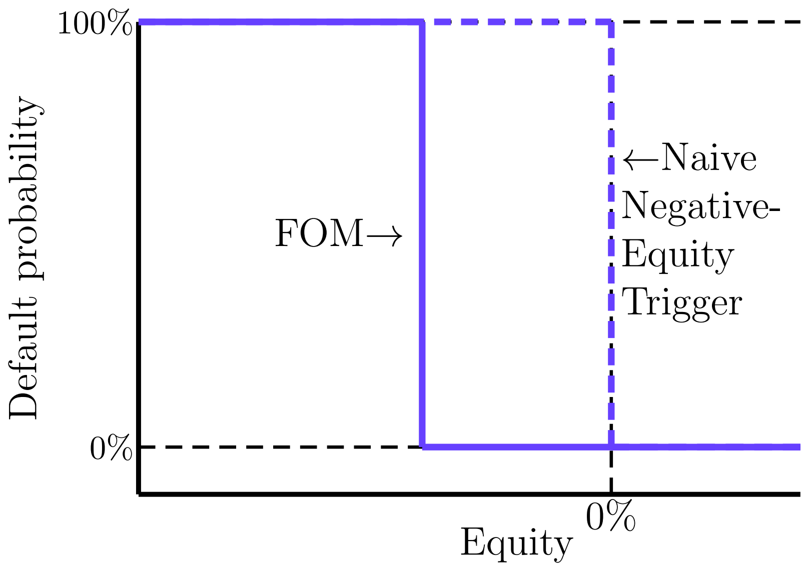

Figure 4. The Default Function in a Frictionless Model. With a naive negative-equity trigger, the borrower defaults as soon as the price of the house falls below the outstanding balance of the mortgage. The Frictionless Option Model (FOM) recognizes that because house prices could recover, there is an option value to waiting to default, so borrowers with modest amounts of negative equity remain current on their mortgages.

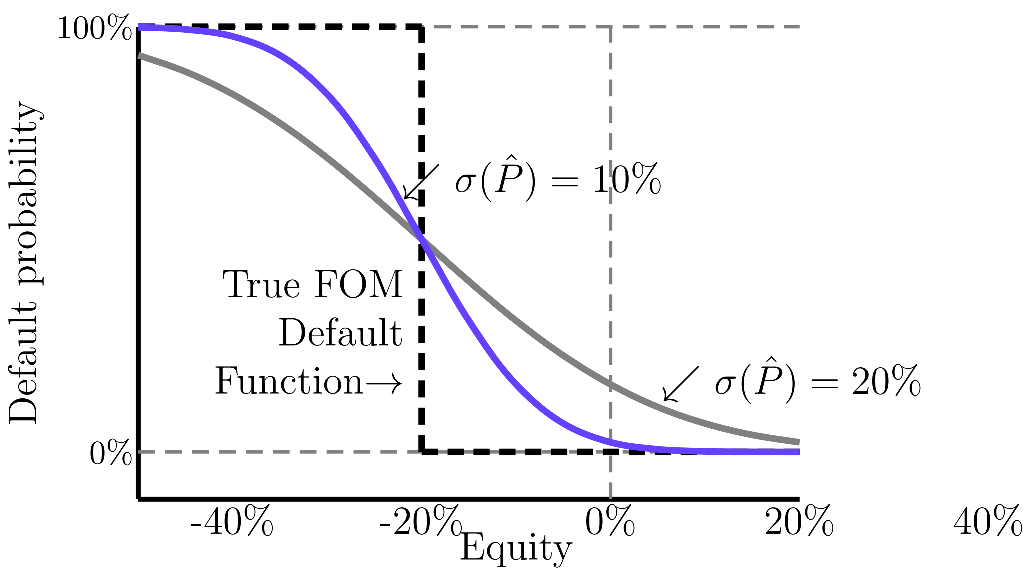

Figure 5. Estimation Issues with the Frictionless Model. If house prices are measured with error, then researchers will observe borrowers who default with equity below the FOM's true threshhold level.

Mortgage Hardship Reason	**Share**
Curtailment of Income	41.8%
Unemployment	*17.4%*
Curtailment of Income	*22.0%*
Business Failure	*2.3%*
Death or Illness in the Family	23.2%
Extreme Financial Stress Other than Loss of Income	14.4%
Excessive Obligation	*11.5%*
Extreme Hardship	*2.5%*
Payment Adjustment	*0.4%*
Marital Difficulties	7.6%
Property Problem or Casualty Loss	1.9%
Inability to Sell or Rent Property	1.3%
Employment Transfer or Military Service	0.8%
All Other Reasons	9.0%

Table 2. Reasons Given by Borrowers for Default. The sample consists of delinquent Freddie Mac borrowers who made contact with their servicer. *Source:* Cutts and Merrill (2008).

Panel A: Comparing Effects of Lower CLTV with Payment Reduction via Lower Interest Rates

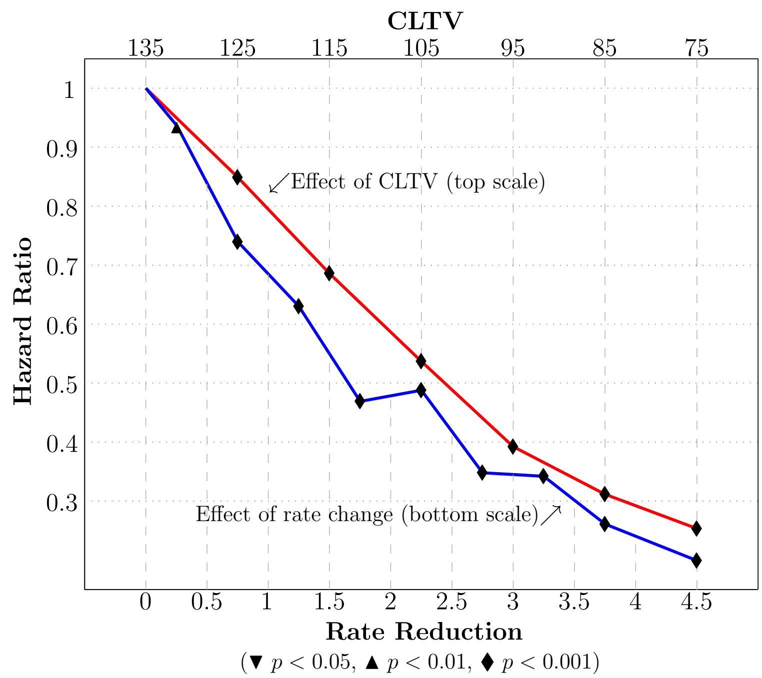

Panel B: Coefficients on Control Variables

5/1	1.369***	Prepaym. penalty active	1.073***	Purpose=Cashout Refi	0.948**
	(0.0297)		(0.0108)		(0.0195)
3/1	1.652***	Log(loan amount)	1.190***	Not owner-occupied	1.061
	(0.0517)		(0.0391)		(0.064)
Initial int. rate	2.516***	Origination LTV	1.043***	Condo	0.840***
	(0.48)		(0.00536)		(0.0415)
(Initial int. rate)^2	0.969*	(Origination LTV)^2	1.000***	12-month HPA	0.983***
	(0.0137)		(2.7E-05)		(0.00168)
FICO/100	0.543***	Full doc.	0.595***	Unemp. rate	1.011*
	(0.0127)		(0.0148)		(0.00434)
Open liens = 2	1.231***	No doc.	1.087***	6mon Δ(Unemp. rate)	1.006
	(0.027)		(0.0172)		(0.00861)
Open liens $\geq$ 3	0.913	Purpose = Refi	0.849***	30(year FRM rate	1.068*
	(0.0515)		(0.0249)		(0.0285)
Baseline hazard strata	Closing q.	Observations	4,790,556		
State dummies	✓	# Loans	138,077		
Log Likelihood	-499283	# Incidents	55,238		

