

Mortgage Default during the U.S. Mortgage Crisis*

Thomas Schelkle[†]

This Version: May 16, 2018

First Version: October 2011

Abstract

Which theory can quantitatively explain the rise in mortgage defaults during the U.S. mortgage crisis? This paper finds that the double-trigger hypothesis, which attributes mortgage default to the joint occurrence of negative equity and a life event such as unemployment, is consistent with the evidence. By contrast, a traditional frictionless default model strongly overpredicts the increase in default rates. This paper provides micro-foundations for double-trigger behavior in a model where unemployment causes liquidity problems for the borrower. This framework implies that mortgage crises may be mitigated at a lower cost by bailing out borrowers instead of lenders.

JEL codes: E21, G21, D11

Keywords: Mortgage default, mortgage crisis, house prices, negative equity

*This paper is a revised version of the first chapter of my Ph.D. thesis at the London School of Economics. I am grateful to Francesco Caselli and Rachel Ngai for many helpful discussions and much valuable advice. I also thank the editor and two anonymous referees, as well as Wouter den Haan, Michael Krause, Alexander Ludwig, Albert Marcet, Xavier Mateos-Planas, Stephan Seiler, Jonathan Temple, and seminar participants at RWTH Aachen, University of Cologne, LSE, University of Manchester, Royal Holloway, ESEM 2011 in Oslo, RES 2012 in Cambridge, SED 2012 in Limassol and GEA 2012 in Cologne for valuable comments. Special thanks go to Jason Friend for help with the data. A previous version of this paper was circulated under the title “Simulating Mortgage Default”.

[†]University of Cologne; Center for Macroeconomic Research; Albertus-Magnus-Platz; 50923 Köln; Germany; Email: schelkle@wiso.uni-koeln.de

1 Introduction

Since the collapse of the house price boom in the United States, residential mortgage delinquencies of both prime and subprime loans have increased substantially. The resulting losses of mortgage-backed securities marked the start of the recent financial and economic crisis. These events highlight the importance of understanding the economic mechanisms that trigger mortgage default and the rise in default rates. Insights into these issues may serve to inform political debates on how to prevent future mortgage crises or mitigate those that have already started. This paper contributes to this research agenda by investigating what type of theoretical mortgage default model can quantitatively explain the observed rise in default rates in the United States between 2002 and 2010.

The paper considers the two main competing theories of mortgage default: the frictionless option-theoretic model and the double-trigger hypothesis. The traditional frictionless literature assumes that borrowers “ruthlessly” default on their mortgages to maximize their financial wealth, cf. for example Kau, Keenan, and Kim (1994), Kau, Keenan, Muller, and Epperson (1992, 1995) and the surveys of Quercia and Stegman (1992), Kau and Keenan (1995) and Vandell (1995). In this framework, negative equity is a necessary but not sufficient, condition for default. Instead, there exists a threshold level of negative equity such that a rational wealth-maximizing agent will exercise the default option. This theory is frictionless in the sense of assuming a perfect credit market for unsecured credit, such that default is unaffected by income fluctuations or liquidity problems.

The other main theory on mortgage default is the double-trigger hypothesis. This theory agrees that negative equity is a necessary condition for default but attributes default to the joint occurrence of negative equity and a life event such as unemployment or divorce. The double-trigger hypothesis is well known among mortgage researchers, cf. the discussions by Gerardi, Shapiro, and Willen (2007), Foote, Gerardi, and Willen (2008), Foote, Gerardi, Goette, and Willen (2009) and the survey of Foote and Willen (2018). However, in contrast to the well established and formalized frictionless default theory, the double-trigger hypothesis is usually only discussed verbally or with stylized models.

These two microeconomic theories are assessed quantitatively using data on prime fixed-rate mortgages with high initial loan-to-value ratios (above 95%).¹ The test exploits variation across cohorts of loans originated between 2002 and 2008. In the data, cohorts of mortgage borrowers who experienced a more adverse path of average house prices, and therefore lower home equity, defaulted much more frequently. This variation in default rates across cohorts does not seem to be driven by other observable loan or borrower characteristics. *Qualitatively*, both theories are consistent with this joint variation in home equity and default rates. However, an estimation and simulation of reduced-form models of the two theories reveals important *quantitative* differences between them.

Specifically, the estimation by a simulated method of moments procedure forces both models to match the default rates of the 2002 cohort, which they can both fit well. Using the estimated parameters representing the default threshold and frequency of life events, respectively, I then assess the ability of the models to predict the default rates of the 2003 to 2008 cohorts. In this out-of-sample exercise, I find that the frictionless theory is excessively sensitive to changes in aggregate house prices and predicts an excessively strong rise in default rates across cohorts. The leftward shift in the home equity distribution caused by the observed fall in aggregate house prices strongly increases the number of borrowers with extreme levels of negative equity. This moves too many borrowers over the estimated default threshold compared to the observed increase in default rates. In contrast, the double-trigger hypothesis offers a surprisingly good fit to the observed rise in default rates. The predictions of this theory are based on borrowers who experience any level of negative equity, of which only a certain share defaults according to the estimated frequency of life events. The double-trigger theory then predicts an increase in default rates across cohorts proportional to the increase in the number of borrowers with any level of negative equity. Observed default rates exhibit such a pattern.

In the second part of the paper, these insights guide the development of a parsimo-

¹There are various data reasons that make a focus on highly leveraged borrowers advantageous, as explained in detail in section 2. However, I also document that the full sample of loans with all possible initial loan-to-value ratios exhibits very similar dynamics, though there are of course level differences. In an extension, I also show explicitly that under plausible assumptions, the results of the analysis generalize to loans with lower initial loan-to-value ratios.

nious structural model of mortgage default. This dynamic stochastic partial equilibrium model micro-founds the double-trigger hypothesis. Borrowers in the model face liquidity constraints and idiosyncratic unemployment shocks such that unemployed borrowers who have exhausted their buffer-stock savings need to make painful cuts to consumption. This magnifies the cost of servicing the mortgage such that unemployment triggers default in a negative equity situation. The model also includes a direct utility flow - from living in the bought house - that prevents employed agents from defaulting after a strong fall of house prices. These features generate double-trigger behavior in the model. The calibrated model can quantitatively explain most of the observed rise in mortgage default as a consequence of falling aggregate house prices.

The structural model is then used to formally analyze two possible mortgage-crisis mitigation policies that may help to stabilize the financial system. If the government desires to neutralize mortgage lenders' losses from default, it could either bail out the lenders or mitigate the liquidity problems of homeowners who would otherwise default. Actual stabilization policy during the crisis arguably encompassed both types of measures. For instance, the Troubled Asset Relief Program (TARP) represents an example of the former, and the Home Affordable Modification Program (HAMP) exemplifies the latter type of policy. The analysis shows that in this model where default is partly driven by liquidity problems, a policy that subsidizes homeowners is about 7 – 10 times cheaper than a bailout of lenders. Though these are only partial equilibrium results, they suggest that there is a great potential to reduce the taxpayers' cost of mitigating a mortgage crisis.

The paper relates to different strands of the theoretical and empirical literature. The structural model of the paper builds on previous theoretical work by Campbell and Cocco (2003, 2015) and Corradin (2014) who also model liquidity constraints in a mortgage framework. Similar models are used to examine bail-out guarantees or the role of mortgage product innovation and falling house prices in the mortgage crisis by Chatterjee and Eyigungor (2015), Corbae and Quintin (2015), Garriga and Schlagenhaut (2009), Jeske, Krueger, and Mitman (2013) and Laufer (2018). Compared to this work, I consider

more parsimonious models, but the contribution of my paper is to compare theoretical default models to empirical observations in much more detail. For instance, Campbell and Cocco (2015) only compare theoretical predictions on differences between fixed- and variable rate mortgages to some broad patterns in the data. Chatterjee and Eyigungor (2015) and Corbae and Quintin (2015) conduct relatively stylized computational experiments with their models and investigate only the fit to data on aggregate default or foreclosure rates. Laufer (2018) compares his model only to aggregate data on loans originated between 2002-2004 in Los Angeles county.