Exponentiated coefficients; Standard errors (clustered at state level) in parentheses
* $p < 0.05$ ** $p < 0.01$ *** $p < 0.001$

Figure 6. Estimated Default Hazard. The sample consists of Alt-A mortgages originated in 2005 and 2006. The top panel shows that reducing the size of the monthly payment has an effect that is comparable in magnitude to a dramatic reduction in the balance of the principal. The lower panel displays various coefficients estimated from the model. *Source:* Fuster and Willen (2015)

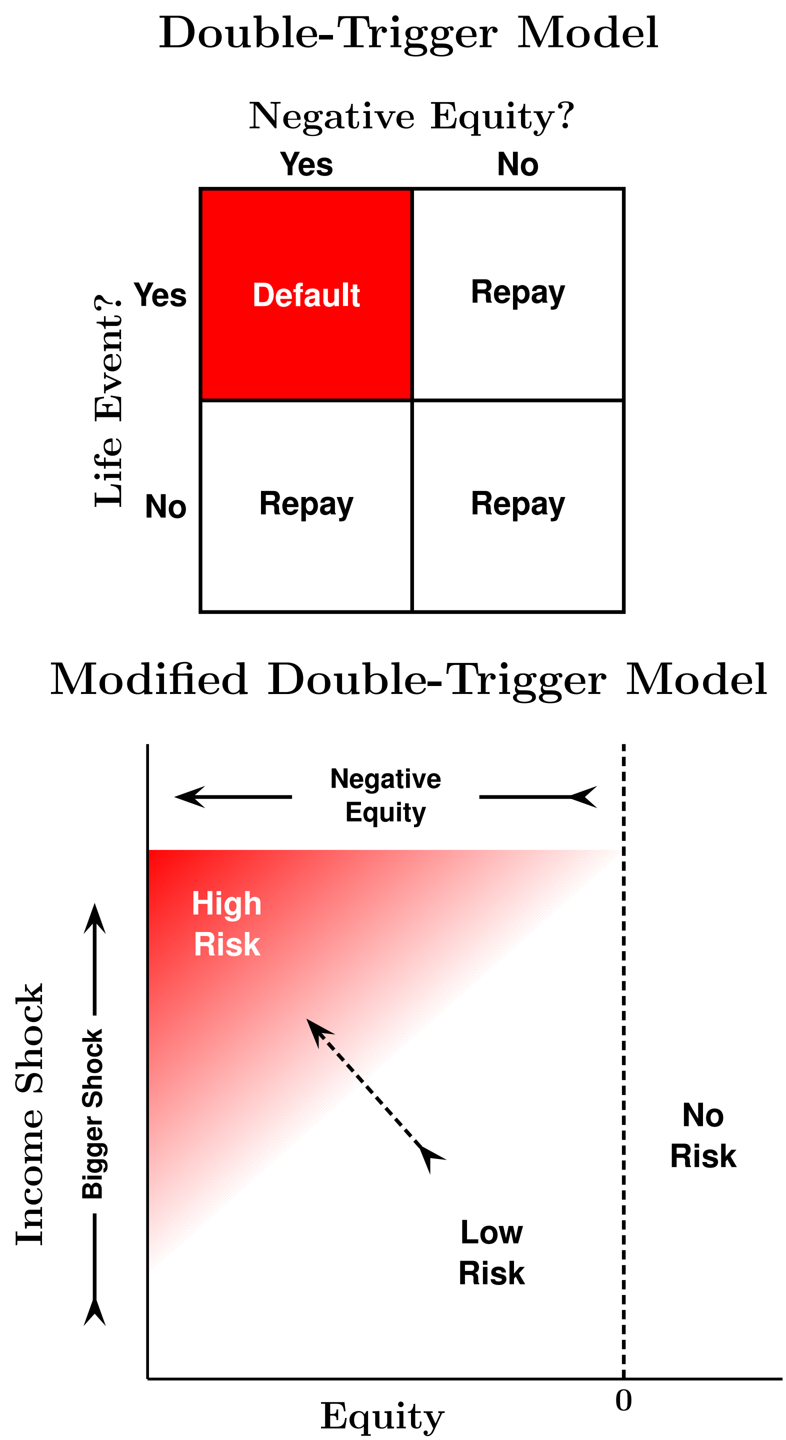

Figure 7. The Double-Trigger Model. The top panel shows the basic double-trigger model, which proposes that borrowers default when they suffer a life event that makes the monthly payment unaffordable and they can't sell their houses because of negative equity. The bottom panel shows a more realistic version of the model in which life events and negative equity interact.

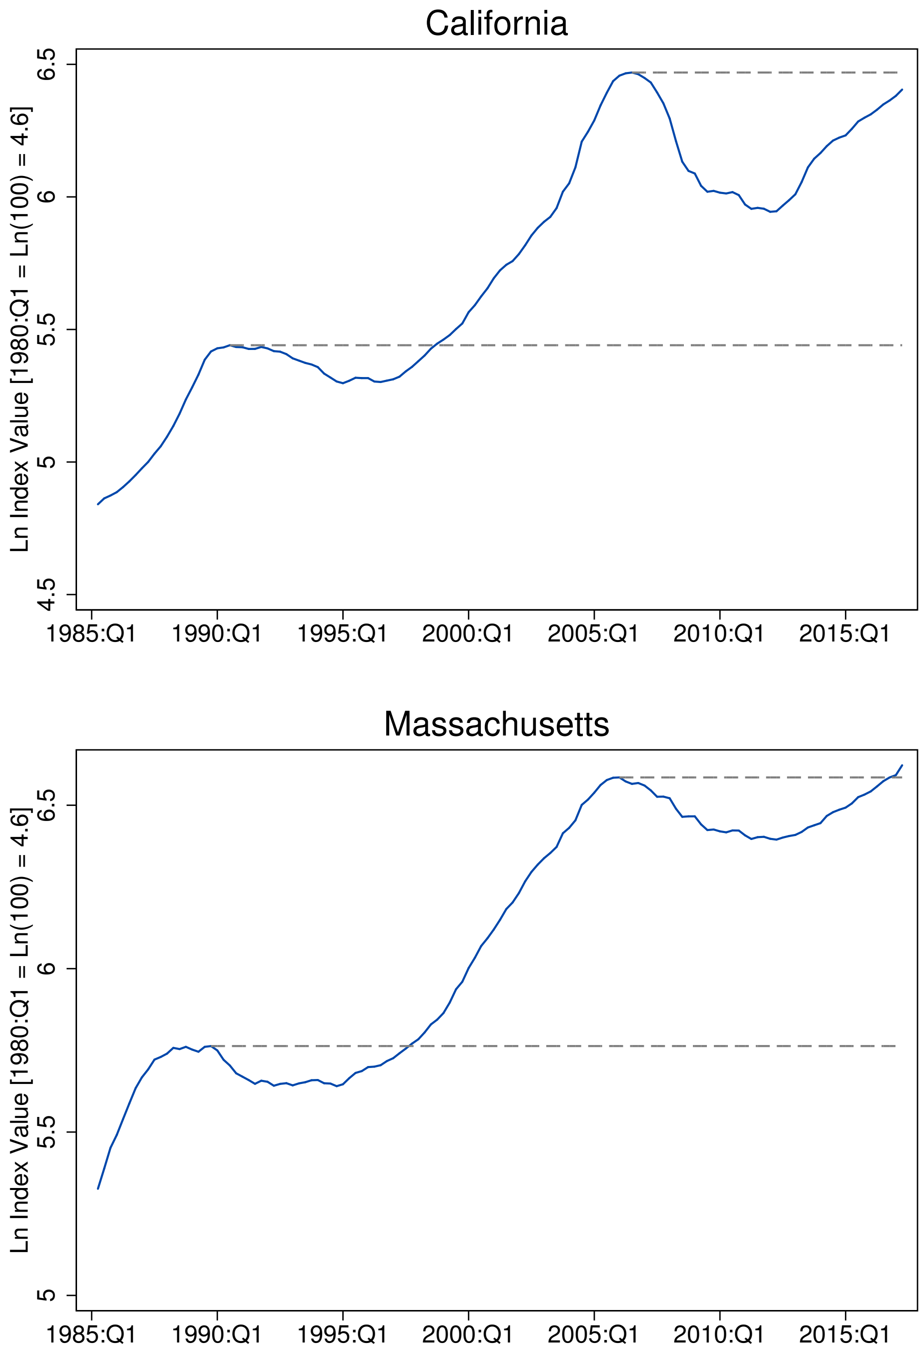

Figure 8. Nominal House-Price Indexes in California and Massachusetts. *Note:* The horizontal dashed lines denote two price peaks for each state. The first peak was reached during the coastal house-price cycle of the late 1980s and early 1990s, and the second was reached during the later national housing cycle. *Source:* Federal Housing Finance Agency.