In contrast, I compare model predictions to detailed monthly data for several cohorts of loans from the whole country. As these cohorts experienced very different house price paths, this provides much richer variation for an empirical assessment of theoretical default models. In these comparisons, I simulate empirically accurate home equity distributions within and between cohorts and apply the theoretical default mechanisms to them. I isolate the effect of falling house prices by keeping the type of loan contract constant in both the model and the data. In contrast, comparisons to aggregate data in prior work may confound the effect of house prices with compositional effects with respect to loan cohorts or contract types. This more demanding comparison to the data reveals that including income shocks and liquidity constraints in a mortgage framework does not automatically lead to an empirically successful model. Instead, only models where agents with substantial negative equity but no liquidity problems do not find it optimal to default will truly exhibit double-trigger behavior and an accurate sensitivity of default rates to changes in aggregate house prices. This is an important finding because capturing the economic default mechanism and accurate house price sensitivity are of great relevance for the use of such models in policy or macroeconomic risk analysis. Finally, I also apply this micro-founded and empirically tested double-trigger model to formally evaluate the cost differences between a bailout of lenders and subsidies to homeowners in the context of mitigating a mortgage crisis.²

²Eberly and Krishnamurthy (2014) provide a qualitative analysis of various anti-foreclosure policies using a simple theoretical model. By contrast, this paper conducts a quantitative investigation using a more realistic calibrated model. Hatchondo, Martinez, and Sánchez (2015) and Laufer (2018) also use calibrated default models to analyze policies preventing default, such as stronger recourse and limits on

This paper is also related to a vast empirical literature that studies the determinants of mortgage default.³ This literature provides a wealth of evidence that negative equity and falling house prices are strong determinants of default. Some studies have also investigated the role of life events as triggers for default. Many studies have found that state unemployment or divorce rates are correlated with default rates. Elul, Souleles, Chomsisengphet, Glennon, and Hunt (2010) provide evidence that variables measuring illiquidity and interactions between illiquidity and negative equity significantly affect default. Gerardi, Herkenhoff, Ohanian, and Willen (2013) show that at the individual level, unemployment and income shocks increase the probability of default. My paper is motivated by these prior empirical results.⁴ However, it uses a very different methodology.

The prior literature sheds light on the theories by documenting merely statistical correlations between life events and defaults. By contrast, the analysis in this paper explicitly includes the economic structure of the competing theories.⁵ I investigate whether

the initial loan-to-value ratio or the possibility for cash-out refinancing. In contrast to these preventive policies, this paper analyzes mitigation policies that are only implemented once a crisis has already started. A major advantage of my paper is that before conducting policy analysis, I first show that the structural model can capture the empirically observed rise in default rates across loan cohorts during the crisis. Foote and Willen (2018) provide a survey of recent research on the effectiveness of different policies to prevent foreclosures.

³Studies within this extensive literature differ by research question, estimation method, data set and results. A detailed literature review that would do justice to these different contributions is unfortunately beyond the scope of this paper. The pre-crisis literature is surveyed by Quercia and Stegman (1992) and Vandell (1995), and an example is the study by Deng, Quigley, and Van Order (2000). The U.S. mortgage crisis has led an enormous increase in empirical work on mortgage default. Foote and Willen (2018) provide a survey of this recent literature. Examples of this empirical research include Amromin and Paulson (2009), Bajari, Chu, and Park (2010), Bhutta, Dokko, and Shan (2017), Demyanyk and Van Hemert (2011), Elul, Souleles, Chomsisengphet, Glennon, and Hunt (2010), Ferreira and Gyourko (2015), Foote, Gerardi, Goette, and Willen (2008), Foote, Gerardi, Goette, and Willen (2009), Foote, Gerardi, and Willen (2008), Fuster and Willen (2013), Gerardi, Herkenhoff, Ohanian, and Willen (2013), Gerardi, Lehnert, Sherlund, and Willen (2008), Gerardi, Shapiro, and Willen (2007), Ghent and Kudlyak (2011), Guiso, Sapienza, and Zingales (2013), Jagtiani and Lang (2011), Mayer, Pence, and Sherlund (2009), Mian and Sufi (2009), Palmer (2013) and Stanton and Wallace (2011), among others.

⁴Another interesting empirical fact is the great heterogeneity in default behavior among borrowers with the same level of negative equity (Quercia and Stegman 1992). A double-trigger model can rationalize this fact because life events, which are unobserved in all standard mortgage data sets, may account for the heterogeneity in default behavior of borrowers with the same level of negative equity.

⁵Bajari, Chu, and Park (2010) also estimate a model based on a simple default rule akin to my reduced-form models in section 3. However, they consider the different triggers separately, i.e. borrowers default either because of negative equity or a life event. I investigate the theoretically more appealing double-trigger hypothesis, which requires the satisfaction of both conditions. Bhutta, Dokko, and Shan (2017) focus on estimating the distribution of equity thresholds that trigger ruthless default after controlling for defaults due to liquidity reasons. By contrast, the reduced-form exercise of my paper aims at discriminating between pure forms of the two main theories and documenting - in a transparent way - what their empirical problems and merits are. In contrast to both of these papers, I also develop a structural model based on optimizing behavior.

the theories exhibit an empirically accurate sensitivity of default rates to changes in aggregate house prices. This reveals the excess sensitivity of a purely negative equity threshold-based default theory and the accurate sensitivity of a double-trigger model. Thus, the paper documents a novel set of facts on the relative merit of the two theories and their ability to explain the aggregate default dynamics during the mortgage crisis.

This paper is structured as follows. Section 2 describes the data and empirical facts regarding mortgages and house prices. Reduced-form models of the two theories are compared to the data in section 3. The structural model is developed in section 4 and parameterized in section 5. The results of the structural model are presented in section 6. The structural model is applied to policy analysis in section 7, and section 8 concludes. An online appendix contains technical details and further results.

2 Data and Empirical Facts

This section presents the data on mortgages and house prices as well as the key facts of default rates and house price paths across cohorts that the paper attempts to explain. Furthermore, it provides evidence on relatively stable loan and borrower characteristics across cohorts.

2.1 Data Sources and Main Facts

Information on mortgage contract characteristics and payment histories in the United States is based on the large loan-level data base of Lender Processing Services (LPS), also known as McDash data. I did not have access to the full loan-level data but obtained information that was aggregated from the full data base. “Aggregate” here simply means that my data contain the average value of a certain variable for all loans in the data base that satisfy a set of conditions that I can specify. The data cover the time period from January 2002 until June 2010 at a monthly frequency and the analysis is focused on loans originated between 2002 and 2008.

I restrict the sample to prime, first, fixed-rate, 30-year mortgages that have a standard

amortization schedule (i.e., are not balloon mortgages). I focus on only one mortgage type because the structural model would have to be recomputed for each different mortgage contract. The selection is motivated by the fact that these are the most common mortgage contracts. Furthermore, mortgage distress was not primarily a subprime phenomenon; it also greatly affected prime borrowers (Ferreira and Gyourko 2015). The data base contains around 23 million loans with these characteristics in the year 2010. I further focus the analysis on loans with an initial loan-to-value ratio (LTV) above 95%, which, depending on the year, represents about 20 – 30% of all outstanding loans and around 20% of newly originated loans that satisfy the above restrictions.

Looking at loans within a narrow range of LTVs allows to generate a more accurate home equity distribution in simulations of the model. This is very important due to the highly non-linear relationship between default decisions and negative equity in the theoretical models, on which the conducted empirical test also relies heavily. Furthermore, loans with a high LTV default most frequently, so it is reasonable to make them the focus of an analysis of mortgage default. However, the main reason for concentrating on this group is a data problem. In the LPS data, only the LTV of the first mortgage is observed, but not the combined LTV of the first and a possible second mortgage.⁶ Since the combined mortgage amount determines a borrower's home equity, the fact that second mortgages are unobserved is a problem. In order to mitigate this data problem, I thus focus on first mortgages with a very high LTV because these borrowers should be least likely to have a second mortgage on their home. Though these considerations make a focus on highly leveraged borrowers desirable, there is evidence that the general patterns and conclusions of the analysis apply more widely. Specifically, online appendix A documents that the full sample of all loans generally exhibits the same dynamics across cohorts as this sample of highly leveraged borrowers, though there are, of course, various level differences. Furthermore, in appendix section B.4, I show that under plausible assumptions regarding second mortgages, the main conclusions of the reduced-form exercise generalize to loans

⁶Elul, Souleles, Chomsisengphet, Glennon, and Hunt (2010) provide evidence that second mortgages are frequent and significantly affect the combined loan-to-value ratio. They report that, on average, 26% of all borrowers have a second mortgage, which adds an average of 15% to the combined LTV. Unfortunately, they do not report a breakdown of these statistics by the LTV of the first mortgage.

with an initial LTV of the first mortgage between 75% and 84%.

For each loan cohort (defined by origination month) over time, the data set contains the number of active loans that are delinquent and the number of completed foreclosures. Following much of the empirical literature, I define a loan to be in default when it is 60 days or more past due, i.e. two payments have been missed. Cumulative default rates for a loan cohort are then constructed as the share of active loans that are 60 days or more delinquent times the share of initial loans that are still active plus the share of initial loans where foreclosure has already been completed.⁷ However, I also show in online appendices B.3 and C.5 that all my substantive results are robust to using an alternative default definition of 120 or more days past due, which represents even more serious delinquency.⁸

Information on house prices comes from the Federal Housing Finance Agency (FHFA). The monthly national and census division-level repeat-purchase house price indices between 1991 and 2010, deflated by the Consumer Price Index (CPI), are used as measures of aggregate real house price movements. The simulations of this paper also contain realistic microeconomic house price distributions based on empirical estimates of their variance by the FHFA, as discussed in section 3.4.

The key empirical facts on mortgage default rates and house prices across loan cohorts are presented in figure 1. Figure 1(a) shows the observed cumulative default rates for loan cohorts originated between 2002 and 2008 grouped by the year of origination.⁹

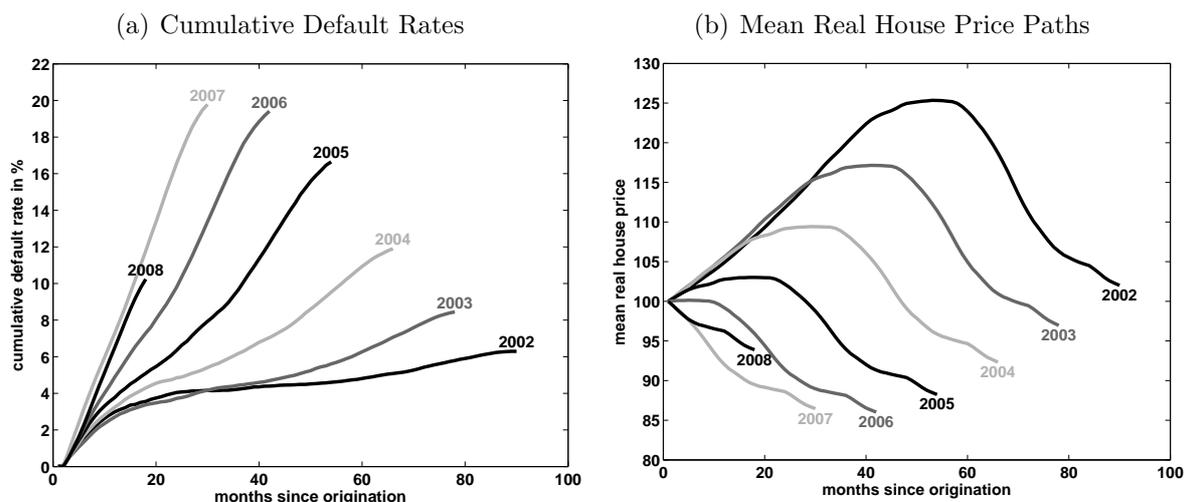
⁷The period of default is backdated by one month to capture the time when the first payment was missed.

⁸The reason for also looking at a 120-days definition is that a considerable fraction of loans that are only 60 days past due will ultimately become current again. Evidence on cure rates is, for example, provided by Adelino, Gerardi, and Willen (2013); and evidence on general transitions between different stages of delinquency by Herkenhoff and Ohanian (2013). However the theoretical literature on mortgage default, which I follow here, models default as a permanent mortgage termination. Thus, there exists a certain tension between the theoretical and empirical literature with respect to the concept of default. The rationale for also looking at a 120-days definition is that more serious stages of delinquency are also much more permanent, as documented empirically in Herkenhoff and Ohanian (2013). Thus, this check addresses the potential concerns regarding the correspondence between theoretical and empirical concepts of default. In this check, I find that my results are robust to using this alternative measure of default. Furthermore, I have also investigated the effect of using a definition of default that requires a loan to be in foreclosure. This approach also generates similar results (which are available upon request) and does not resolve the empirical problems of a frictionless option model, as documented in section 3.

⁹In the data set and all the model simulations of the paper, loan cohorts are defined by month of origination. However, in all the graphs of the paper, I group loan cohorts by year of origination, and the curves shown for an origination year are averages of the underlying twelve cohorts defined by origination

Figure 1(b) presents the mean real house price paths for these cohorts of loans. These mean house price paths accurately account for the geographical composition across census divisions of the different loan cohorts. One observes that mortgage borrowers who experienced a more adverse path of average house price growth rates defaulted much more frequently. Explaining this variation quantitatively - and using it to discriminate between the mentioned theories - are the main aims of the paper.

Figure 1: Cumulative Default Rates and House Prices by Loan Cohort



A general problem of the empirical literature on mortgage default is the lack of loan-level data that link borrowers' repayment behavior to their individual house prices and life events. The implications of unobserved individual unemployment spells for empirical investigations of mortgage default are also discussed by Gyourko and Tracy (2014). Though some studies cited in the introduction have made partial progress in this direction, such an ideal data set does still not exist. Given these general data problems, this paper takes a different approach that is characterized by two main features. First, it investigates at a more aggregate level - of whole cohorts of loans - how shifts in home equity distributions induced by changes to average house prices lead to changes in default rates. Second, it exploits the economic structure of different theories and how their predictions on default specifically depend on home equity. This approach provides novel evidence regarding the house price sensitivity embedded in the competing theories and the ability

month.

of those theories to explain the aggregate dynamics of default during the mortgage crisis.

2.2 Loan Characteristics at Origination

Before proceeding to the analysis, I briefly discuss one alternative explanation for the rise in default rates observed in figure 1(a). This explanation is that lending standards and loan quality deteriorated sharply before the mortgage crisis. Thus, I first present evidence that average loan quality is fairly stable across cohorts in my data set. One should also keep in mind that I only consider prime fixed-rate mortgages. Therefore, any compositional shifts in the mortgage market towards subprime or variable rate mortgages do not - by construction - confound my analysis.

One concern is that the loan-to-value ratio (LTV) might have increased over time, leaving a smaller buffer before borrowers experienced negative equity. I only consider loans that have an LTV above 95% and thus limit this possibility to shifts within that class of loans. Within this class, the average LTV is basically constant across cohorts and only fluctuates mildly around the average value of 98.2%, as seen in the first row of table 1. In section 3, I even control for observed changes across cohorts in the within-cohort distribution of LTVs and find that these are irrelevant.

The second row of table 1 reports the average FICO credit score of the different loan cohorts at the time of origination. These are very stable as well. To the extent that these credit scores are good measures of creditworthiness, a significant deterioration in loan quality is not observable here.

Table 1 also contains information on the average mortgage rate that different cohorts face. A higher mortgage rate might make the loan less attractive to the borrower. There is some variation in this variable across cohorts. However, mortgage rates and default rates seem to be fairly uncorrelated across cohorts.

The average debt-to-income (DTI) ratio, representing the required mortgage payment as a share of gross income, is presented in the last row of table 1.¹⁰ This ratio has increased

¹⁰The data on the DTI is the only mortgage variable in the whole paper that is based on a somewhat different loan selection. The reason is that the DTI was not available in the tool that was used to aggregate and extract information from the LPS loan-level data set. Instead, LPS provided me with a separate tabulation where it was not possible to use the same selection criteria. Specifically, the DTI information is for the same LTV class as the rest of the data but it does not cover only prime,

over time, indicating that borrowers in later cohorts need to devote more of their gross income to service the mortgage. However, the increase was quite modest.

Table 1: Average Loan Characteristics at Origination by Loan Cohort

Cohort	2002	2003	2004	2005	2006	2007	2008	All
LTV in %	98.2	98.3	98.2	98.3	98.4	98.1	97.8	98.2
FICO score	676	673	669	670	668	670	678	672
Mortg. rate in %	6.9	6.0	6.1	6.0	6.6	6.7	6.2	6.4
DTI in %	39	39	40	40	40	42	42	40

These statistics show that in my data set of prime fixed-rate mortgages with an LTV above 95%, there is no evidence in favor of a strong deterioration of lending standards over time. Furthermore, in online appendix A, I document that this stability of loan characteristics across cohorts applies more widely to the full loan sample of all possible initial LTVs, though there are of course level differences. Nevertheless, these conclusions may be somewhat specific to the prime market.¹¹ However, when thinking about the entire mortgage market, one should keep in mind that the prime market is much larger than the subprime market, which even in its heyday constituted only about 20% of new loan originations (Joint Center for Housing Studies at Harvard University 2008, p. 4). Overall, this evidence casts doubts on explanations of the mortgage crisis that rely solely on lax lending standards. Instead, this paper shows that the decline in house prices can explain the increase in default rates within a suitable theoretical model.

fixed-rate, 30-year mortgages. However the vast majority of loans in the LPS data are prime, fixed-rate mortgages, and the modal maturity of these loans is 30 years, so this information should at least be a good approximation of the actual loan pool I consider.