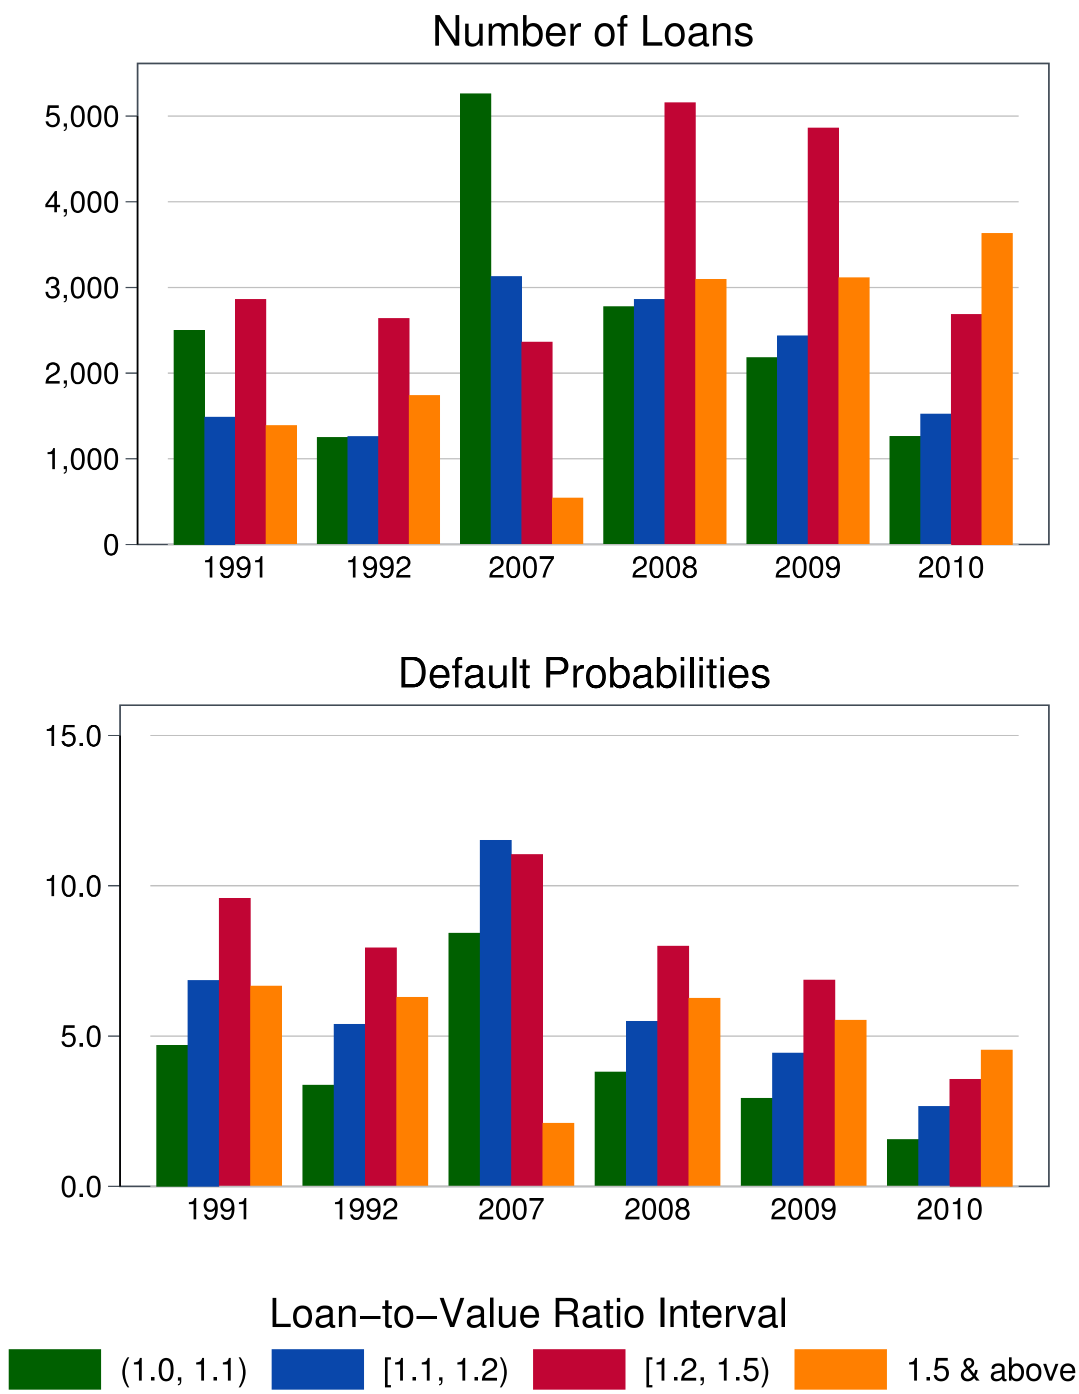

Figure 9. Massachusetts Default Probabilities in Two Housing Cycles. *Note:* The default probabilities in the lower panel refer to the fraction of mortgages at the end of the fourth quarter of the year that default at some point during the following three years. *Source:* Authors' calculations using public deed-registry data from the Warren Group.

7/18/2008	17.1
10/27/2008	17.3
11/17/2008	19.4
10/28/2009	20.2
2/22/2010	23.3
12/28/2010	22.2
8/24/2012	22.3

Table 3. J.P. Morgan Forecasts for the Cumulative Losses on Deals in the 2006-1 ABX. *Note:* For each date, the table displays the forecast that corresponded to the best house price outcome, as those, ironically, were most realistic *ex post*. The table shows that researchers were not surprised by the behavior of subprime borrowers during the crisis.

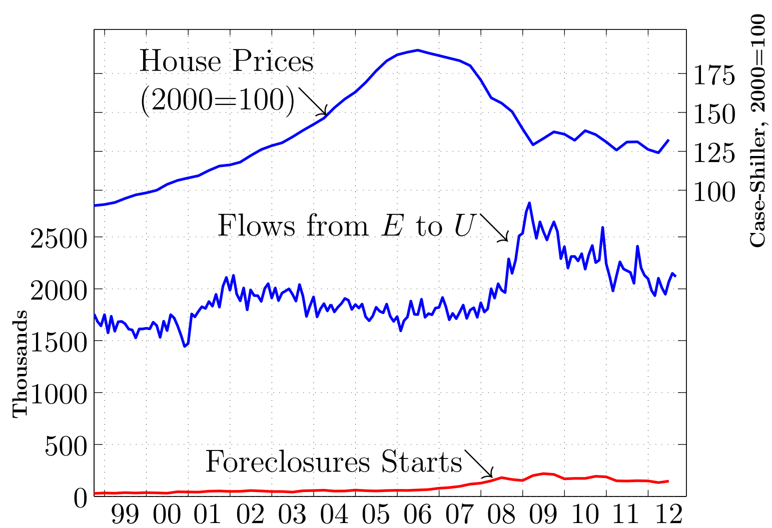

Figure 10. Flows from Employment (E) into Unemployment (U) versus Foreclosure Starts and House Prices. *Sources:* Current Population Survey, Mortgage Bankers Association, and Case-Shiller.

	Unemployment Rate (%)	CLTV (%) > 120	101–120	81–100	≤ 80
Owner-Occupied	> 12.0	2.21	1.01	0.61	0.23
	10.1-12.0	1.77	0.90	0.55	0.18
	8.1-10.0	1.81	0.83	0.52	0.22
	≤ 8.0	0.86	0.66	0.51	0.24
Non Owner-Occupied	> 12.0	1.16	0.48	0.18	0.13
	10.1-12.0	1.20	0.54	0.52	0.16
	8.1-10.0	1.06	0.65	0.36	0.17
	≤ 8.0	0.88	0.59	0.36	0.19

Table 4. Monthly Default Transition Rates for Prime Mortgage Borrowers. Compare to the "Modified Double-Trigger Model" in the lower panel of Figure 7. *Sources:* Goodman et al. (2010), using data from LoanPerformance, Bureau of Labor Statistics, and Amherst Securities.

		2007	2008	2009	2010	All
Fixed-Rate Mortgage Share		38%	48%	62%	74%	59%
Prior to delinquency spell that led to foreclosure...						
% of loans with...	Reset	18%	20%	18%	11%	17%
	Payment increase	12%	17%	11%	9%	12%
	Payment reduction	0%	0%	4%	8%	4%
	No change since orig.	88%	82%	85%	83%	84%
Private-Label		68%	54%	37%	23%	41%
# obs in thous.		374	641	874	756	2,646

Table 5. Payment Changes and Defaults among Various Mortgage Types. *Note:* A small minority of borrowers who eventually lost their homes to foreclosure experienced a payment increase before they first became delinquent. Payment increases preceded initial delinquency for only 12 percent of borrowers. Eighty-four percent of borrowers who eventually lost their homes were making the same payment at the time of initial delinquency as when they first took out their loans. *Source:* Foote, Gerardi, and Willen (2012) using data from Lender Processing Services.

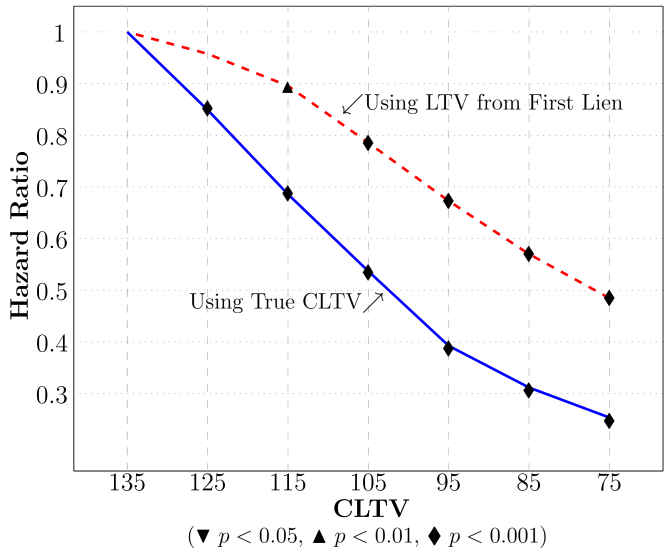

Figure 11. The Benefits of Using a Matched Sample of Public Records and Loan-Level Data. *Note:* The figure shows the relative effect of different levels of equity on the default hazard, using equity measured with only first-lien data and a matched sample that allows the calculation of a combined loan-to-value ratio (CLTV). *Source:* Fuster and Willen (2015).