¹¹For example, Demyanyk and Van Hemert (2011) present evidence that loan quality deteriorated in the subprime market. However, Palmer (2013) finds that the fall of house prices is also the dominant cause of the rise of mortgage defaults in the subprime market. Similar to the statistics presented above, Amromin and Paulson (2009) note that it is difficult to detect a deterioration in loan quality in the prime market. A particular advantage of my descriptive statistics is that they are based on all loans in the LPS data base that satisfy my sample selection criteria. Other empirical studies using LPS data typically work with a 1% random sample, such that their descriptive statistics are based on fewer observations.

3 Reduced Form Models

This section presents evidence on mortgage default based on estimating and simulating two highly stylized models. These models represent the simplest possible reduced forms of a frictionless option-theoretic model (the “threshold” model) and the double-trigger hypothesis (the “shock” model). A key feature of the simulation is that it includes empirically accurate home equity distributions within and between cohorts, emanating from variation in initial mortgage balances and house prices. The aim is to discriminate between pure forms of these theories in a relatively general way that is independent of the exact specification of the respective structural model. Several recent structural default models cited in the introduction are, in a sense, hybrids of these two theories. Thus, one may also view this section as a general empirical inspection of the two key economic mechanisms embedded in recent default models. The insights into the empirical performance of these default mechanisms then guide the development of a suitable structural economic model in the following section.

3.1 Model Setup

This paper considers individual borrowers who took out fixed-rate 30-year mortgages. Each loan cohort defined by origination date consists of many borrowers who are indexed by $i = 1, \dots, N$ and observed in periods $t = 1, \dots, T$ after loan origination. Borrowers make a single decision each period and can either service the mortgage or default on the loan and “walk away” from the house. Denote the default decision of an individual borrower i in month t after origination by a set of dummy variables d_{it} . The variables d_{it} take the value 1 once the borrower has defaulted, and they take the value 0 in all periods prior to default. Thus, it is sufficient to present default decision rules in period t for situations when the borrower has not yet defaulted.

For a fixed-rate mortgage, the nominal mortgage balance M_{it} of borrower i evolves deterministically over time according to

$$M_{i,t+1} = (1 + r^m)M_{it} - m_i \tag{1}$$

where r^m is the monthly mortgage rate, which is constant across individuals, and m_i are fixed nominal monthly payments. As in the real world, these fixed payments mostly cover mortgage interest early on in the contract; later, a larger share goes towards the repayment of principal. The payments m_i are determined at the beginning of the contract and satisfy

$$m_i = \left[\sum_{t=1}^T \frac{1}{(1+r^m)^t} \right]^{-1} M_{i0} \quad (2)$$

where M_{i0} is the initial loan amount and the loan has a maturity of $T = 360$ months. The initial loan amount is a function of the initial loan-to-value ratio LTV_i and initial house price P_{i0} and given by $M_{i0} = LTV_i \times P_{i0}$. Here, borrowers are heterogeneous with respect to the LTV. It is assumed that agents make decisions based on real variables. Thus, it is useful to define the real mortgage balance as $M_{it}^{real} = \frac{M_{it}}{\Pi_t}$ where Π_t is the CPI and $\Pi_0 = 1$. This assumption does not affect the results, and the conclusions are identical when decisions are based on nominal variables.

The real house price P_{it} of a homeowner evolves stochastically over time, as described in section 3.4 below. P_{i0} is normalized to 100.

The real home equity of a borrower is given by $P_{it} - M_{it}^{real}$, i.e. the difference between the value of the house and the outstanding mortgage amount. The two default theories differ in how their predictions depend on home equity.

3.2 The Threshold Model

The first model assumes that borrowers with negative equity default on their mortgage the first time that the real value of equity falls below a certain threshold value. Therefore, I call this the “threshold model”. Here, I adopt the simplest possible specification with a threshold that is proportional to the initial house price and constant over time, given by ϕP_{i0} where $\phi < 0$.¹² If, in period t , the borrower has not yet defaulted, then the default

¹²One could also argue for a specification with a threshold ϕP_{it} such that the equity threshold is proportional to the current house price. This is just a minor difference in functional form, and all the main results below are basically identical when using such a specification instead. The only difference is that the estimated value of ϕ is then slightly larger in absolute magnitude.

decision in that period is described by

$$d_{it} = \begin{cases} 1, & \text{if } P_{it} - M_{it}^{real} < \phi P_{i0} \\ 0, & \text{otherwise} \end{cases} \quad (3)$$

This is a simple reduced form of a frictionless option model. The corresponding structural model would derive the threshold parameter ϕ from optimizing behavior. For example, the borrower might trade off the expected future capital gains on the house for the mortgage payments in excess of rents. Thus, the parameter ϕ should depend heavily on expected future house prices. The fact that such expectations are hard to pin down is one important reason why it has been difficult to compare predictions of such a theory to empirical default behavior. Here, I remain agnostic about the exact trade-off and the value of ϕ and instead estimate it from the data.

3.3 The Shock Model

The second model assumes that borrowers with any level of negative equity only default on their mortgage when they also receive a default shock in that period. I call this the “shock model”. Again, I adopt the simplest possible specification. The probability of experiencing a default shock $\psi \in [0, 1]$ is constant, and default shocks are independently and identically distributed over time. If the borrower has not yet defaulted, the default decision in period t is determined by

$$d_{it} = \begin{cases} 1, & \text{if } P_{it} - M_{it}^{real} < 0 \text{ and the default shock occurs} \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

This is a reduced-form of a double-trigger model. Here, the default shock represents a life event that triggers default when combined with negative equity. The parameter ψ represents the probability that a life event occurs, and is estimated from the data. Possible examples of such a life event could be unemployment or divorce, but I again preserve generality here and remain agnostic about the exact nature of these events. Section 4 then provides a micro-founded double-trigger model where unemployment acts as the life event.

3.4 Simulation of House Prices

This section describes how house prices are modeled and simulated. This information also applies to the structural model in the following section. The general aim is to accurately capture realistic house price distributions within and between cohorts in the simulation based on the empirical procedures and estimates of the FHFA.

Throughout the paper, the evolution of the real house price P_{it} of an individual house i in period t is modeled as

$$\ln(P_{it}) = \ln(P_{i,t-1}) + g_t^{agg} + g_{it}^{ind} \quad (5)$$

where the house price growth rate has two components, an aggregate component g_t^{agg} that is common to all houses and an individual component g_{it}^{ind} specific to the individual house. Including an individual component is important in order to accurately capture the house price variation among borrowers within a cohort. Without this feature theoretical models cannot explain any default during times of positive aggregate house price growth. This formulation is consistent with the approach used by the FHFA to estimate the house price index, cf. the description in Calhoun (1996).¹³

In equation (5), a census division index was suppressed for convenience. However, the aggregate trend represented by g_t^{agg} and the moments of g_{it}^{ind} are in fact specific to the census division in which the house is located. Thus, this paper uses information on house prices at the census division level and the regional composition of loan cohorts to simulate house prices accurately. When drawing house prices, the simulation draws are allocated across census divisions such that in each cohort the simulated sample has the same regional composition as in the mortgage data. The aggregate component g_t^{agg} represents the growth rate of the census division real house price index. In the simulation, this component is taken directly from the FHFA data deflated by the CPI. The aggregate component generates the variation in mean house price paths across loan cohorts.

The individual component g_{it}^{ind} is unobserved. However, the FHFA provides estimates

¹³I use a slightly different notation relative to the FHFA because I want to use this equation in a dynamic optimization problem and simulations. In order to see how it is related, rewrite equation (5) as $\ln(P_{it}) = \ln(P_{i,0}) + \sum_{\tau=1}^t g_{i\tau}^{agg} + \sum_{\tau=1}^t g_{i\tau}^{ind}$ where $\ln(P_{i,0}) + \sum_{\tau=1}^t g_{i\tau}^{agg} = \beta_i + N_i$ and $\sum_{\tau=1}^t g_{i\tau}^{ind} = H_{it}$ give equation (1) in Calhoun (1996).

of the variance, and I use these to simulate a realistic microeconomic house price distribution. Specifically, it is assumed that the individual component g_{it}^{ind} is independent over time and individuals and that it is normally distributed with mean zero and variance V_t . The variance of g_{it}^{ind} depends on the time since the house was bought. This is a realistic feature of the data and is based on estimates of the FHFA. For simplicity, the following exposition assumes that the house was bought in period 0 such that t is also the time since purchase. Using my own notation, the FHFA specifies a quadratic formula in time for the variance of the total individual part of the house price change since purchase, given by

$$\text{Var} \left(\sum_{\tau=1}^t g_{i\tau}^{ind} \right) = \frac{\kappa}{3}t + \frac{\lambda}{9}t^2 \quad (6)$$

where an adjustment has been made for the fact that this paper operates at a monthly instead of a quarterly frequency. By the independence assumption, the variance of g_{it}^{ind} is thus given by

$$V_t = \text{Var} (g_{it}^{ind}) = \text{Var} \left(\sum_{\tau=1}^t g_{i\tau}^{ind} \right) - \text{Var} \left(\sum_{\tau=1}^{t-1} g_{i\tau}^{ind} \right) = \frac{\kappa}{3} + \frac{\lambda}{9}(2t - 1). \quad (7)$$

The FHFA provides estimates of κ and λ at the census division level that I use to generate realistic distributions around the division level aggregate trends. The estimates of κ are positive, and those of λ are negative and small in absolute magnitude. This implies that the variance of $\sum_{\tau=1}^t g_{i\tau}^{ind}$ increases less than linearly with time and that the variance of a single g_{it}^{ind} decreases over time.¹⁴

3.5 Model Simulation, Estimation and Test

Conditional on the respective model parameters ϕ and ψ , both models can be simulated for subsequent cohorts of loans originated each month between 2002 and 2008. For each cohort, I draw 25,000 individual histories of house prices, and for the shock model, I also draw default shock histories. When computing the mortgage balance, the mortgage rate is kept constant within a cohort and set equal to the respective cohort average. However,

¹⁴On average, across census divisions, the estimates of κ and λ imply that the shock in the first month g_{i1}^{ind} has a standard deviation of about 2.49%, while after five years the standard deviation of g_{i60}^{ind} is around 2.37%. Hence, the standard deviation of g_{it}^{ind} decreases relatively slowly over time.

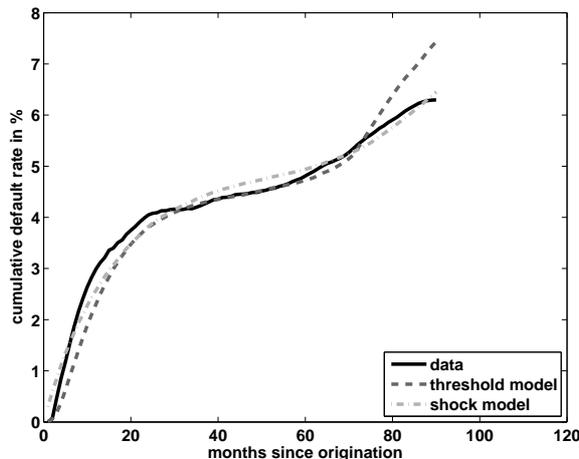
borrowers within a cohort are heterogeneous with respect to the initial LTV, which varies in steps of one percentage point between 95% and 104% (the few loans with a higher LTV are subsumed in the 104% LTV bin). The frequency of these different loan-to-value ratios at origination varies across cohorts, as observed in the mortgage data. This means that possible changes to the average mortgage rate and the LTV distribution across cohorts are taken into account in the simulation. Data on the path of inflation rates from the CPI are used to compute the real mortgage balance. The decision rules are then applied to these shock histories and paths of the real mortgage balance. Therefore, the simulation generates empirically accurate home equity distributions within and between cohorts by including this variation in initial mortgage balances, mortgage rates and house prices.

The idea of the estimation and test procedure is to estimate the unknown model parameters using only the default data of the cohort originated in 2002. The estimation employs a simulated method of moments procedure (McFadden (1989), Pakes and Pollard (1989), Duffie and Singleton (1993)). The respective parameters ϕ and ψ are chosen such that the cumulative default rates for the 2002 cohort simulated from the model match - as well as possible - those observed in the data, cf. online appendix B.1 for details. Keeping the parameter values estimated for the 2002 cohort fixed, the test is based on out-of-sample predictions of the two models for the cohorts originated between 2003 and 2008. The test consists of informally comparing simulated and empirically observed default rates for these remaining cohorts and checking which estimated model provides a better fit to the data.

3.6 Results

For the threshold model, the negative equity default threshold ϕ is estimated as -11.1% . This means that borrowers default as soon as they have a real value of negative equity of 11.1% of the initial house price. In contrast, for the shock model, the default shock probability ψ is estimated to be 1.05% such that each period, 1.05% of those borrowers with negative equity default on their loan. The fit of the two models to the cumulative default rate of the 2002 cohort is shown in figure 2. Both models fit this data very well.

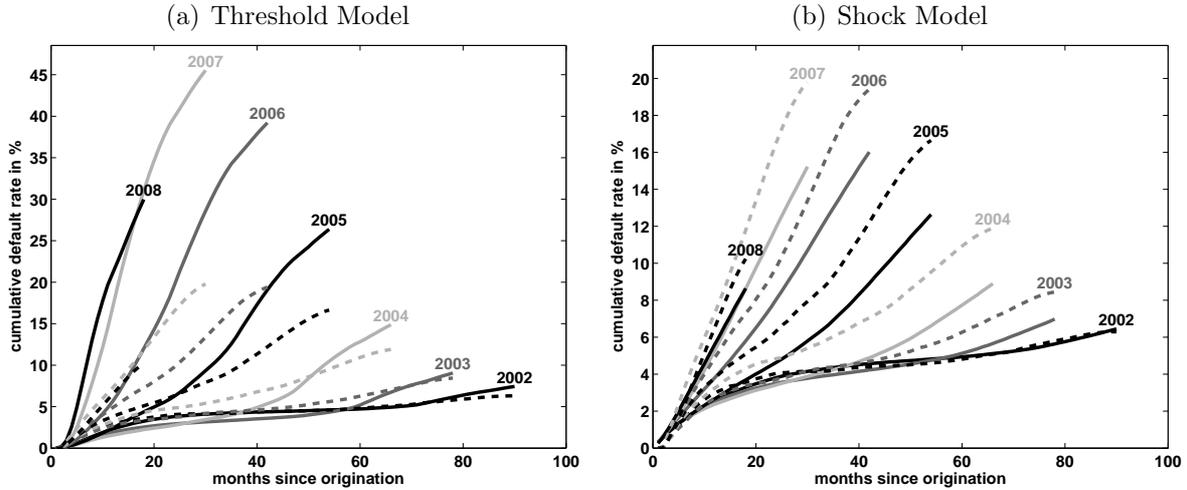
Figure 2: Cumulative Default Rate of 2002 Cohort: Models vs. Data



The next step is to test the two estimated models by checking how well they perform in predicting out-of-sample. Figure 3(a) shows the fit of the threshold model to the full sample of all cohorts between 2002 and 2008. The equivalent fit of the shock model is presented in figure 3(b). It turns out that the threshold model has severe empirical problems. When it is forced to match default rates of the 2002 cohort, it overpredicts default rates for the later cohorts in the simulation period by at least one order of magnitude. The threshold model is excessively sensitive to the shifts in the mean of the house price distribution observed in the data. By contrast, the shock model provides a good fit to the broad dynamics in the data. The impression derived from visually inspecting the graphs is also confirmed by simple measures of goodness-of-fit of the model predictions to the data. For instance, the root mean squared error (mean absolute error) of the out-of-sample forecast is 4.1 (3.2) times higher for the threshold model than for the shock model.