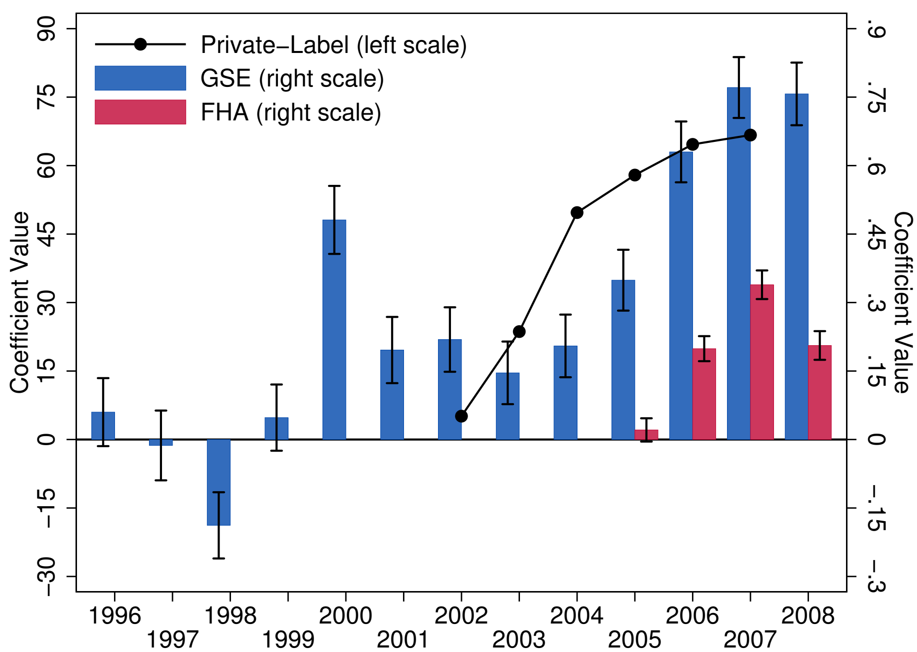

Figure 12. Cohort Effects in Default Regressions. *Sources:* Demyanyk and Van Hemert (2011) for private-label coefficients and Lam, Dunsky, and Kelly (2013) for GSE and FHA coefficients.

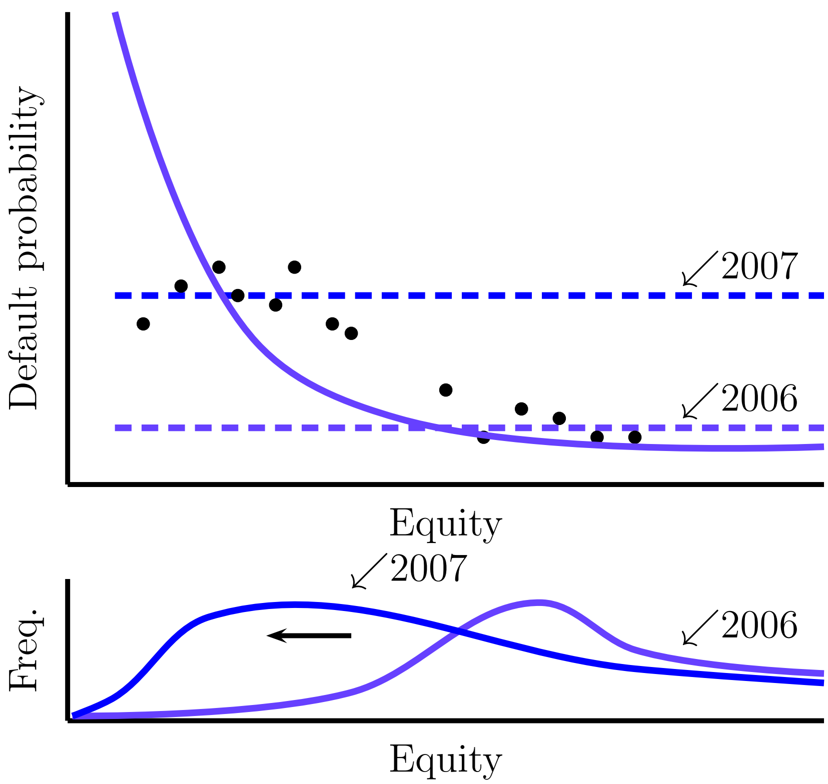

Figure 13. The Identification Problem in Estimating the Default Function. *Note:* According to the "true" model, default is independent of the level of equity, so the default functions in the upper panel are horizontal. But an exogenous increase in foreclosures could lead to a fall in house prices, so that the distribution of equity in the lower panel shifts to the left. Researchers could mistakenly conclude that reduced equity causes foreclosure, as suggested by the downward-sloping default function in the upper panel.

	(1) **"Can Pay"** $c < y - m$		(2) $c > y - m > c(VA)$		(3) **"Can't Pay"** $y - m < c(VA)$		(4) Total
	#	share	#	share	#	share	#
	(i)	(ii)=(i)/(vii)	(iii)	(iv)=(iii)/(vii)	(v)	(vi)=(v)/(vii)	(vii)
Panel A: All Households							
Default	74	0.377	65	0.333	57	0.291	196
Population	5,173	0.699	1,704	0.230	531	0.072	7,404
Default Rate (# Def./# Pop.)	0.014		0.038		0.107		0.027
Panel B: Households w/LTV > 90							
Default	47	0.409	41	0.352	28	0.239	115
Population	1,117	0.664	428	0.254	140	0.083	1,684
Default Rate	0.042		0.095		0.197		0.069
Panel C: Households w/LTV < 90							
Default	27	0.330	25	0.306	29	0.364	81
Population	4,056	0.709	1,277	0.223	391	0.068	5,720
Default Rate	0.007		0.019		0.075		0.014

Table 6. Can-Pay and Can't-Pay Defaults in the Panel Survey of Income Dynamics (PSID). *Note:* This table displays statistics on strategic-default measures calculated from the PSID. Income y is defined as average monthly after-tax family income. If the head of household is unemployed as of the survey date, then the head's labor earnings for that month are set to zero (likewise for the spouse). If the head is recently divorced, then spousal labor earnings are set to zero. Consumption c is defined as the monthly average of reported expenditures, and m is the monthly mortgage payment across all mortgages, plus associated property taxes and insurance. Department of Veterans Affairs (VA) subsistence consumption $c(VA)$ is defined using the VA residual-income concept, which is based on the household's region and number of children. *Source:* Gerardi et al. (2017).