Admittedly, the shock model generates an increase in default rates that is slightly too low. This could indicate that there is a small difference between the quality of borrowers in the 2002 cohort and that of borrowers in the other cohorts, though they appear to be similar based on observed characteristics. Such a conclusion is also supported by the very good fit to all but the 2002 cohort that one obtains when one estimates the model on all cohorts simultaneously instead of just the first cohort, cf. the in-sample fit presented in

Figure 3: Cumulative Default Rates of all Cohorts: Models (solid lines) vs. Data (dashed lines)

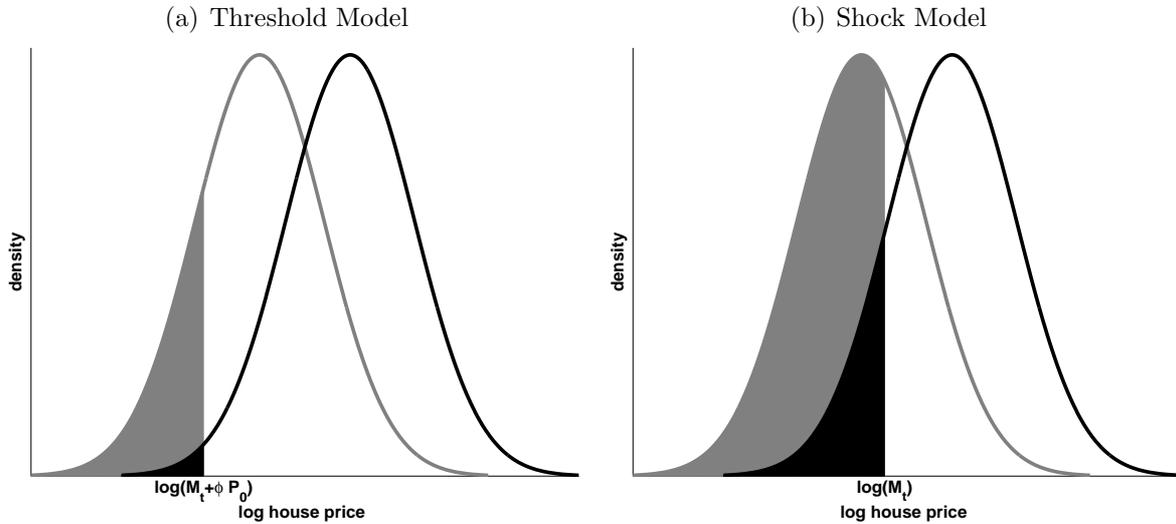


online appendix B.2. Another potential explanation is that life events, and in particular unemployment, became more frequent during this time period.

The explanation for the difference between models is illustrated in figure 4. The figure shows the leftward shift of the (log) house price distribution from the black to the gray distribution due to a fall in aggregate house prices. The shock model predicts that out of all borrowers with negative equity, i.e. those with a (log) house price below the (log) mortgage balance, a fraction ψ defaults each period. In reaction to a leftward shift in the house price distribution, the shock model predicts that the default rate should increase in proportion to the increase in the number of borrowers who experience any level of negative equity. In figure 4(b), this is represented by the ratio of the sum of the black and gray areas to the original black area. It turns out that observed default rates approximately exhibit this pattern. However, the threshold model is concerned with the far left tail of the house price and equity distribution. It predicts that all borrowers with a house price below the sum of the mortgage balance and ϕ times the initial house price default (remember that ϕ is negative). When the house price distribution shifts left, the threshold model predicts that default rates should increase as much as the number of borrowers with extreme levels of negative equity, represented by the shift from the black area to the sum of the black and gray areas in figure 4(a). However, the number of borrowers with such an extreme level of negative equity increases much faster than

observed default rates. This difference between models drives the empirical results.

Figure 4: Illustration of Reaction of Models to Fall in Aggregate House Prices



I have conducted a large number of robustness checks to scrutinize these results, which are reported in detail in online appendices B.2, B.3 and B.4. These checks include estimating the models on cohorts other than the 2002 one, replacing the out-of-sample test with an in-sample test, abstracting from within-cohort and cross-cohort heterogeneity in initial LTVs and mortgage rates, assuming a different distribution of individual house price shocks, allowing threshold and shock parameters to vary over the course of the loan, using an alternative definition of default, and extending the analysis to loans with a lower initial LTV. I find that the results are robust across all these specifications.¹⁵

Two conclusions can be drawn from the results of this section. First, an empirically successful structural model cannot rely exclusively on a single-trigger mechanism. Instead, some feature other than house price shocks must play a role. This motivates the development of a structural double-trigger model in the following section.

A potential criticism of such a conclusion is that a change in the threshold parameter ϕ across cohorts (which I kept constant) may be able to resolve the documented problems

¹⁵I have focused on an out-of-sample exercise here because it documents the differing sensitivities of the theories to changes in aggregate house prices in a simple and intuitive way. However, there are also reasons to consider the performance of the theories in an in-sample test, which examines the fit of the two models when they are estimated on the data from all cohorts. Such an exercise is informative regarding the best possible fit to the data that both models can achieve. An in-sample investigation confirms or even strengthens the out-of-sample results, cf. figure B.1 in appendix B.2. The threshold model still has considerable problems with matching the data, even under these most favorable circumstances. In contrast, the shock model achieves a very good fit to the data.

of single-trigger models. As the threshold model fails by overpredicting the default rates of later cohorts, this would require rational single-trigger borrowers to become more reluctant to default during the crisis. Accordingly, borrowers would need to have more optimistic expectations of future house prices during the housing bust than the housing boom. However, such a change in expectations seems implausible given the economy entered a deep recession at that time. Furthermore it is also inconsistent with survey evidence on house price expectations (Case, Shiller, and Thompson 2012), which indicates that expectations became more pessimistic during the time period considered here. Thus, a threshold model with a time-varying expectations process would have even greater difficulties explaining the default behavior observed during the crisis.

In a similar vein, one may worry that ϕ could vary across cohorts due to changes in interest rates. Indeed, the traditional option theoretic literature predicts that higher interest rates make the borrower more reluctant to default, such that default only occurs at an even more extreme level of negative equity, cf. for example Kau, Keenan, Muller, and Epperson (1992). To shed light on this possibility, I construct interest rate paths for different cohorts from data on the annual nominal interest rates of ten-year government bonds. These data show that at a given time since origination, the 2007 and 2008 cohorts faced lower nominal interest rates than did the 2002 cohort. However, according to the theory, this interest rate variation should make the later cohorts even more willing to default, such that their default threshold ϕ should be closer to zero compared to that of the 2002 cohort. Because the benchmark threshold model already overpredicts default rates for the later cohorts, this evidence suggests that changes to interest rates cannot resolve the empirical problems of frictionless option theory and instead would worsen them.

Another concern may be that the reduced-form nature of the threshold model somehow biased the analysis against frictionless option theory, even though the threshold model does preserve its key theoretical prediction. For this reason, I have also investigated a structural single-trigger model that explicitly solves the stochastic optimization problem of a rational borrower in a frictionless world. This analysis shows that when

comparing a structural single-trigger model to the structural double-trigger model of the next section, the conclusions across models are the same as when comparing reduced form models, cf. online appendix C.6.

A second conclusion is that for a double-trigger model, the increase in the fraction of borrowers with negative equity caused by the mean shift in house prices is sufficient to broadly explain the rise in default rates in this data set. Together with the evidence on the stability of loan characteristics in section 2, this supports the view that the fall of aggregate house prices is key to understanding the observed rise in default rates.

4 Structural Model

This section introduces a micro-founded partial equilibrium model of double-trigger behavior. In this model, unemployment acts as the life event that may trigger default in combination with negative equity. The aim of the structural model is to analyze whether and how unemployment may play this role, to assess how well such a micro-founded model explains the rise in default rates during the crisis and to subsequently use the model for policy analysis in section 7. In the model, a homeowner who bought a house with a fixed-rate mortgage each period chooses non-housing consumption and decides whether to stay in the house and service the mortgage or leave the house and terminate the mortgage. The mortgage can be terminated either by selling the house or by defaulting on the loan by “walking away”. The homeowner faces uncertainty regarding the future price of the house, unemployment shocks and a borrowing constraint for unsecured credit. It is also assumed that the homeowner derives a direct utility flow from the bought house. One period corresponds to one month. Throughout this section, an individual index i is suppressed for convenience.

4.1 Mortgage Contract

The household took out a fixed rate mortgage with outstanding nominal balance M_0 and nominal mortgage rate r^m to finance the purchase of a house of price P_0 in period 0. Mortgage interest and principal have to be repaid over T periods in equal installments of

nominal value m that are fixed at the beginning of the contract and satisfy equation (2). Over time, the outstanding nominal mortgage balance M_t evolves according to equation (1) as long as the household services the mortgage.

4.2 Preferences and Choices

Preferences are specified as in Campbell and Cocco (2003) but allow a direct utility benefit of owning a house. Household decisions over the length of the mortgage contract are determined by maximizing the expected utility given by

$$U = E_0 \sum_{t=1}^T \beta^{t-1} \left(\frac{C_t^{1-\gamma}}{1-\gamma} + \theta \mathcal{I}(\text{own}_t) \right) + \beta^T \frac{W_{T+1}^{1-\gamma}}{1-\gamma} \quad (8)$$

which is derived from consumption C_t in periods 1 to T and remaining wealth W_{T+1} at the end of the contract.¹⁶ The flow utility function from consumption is assumed to be of the CRRA form, where γ denotes the parameter of relative risk aversion and the inverse of the intertemporal elasticity of substitution. β is the time discount factor. $\mathcal{I}(\text{own}_t)$ is an indicator variable that is one if the agent owns a home in period t and zero otherwise. θ is a direct utility benefit from being a homeowner. This could reflect, for example, an emotional attachment to the specific house or the benefit that an owner cannot be asked to move out by a landlord, as may happen to a renter. This direct utility flow is a very important feature to generate double-trigger behavior because it prevents employed agents from defaulting after a strong fall of house prices, as explained in more detail in section 6.3 below.