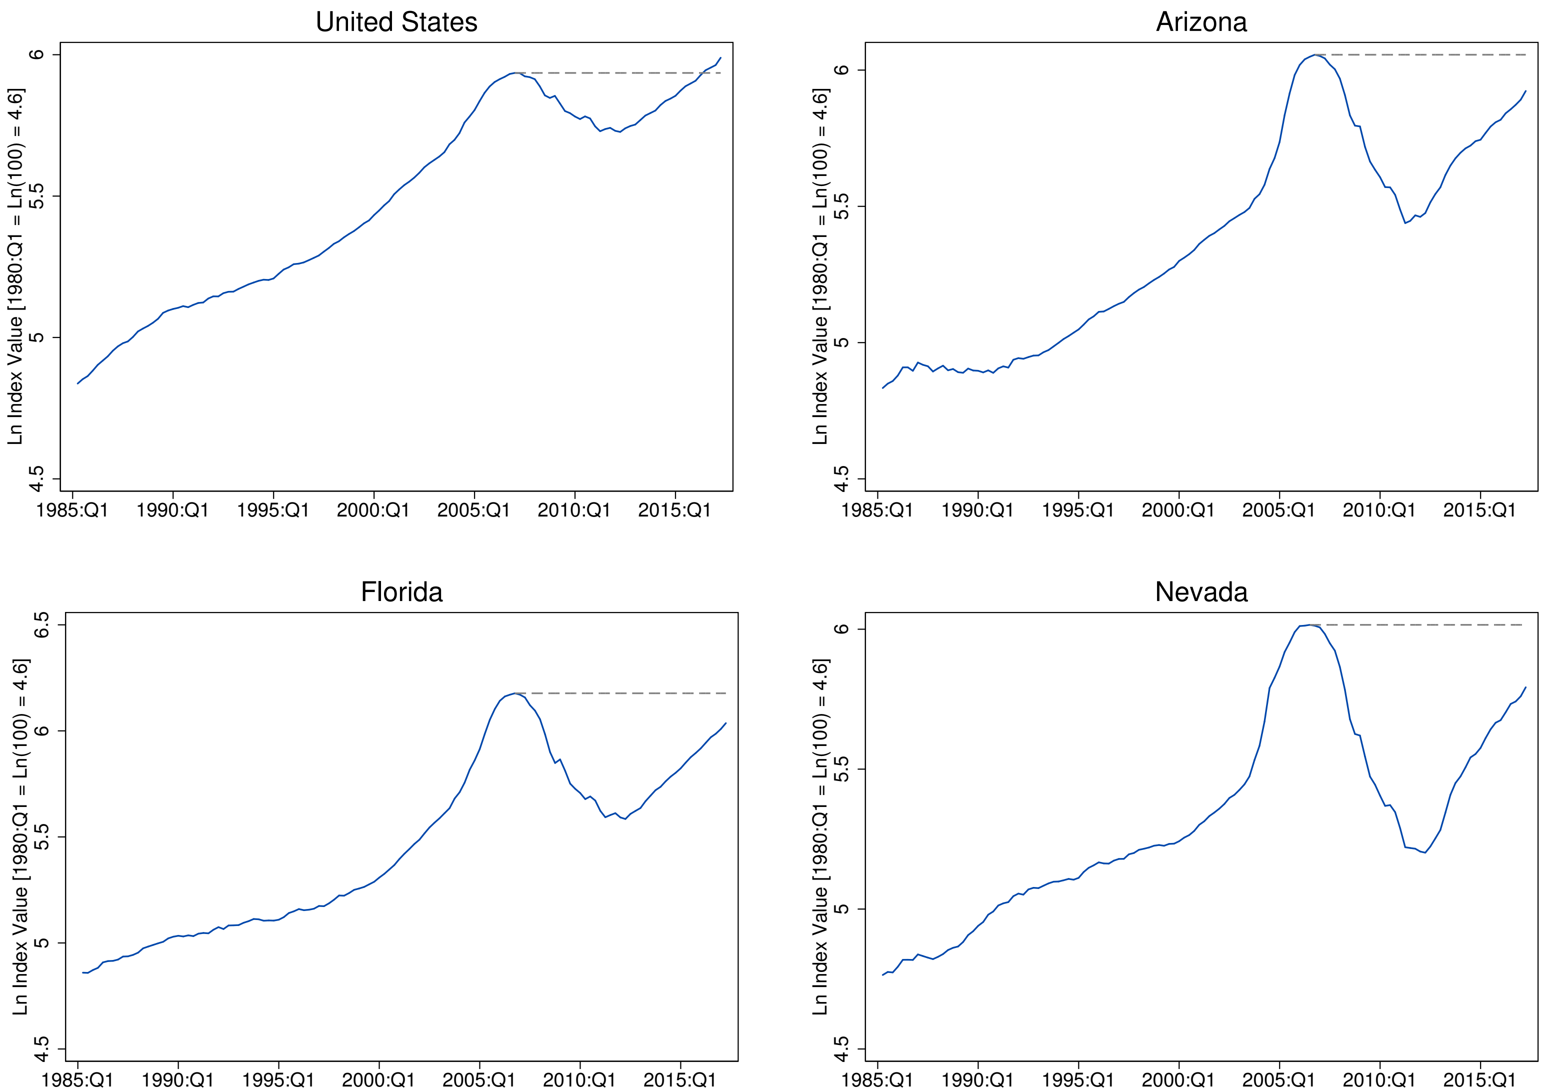

Figure 14. Nominal House Prices in the United States and Three "Sand States." See Figure 8 for price data for California, the fourth Sand State. *Source:* Federal Housing Finance Agency.

References

Adelino, Manuel, Kristopher Gerardi, and Paul S. Willen. 2013. "Why Don't Lenders Renegotiate More Home Mortgages? Redefaults, Self-Cures and Securitization." *Journal of Monetary Economics* 60(7): 835–853.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino. 2016. "Loan Originations and Defaults in the Mortgage Crisis: The Role of the Middle Class." *Review of Financial Studies* 29(7): 1635.

Adelson, Mark. 2003. "A Journey to the Alt-A Zone: A Brief Primer on Alt-A Mortgage Loans." Nomura Fixed Income Research. Available at http://www.markadelson.com/pubs/Journey_to_the_Alt-A_Zone.pdf.

Albanesi, Stefania, Giacomo De Giorgi, and Jaromir Nosal. 2017. "Credit Growth and the Financial Crisis: A New Narrative." Working Paper No. 23740. Cambridge, MA: National Bureau of Economic Research.

Anenberg, Eliot, and Edward Kung. 2014. "Estimates of the Size and Source of Price Declines Due to Nearby Foreclosures." *American Economic Review* 104(8): 2527–2551.

Aragon, Diego, Andrew Caplin, Sumit Chopra, John V. Leahy, Yann LeCun, Marco Scoffier, and Joseph Tracy. 2010. "Reassessing FHA Risk." Working Paper No. 15802. Cambridge, MA: National Bureau of Economic Research.

Avery, Robert B., Raphael W. Bostic, Paul S. Calem, and Glenn B. Canner. 1996. "Credit Risk, Credit Scoring, and the Performance of Home Mortgages." *Federal Reserve Bulletin* 82(July): 621–648.

Bhutta, Neil, Jane Dokko, and Hui Shan. 2017. "Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust." Forthcoming in the *Journal of Finance*.

Bubb, Ryan, and Alex Kaufman. 2014. "Securitization and Moral Hazard: Evidence from Credit Score Cutoff Rules." *Journal of Monetary Economics* 63(April): 1–18.

Campbell, John Y., and João F. Cocco. 2015. "A Model of Mortgage Default." *Journal of Finance* 70(4): 1495–1554.