In each period, the homeowner has to decide how much to consume and whether to stay in or leave the house. If the agent wants to leave, this can be done by either selling the house (and repaying the current mortgage balance) or by defaulting on the loan by “walking away”.¹⁷

¹⁶Following Campbell and Cocco (2003), the specification in equation (8) implicitly assumes that the borrower maximizes utility only over the course of the mortgage contract because the continuation value is largely arbitrary. An obvious alternative is to extend the utility function to the remaining lifetime of the borrowers. One complication here is that I do not have any demographic information on the borrowers in my data set. However, I have experimented with adding further time periods after the end of the mortgage contract and also with including a retirement period. These additions had no significant effect on the results.

¹⁷The model abstracts from mortgage termination through refinancing for computational reasons,

4.3 Constraints

The dynamic budget constraint depends on the borrower's house tenure choice. For a homeowner who stays in the house, it is given by

$$A_{t+1} = (1 + r) \left(A_t + Y_t - \frac{m}{\Pi_t} + \tau r^m \frac{M_t}{\Pi_t} - C_t \right) \quad (9)$$

where A_t denotes real asset holdings and Y_t denotes real net labor income in period t . The real interest rate on savings r is assumed to be constant over time. m is the nominal payment to service the mortgage. However, the nominal mortgage interest $r^m M_t$ is tax deductible and τ is the tax rate. All nominal variables need to be deflated by the current price level for consumption goods Π_t to arrive at a budget constraint in terms of real variables. The presence of Π_t generates the “mortgage tilt effect”. This means that due to inflation, the real burden of the mortgage is highest at the beginning of the contract and then declines over time. It is assumed that the inflation rate π is constant over time, and Π_t thus evolves according to $\Pi_{t+1} = (1 + \pi)\Pi_t$ with $\Pi_0 = 1$.

In the event that the house is sold at the current real price P_t , the homeowner needs to repay the current outstanding nominal mortgage balance M_t and can pocket the rest. The budget constraint of a seller reads as

$$A_{t+1} = (1 + r) \left(A_t + Y_t - R + P_t - \frac{M_t}{\Pi_t} - C_t \right). \quad (10)$$

Here, R is the real rent for a property of the same size. It is assumed that an agent who terminates the mortgage through prepayment or default needs to rent an equivalent house for the rest of his life, which is a common assumption in the literature.¹⁸ The resulting parsimonious specification simplifies the computational solution of the model considerably. However, the assumption also captures the economically important fact that in the real world, a defaulting borrower is closed out of the mortgage market for an

which is a common simplification of the related theoretical literature cited in the introduction. Otherwise, the mortgage balance becomes a separate state variable. This is unlikely to be a major limitation in the context of default because refinancing is only feasible when the borrower has positive equity in the house or substantial other liquid assets. Thus, refinancing does not directly compete with the default decision in a situation of negative equity and low liquid wealth.

¹⁸Thus, a change in housing status from owning to renting is irreversible. This assumption also rules out downsizing the house after a default, which could play a role in the default decision of borrowers in the real world.

extended period of the time and experiences a strong decline in his credit rating. From the borrower’s point of view, this is one of the costs of defaulting. In the absence of such costs, a rational borrower would find it optimal to default already at very small levels of negative equity independently of his liquidity position, which would lead to clearly counterfactual predictions.

Real rents are assumed to be proportional to the initial house price and then constant over time as

$$R = \alpha P_0. \tag{11}$$

This specification involves both a highly realistic feature of rents and an approximation. The realistic feature is that during the period of study, real rents remained almost constant, while real house prices first increased and then decreased enormously. The specification implies that after origination, the rent-price ratio decreases when real house prices increase. Such a negative relationship between the rent-price ratio and real house prices exists in the data provided by Davis, Lehnert, and Martin (2008), not only during the recent period but also at least since 1975. In this paper, I take these observations as given and specify the exogenous variables of the model accordingly. Explaining this pattern is an important area for future research. However, a fully realistic specification would also require to make α cohort-specific. Instead, for computational reasons, I use an approximation such that α is constant across cohorts and is calibrated to a suitable average.

In contrast, if the agent decides to default on the mortgage by “walking away”¹⁹ or is already a renter, the budget constraint is given by

$$A_{t+1} = (1 + r)(A_t + Y_t - R - C_t). \tag{12}$$

It is assumed that for reasons not explicitly modeled here, the household faces a

¹⁹The specification assumes a non-recourse loan, which is a common assumption for the U.S. mortgage market even though formally there are recourse laws in some states. However, the empirical study of Ghent and Kudlyak (2011) finds that recourse only deters those borrowers from defaulting who own relatively high-value properties (above \$200,000 in real 2005 terms). As my data set contains borrowers with much lower average house values, the neglect of recourse does not seem to be a major concern.

borrowing constraint for unsecured credit, given by

$$A_{t+1} \geq 0. \tag{13}$$

Together with the budget constraints above, this implies that the amount of resources available for consumption in a given periods depend on current wealth and the housing tenure choice. By modeling borrowing constraints, the model builds on the buffer-stock saving framework of Zeldes (1989), Deaton (1991) and Carroll (1997).

For a homeowner, remaining wealth at the end of the contract is given by $W_{T+1} = A_{T+1} + Y_{T+1} + P_{T+1}$ and for a renter, by $W_{T+1} = A_{T+1} + Y_{T+1}$.

4.4 Labor Income Process

The household’s real net labor income Y_t is subject to idiosyncratic unemployment shocks and is exogenously given by

$$Y_t = \begin{cases} (1 - \tau)Y_0 & \text{if employed} \\ \rho(1 - \tau)Y_0 & \text{if unemployed} \end{cases} \tag{14}$$

where Y_0 is initial real gross income, τ is the tax rate and ρ is the net replacement rate of unemployment insurance. Over time, employment status evolves according to a Markov transition process with the two states “employed” and “unemployed” and constant job separation and job finding probabilities. Employed agents lose their job with probability s and stay employed with probability $(1 - s)$. Unemployed agents find a job with probability f and stay unemployed with probability $(1 - f)$. Here, I focus on income fluctuations due to unemployment risk because unemployment involves a severe decline in labor income from one month to another, making it a very plausible cause of short-run liquidity problems. This also allows to relate the model closely to the double-trigger hypothesis and the empirical evidence that default is correlated with state unemployment rates.²⁰

²⁰My formulation abstracts from deterministic changes to labor income over the life-cycle and keeps the labor income of employed and unemployed agents constant over time. One reason for this is, again, the lack of demographic information on the borrowers in my data set. In any case, these borrowers belong to the lower half of the income distribution, and people in lower income classes tend to have relatively flat income profiles. Nevertheless, if income during unemployment rises over time, then this prolongs the

4.5 House Price Process

Real house prices are exogenous and evolve over time, as specified in section 3.4 and equation (5). It is assumed that homeowners view the aggregate component g_t^{agg} of house price appreciation to be stochastic and distributed according to an i.i.d. normal distribution with mean μ and variance σ^2 . This process for the aggregate house price component is only used for forming agents' expectations. In the simulation, the realizations of g_t^{agg} are those observed in the data. For the individual component, agents know that g_t^{ind} is distributed normally with mean zero and time-varying variances that depend on the parameters κ and λ as specified in section 3.4. In order to reduce the computational burden when computing policy functions, the parameters μ , σ , κ and λ are not varied across the nine census divisions. Instead, they are set equal to national averages, cf. section 5.2 on the calibration. However, the house price realizations in the simulation of the model are generated from the division-specific data and distributions.

4.6 Initial Conditions

The homeowner solves the dynamic stochastic optimization problem conditional on initial asset holdings A_0 , initial employment status, an initial loan-to-value ratio $LTV = \frac{M_0}{P_0}$ and a debt to (gross) income ratio $DTI = \frac{m}{Y_0}$, which refers to the ratio of the monthly mortgage payment to gross income. I assume that borrowers were employed when they obtained their loans. With respect to initial assets A_0 , I use the computed policy functions to set initial assets equal to the buffer-stock desired by a borrower in period 1 who is employed and faces a house value equal to P_0 . Thus, I shut down possible effects from borrowers first converging to their desired buffer-stock and being more vulnerable to income shocks during the time immediately after origination. The initial house price P_0 is normalized to 100. LTV and DTI then uniquely determine M_0 and Y_0 .