Campbell, John Y, Stefano Giglio, and Parag Pathak. 2011. "Forced Sales and House Prices." *American Economic Review* 101(5): 2108–2131.

Corradin, Stefano. 2014. "Household Leverage." *Journal of Money, Credit and Banking* 46(4): 567–613.

Cutts, Amy Crews., and William A. Merrill. 2008. "Interventions in Mortgage Default: Policies and Practices to Prevent Home Loss and Lower Costs." In *Borrowing to Live: Consumer and Mortgage Credit Revisited*, Nicolas P. Retsinas and Eric S. Belsky, eds., 203–254. Washington, DC: Brookings Institution Press.

Demyanyk, Yuliya, and Otto Van Hemert. 2011. "Understanding the Subprime Mortgage Crisis." *Review of Financial Studies* 24(6): 1848–1880.

Deng, Yongheng, John M. Quigley, and Robert Van Order. 2000. "Mortgage Terminations, Heterogeneity and the Exercise of Mortgage Options." *Econometrica* 68(2): 275–307.

Downing, Chris, Richard Stanton, and Nancy Wallace. 2005. "An Empirical Test of a Two-Factor Mortgage Valuation Model: How Much Do House Prices Matter?" *Real Estate Economics* 33(4): 681–710.

Eberly, Janice C., and Arvind Krishnamurthy. 2014. "Efficient Credit Policies in a Housing Debt Crisis." *Brookings Papers on Economic Activity* 2014(2): 73–136.

Elul, Nicholas S. Souleles Souphala Chomsisengphet Dennis Glennon, Ronel, and Robert Hunt. 2010. "What 'Triggers' Mortgage Default?" *American Economic Review: Papers & Proceedings* 100(2): 490–494.

Epperson, James F., James B. Kau, Donald C. Keenan, and Walter J. Muller III. 1985. "Pricing Default Risk in Mortgages." *Real Estate Economics* 13(3): 261–272.

Experian, and Oliver Wyman. 2009. "Understanding Strategic Default in Mortgages Part I." *Experian-Oliver Wyman Market Intelligence Report.* 2009 Topical Report Series. Available at https://www.experian.com/assets/decision-analytics/reports/strategic-default-report-1-2009.pdf.

Foote, Christopher, Kristopher S. Gerardi, Lorenz Goette, and Paul S. Willen. 2009. "Reducing Foreclosures: No Easy Answers." In *NBER Macroeconomics Annual 2009*, Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds., vol. 24, 89–138. Chicago, IL: University of Chicago Press.

Foote, Christopher L., Kristopher Gerardi, and Paul S. Willen. 2008. "Negative Equity and Foreclosure: Theory and Evidence." *Journal of Urban Economics* 64(2): 234–245.

Foote, Christopher L., Kristopher S. Gerardi, and Paul S. Willen. 2012. "Why Did So Many People Make So Many *Ex Post* Bad Decisions: The Causes of the Foreclosure Crisis." In *Rethinking the Financial Crisis*, Alan S. Blinder, Andrew W. Lo, and Robert M. Solow, eds., 136–186. The Russell Sage Foundation and the Century Foundation: New York, N.Y. and Washington, D.C.

Foote, Christopher L., Lara Loewenstein, and Paul S. Willen. 2016. "Cross-Sectional Patterns of Mortgage Debt During the Housing Boom: Evidence and Implications." Federal Reserve Bank of Boston Working Paper No. 16-12. Available at https://www.bostonfed.org/-/media/Documents/Workingpapers/PDF/wp1612.pdf.

Foster, Chester, and Robert Van Order. 1984. "An Option-Based Model of Mortgage Default." *Housing Finance Review* 3(4): 351–372.

Frankel, Jeffrey. 2017. "Why Financial Markets Underestimate Risk." *Project Syndicate* Editorial, Sept 25. Available at https://www.project-syndicate.org/commentary/financial-markets-crisis-risk-by-jeffrey-frankel-2017-09.

Fuster, Andreas, and Paul Willen. 2015. "Payment Shocks, Negative Equity and Mortgage Default." Federal Reserve Bank of New York Staff Reports No. 582. Available at https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr582.pdf.

Galster, George C., Peter Tatian, and Robin Smith. 1999. "The Impact of Neighbors Who Use Section 8 Certificates on Property Values." *Housing Policy Debate* 10(4): 879–917.

Gerardi, Kristopher, Kyle F. Herkenhoff, Lee E. Ohanian, and Paul S. Willen. 2017. "Can't Pay or Won't Pay? Unemployment, Negative Equity, and Strategic Default." Working Paper No. 21630. Cambridge, MA: National Bureau of Economic Research. Forthcoming in *Review of Financial Studies*.

Gerardi, Kristopher, Lauren Lambie-Hanson, and Paul S. Willen. 2013. "Do Borrower Rights Improve Borrower Outcomes? Evidence from the Foreclosure Process." *Journal of Urban Economics* 73(1): 1–17.

Gerardi, Kristopher, Eric Rosenblatt, Paul Willen, and Vincent Yao. 2015. "Foreclosure Externalities: New Evidence." *Journal of Urban Economics* 87(May): 42–56.

Gerardi, Kristopher, Adam Hale Shapiro, and Paul Willen. 2008. "Subprime Outcomes: Risky Mortgages, Homeownership Experiences, and Foreclosures." Federal Reserve Bank of Boston Working Paper No. 07-15. Available at https://www.bostonfed.org/-/media/Documents/Workingpapers/PDF/wp0715.pdf.

Gerardi, Kristopher, Paul Willen, Shane M. Sherlund, and Andreas Lehnert. 2008. "Making Sense of the Subprime Crisis." *Brookings Papers on Economic Activity* 2008(2): 69–159.

Gerardi, Kristopher S., Christopher L. Foote, and Paul S. Willen. 2011. "Reasonable People Did Disagree: Optimism and Pessimism about the U.S. Housing Market Before the Crash." In *The American Mortgage System: Crisis and Reform*, Marvin Smith and Susan M. Wachter, eds., 26–59. Philadelphia, PA: University of Pennsylvania Press.

Ghent, Andra C., and Marianna Kudlyak. 2011. "Recourse and Residential Mortgage Default: Evidence from U.S. States." *Review of Financial Studies* 24(9): 3139–3186.

Glaeser, Edward, Bruce Sacerdote, and Jose Scheinkman. 1996. "Crime And Social Interactions." *Quarterly Journal of Economics* 111(2): 507–548.

Goldberg, Lawrence, and Charles A. Capone, Jr. 1998. "Multifamily Mortgage Credit Risk: Lessons from Recent History." *Cityscape* 4(1): 93–113.

Goodman, Laurie S., Roger Ashworth, Brian Landy, and Ke Yin. 2010. "Negative Equity Trumps Unemployment in Predicting Defaults." *Journal of Fixed Income* 19(4): 67–72.

Goodstein, Ryan, Paul Hanouna, Carlos Ramirez, and Christof Stahel. 2012. "Are Foreclosures Contagious?" SSRN Working Paper. Available at https://dx.doi.org/10.2139/ssrn.2024794.

Graham, Bryan S., and Jinyong Hahn. 2005. "Identification and Estimation of the Linear-in-Means Model of Social Interactions." *Economics Letters* 88(1): 1–6.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2013. "The Determinants of Attitudes towards Strategic Default on Mortgages." *Journal of Finance* 68(4): 1473–1515.

Guren, Adam, and Tim McQuade. 2016. "How Do Foreclosures Exacerbate Housing Downturns?" Working Paper. Available at http://web.stanford.edu/%7Etmcquade/foreclosures.pdf.

Gyourko, Joseph, and Joseph Tracy. 2014. "Reconciling Theory and Empirics on the Role of Unemployment in Mortgage Default." *Journal of Urban Economics* 80(March): 87–96.

Harding, John P., Eric Rosenblatt, and Vincent W. Yao. 2009. "The Contagion Effect of Foreclosed Properties." *Journal of Urban Economics* 66(3): 164–178.

Hartley, Daniel. 2014. "The Effect of Foreclosures on Nearby Housing Prices: Supply or Dis-amenity?" *Regional Science and Urban Economics* 49(November): 108–117.

Herzog, John P., and James S. Earley. 1970. *Home Mortgage Delinquency and Foreclosure*. Cambridge, MA: National Bureau of Economic Research.

Immergluck, Dan, and Geoff Smith. 2006. "The External Costs of Foreclosure: The Impact of Single-Family Mortgage Foreclosures on Property Values." *Housing Policy Debate* 17(1): 57–79.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante. 2017. "The Housing Boom and Bust: Model Meets Evidence." Working Paper No. 23694. Cambridge, MA: National Bureau of Economic Research.

Kau, James B., and Donald C. Keenan. 1999. "Patterns of Rational Default." *Regional Science and Urban Economics* 29(6): 765–785.

Kau, James B., Donald C. Keenan, and Taewon Kim. 1994. "Default Probabilities for Mortgages." *Journal of Urban Economics* 35(3): 278–296.

Kau, J.B., D.C. Keenan, W.J. Muller, and J.F. Epperson. 1992. "A Generalized Valuation Model for Fixed-Rate Residential Mortgages." *Journal of Money, Credit and Banking* 24(3): 279–299.

Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig. 2010. "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans." *Quarterly Journal of Economics* 125(1): 307–362.

Keys, Benjamin J., Tomasz Piskorski, Amit Seru, and Vikrant Vig. 2012. "Mortgage Financing in the Housing Boom and Bust." In *Housing and the Financial Crisis*, Edward Glaeser and Todd Sinai, eds. Chicago, IL: University of Chicago Press.

Kimball, Zachary K., and Paul S. Willen. 2012. "U.S. Mortgage and Foreclosure Law." In *The New Palgrave Dictionary of Economics*, Steven N. Durlauf and Lawrence E. Blume, eds. Basingstoke, UK: Palgrave Macmillan.

Knoll, Michael S. 2002. "Put-Call Parity and the Law." *Cardozo Law Review* 24(1): 61–92.

Lam, Ken, Robert M. Dunsky, and Austin Kelly. 2013. "Impacts of Down Payment Underwriting Standards on Loan Performance: Evidence from the GSEs and FHA Portfolios." Federal Housing Finance Agency Working Paper No. 13-3. Available at https://www.fhfa.gov/policyprogramsresearch/research/paperdocuments/2013-12_workingpaper_13-3-508.pdf.

Lambie-Hanson, Lauren. 2015. "When Does Delinquency Result in Neglect? Mortgage Distress and Property Maintenance." *Journal of Urban Economics* 90(November): 1–16.

Laufer, Steven. 2017. "Equity Extraction and Mortgage Default." Forthcoming in the *Review of Economic Dynamics*.

Lin, Zhenguo, Eric Rosenblatt, and Vincent W. Yao. 2009. "Spillover Effects of Foreclosures on Neighborhood Property Values." *Journal of Real Estate Finance and Economics* 38(4): 387–407.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias on Predicting Future Utility." *Quarterly Journal of Economics* 118(4): 1209–1248.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies* 60(3): 531–542.

Mayer, Christopher, Karen Pence, and Shane M. Sherlund. 2009. "The Rise in Mortgage Defaults." *Journal of Economic Perspectives* 23(1): 27–50.

Mian, Atif. 2010. "Comment on 'Reducing Foreclosures: No Easy Answers'." In *NBER Macroeconomics Annual 2009*, Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds., vol. 24, 149–156. Chicago, IL: University of Chicago Press.

Mian, Atif, and Amir Sufi. 2009. "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis." *Quarterly Journal of Economics* 124(4): 1449–1496.

Mian, Atif, and Amir Sufi. 2010. "Household Leverage and the Recession of 2007–09." *IMF Economic Review* 58(1): 74–117.

Rogers, Wiliam H., and William Winter. 2009. "The Impact of Foreclosures on Neighboring Housing Sales." *Journal of Real Estate Research* 31(4): 455–479.

Schelkle, Thomas. 2014. "Mortgage Default during the U.S. Mortgage Crisis." University of Cologne, Department of Economics Working Paper No. 72. Available at http://ockenfels.uni-koeln.de/fileadmin/wiso_fak/stawi-ockenfels/pdf/wp_series_download/wp0072.pdf.

Schuetz, Jenny, Vicki Been, and Ingrid Gould Ellen. 2008. "Neighborhood Effects of Concentrated Mortgage Foreclosures." *Journal of Housing Economics* 17(4): 306–319.

Swagel, Phillip. 2009. "The Financial Crisis: An Inside View." *Brookings Papers on Economic Activity* 2009(1): 1–78.

Tirupattur, Vishwanath, Oliver Chang, and James Egan. 2010. "Understanding Strategic Defaults." Morgan Stanley Research. ABS Market Insights. Available at https://www.transunion.com/docs/rev/business/financialservices/FS_MorganStanleyStrategicDefaultResearch.pdf.

Tracy, Joseph, and Joshua Wright. 2012. "Payment Changes and Default Risk: The Impact of Refinancing on Expected Credit Losses." Federal Reserve Bank of New York Staff Report No. 562. Available at https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr562.pdf.

Vandell, Kerry D. 1995. "How Ruthless Is Mortgage Default? A Review and Synthesis of the Evidence." *Journal of Housing Research* 6(2): 245–264.

Wallison, Peter J., and Edward J. Pinto. 2012. "Free Fall: How Government Policies Brought Down the Housing Market." Financial Services Outlook, American Enterprise Institute. April 26. Available at http://www.aei.org/publication/free-fall-how-government-policies-brought-down-the-housing-market/.

White, Brent T. 2010. "Underwater and Not Walking Away: Shame, Fear and the Social Management of the Housing Crisis." *Wake Forest Law Review* 45: 971–1023.