4.7 Computation, Simulation and Test

The model is solved computationally for the optimal policy functions. The borrower's finite horizon optimization problem is characterized by four state variables (liquid wealth

period until buffer-stock savings are exhausted and default occurs.

$X_t = A_t + Y_t$, employment status L_t , house price P_t and time t) and two choice variables (consumption C_t and the mortgage termination choice). Note that for a fixed-rate mortgage, the mortgage balance M_t evolves deterministically over time and is thus captured by the state variable t . The solution proceeds backwards in time and iterates on the value function (the value functions are provided in online appendix C.1). The continuous state and control variables are discretized, and the utility maximization problem in each period is solved by grid search. Expected values of future variables are computed by Gaussian Quadrature. Between grid points, the value function is evaluated using cubic interpolation.

Given the policy functions, the model is simulated for subsequent cohorts of loans originated each month between January 2002 and December 2008. For each cohort, I draw 25,000 individual histories of house prices as explained in section 3.4 and employment histories from the Markov process of section 4.4.

The general idea of the performed computational exercise is the same as in the reduced-form section. I use only the default data from the 2002 cohort (and other data sources) to determine model parameters, as explained in section 5. The test of the model then consists again of informally comparing the out-of-sample model predictions for default rates of the 2003 to 2008 cohorts to the actual observations.

5 Parametrization

The structural model is parameterized in two steps. First, the mortgage contract, house price expectations, rents, labor income, interest and inflation rates are calibrated to data on the respective variables, i.e. to data other than default rates. The preference parameters are divided into one set that is predetermined and another that is estimated, such that the model fits the cumulative default rates of the 2002 loan cohort. All parameter values are summarized in table 2 below.

5.1 Mortgage Contract Characteristics

This paper restricts its attention to 30-year ($T = 360$ months) fixed-rate mortgages. I use average characteristics at origination of the loans in my data set to determine the loan-to-value ratio, mortgage rate and debt-to-income ratio. The average initial loan-to-value ratio of these loans is 98.2%, so I set $LTV = 98.2\%$. The nominal mortgage rate r^m is set to 6.4% per annum, which is the average mortgage rate for newly originated loans in my data set. The debt-to-income ratio DTI is set to 40% as in the data.

5.2 House Price Expectations

As explained above, when computing policy functions, the parameters μ , σ , κ and λ are not varied across the nine census divisions. The monthly house price index from the FHFA at the national level between 1991 and 2010 deflated by the Consumer Price Index (CPI) is used to estimate the parameters μ and σ of the aggregate component. I find that at a monthly frequency $\mu = 0.065\%$ and $\sigma = 0.55\%$. These values imply expected yearly aggregate real house price growth of about 0.8% and a yearly standard deviation of 1.9%. This means that agents in the model have expectations of real aggregate house price growth that, on average, were correct in the years 1991 to 2010 as far as the mean and standard deviation are concerned.

The parameters κ and λ are determined as a simple average of those estimated by the FHFA for each of the nine census divisions. This gives $\kappa = 0.00187$ and $\lambda = -4.51E - 6$ and implies that the individual house price growth shock g_{it}^{ind} in the first month after house purchase is expected to have a standard deviation around 2.5%.

5.3 Income Process

The average tax rate τ is set to 16% and the net replacement rate of unemployment insurance ρ to 62%. This is based on the OECD Tax-Benefit calculator²¹ for the United States. Specifically, the average loan amount, mortgage rate and debt-to-income ratio are used to determine the average gross income of the borrowers in the data set. Based on gross income, the calculator reports the net income in work and out of work, which

²¹<http://www.oecd.org/social/soc/benefitsandwagestax-benefitcalculator.htm>

then determine the average tax and net replacement rates. These calculations take taxes, social security contributions, in-work and unemployment benefits into account. Precise numbers, especially for the tax rate, also depend on the demographics of the household. I have used the average values for a married couple with one earner and no children.

Data from the Bureau of Labor Statistics on the national unemployment rate and median unemployment duration are used to compute time-series of monthly job-finding and -separation probabilities. This is done using steady state relationships. In a steady state, the median duration of unemployment d and the unemployment rate u should satisfy $(1-f)^d = 0.5$ and $u = \frac{s}{s+f}$. These two equations are then solved for the time-series of f_t and s_t , implied by the time-series of u_t and d_t .²² I then set $s = 1.8\%$ and $f = 31\%$, which are the average values of the computed monthly finding and separation probabilities during the period 1990 to 2010. These values imply a steady state unemployment rate around 5.5%. I keep these values of s and f constant in the simulation of the model. However, using the observed changing values when simulating the model yields very similar results (without recomputing policy functions).

5.4 Other Prices

Nominal interest rates for 1-year Treasuries and changes to the Consumer Price Index (CPI) are used to compute real interest rates and inflation rates. Based on these data between 1990 and 2010, the real interest rate r is set equal to 1.4% per year. The inflation rate π is set to 2.4% annually, which is the average value during the simulation period. The initial rent-price ratio parameter α is set equal to 4.0% on a yearly basis, which is the average rent-price ratio between 2002 and 2008 in the data provided by Davis, Lehnert, and Martin (2008).

²²As a check on this procedure, I predict the unemployment rate from the dynamic equation of unemployment $u_{t+1} = u_t + s_t(1 - u_t) - f_t u_t$ using the computed time series of finding and separation probabilities as inputs. It turns out that this provides an excellent fit to the path of the actual unemployment rate.

Table 2: Model Parameters

Mortgage contract	Contract length in months	T	360
	Mortgage rate (yearly)	r^m	6.4%
	Initial loan-to-value ratio	LTV	98.2%
	Initial debt-to-income ratio	DTI	40%
House price process	Mean of aggregate component	μ	0.065%
	Std. dev. of aggregate component	σ	0.55%
	Linear coefficient in indiv. variance	κ	0.00187
	Quadratic coefficient in indiv. variance	λ	-4.51E-6
Income process	Job separation probability	s	1.8%
	Job finding probability	f	31%
	Tax rate	τ	16%
	Net replacement rate of UI	ρ	62%
Other prices	Real interest rate (yearly)	r	1.4%
	Inflation rate (yearly)	π	2.4%
	Rent-price ratio (yearly)	α	4.0%
Preferences	CRRA coefficient	γ	5
	Discount factor (yearly)	β	0.9
	Utility benefit of owning	θ	0.28

5.5 Preferences

In order to reduce the computational burden and due to identification concerns, I first choose reasonable values for β and γ and then estimate only θ . For the intertemporal elasticity of substitution, which is the inverse of γ , Guvenen (2006) reviews empirical estimates ranging from around 1 to 0.1, which implies values of γ ranging from 1 to 10. Furthermore, he argues that conflicting estimates can be reconciled if the rich have a high and the poor have a low elasticity. I choose $\gamma = 5$, which is in the middle of this range. Following Guvenen's reasoning, one could also argue for higher values because borrowers in my data set belong to the lower half of the income distribution. For β , I choose a value of 0.9 at an annual frequency which may be a bit on the low side. However, adapting Guvenen's argument to β , the reason is that I am analyzing borrowers who make a very small down payment, perhaps because they are very impatient. The agents who are net savers could then have a higher discount factor. Online appendix C.2 contains a

sensitivity analysis with respect to β and γ .²³

Given values of β and γ , the parameter θ , representing the direct utility benefit from owning the house, is estimated by the simulated method of moments. Again, the parameter is chosen such that cumulative default rates simulated from the model match those observed for the 2002 cohort, and the remaining data are used to test the ability of the estimated model to predict out-of-sample. This yields an estimate for θ of 0.28. Relative to a situation of $\theta = 0$, a homeowner at the initial conditions values this direct utility flow of $\theta = 0.28$ about as much as a permanent increase in income of 1.5%. The real present discounted value of such a hypothetical permanent increase in income corresponds to about 6% of the initial house price. Thus, the estimated direct utility benefit is of a modest - but still significant - magnitude here.^{24,25}

²³Results from the sensitivity analysis can be summarized as follows. The model relies on a sufficiently high value of γ and low value of β to generate double-trigger behavior. The reason is that a low willingness to substitute intertemporally and a high impatience to consume today worsen the liquidity problem caused by unemployment. As a result, being employed and being unemployed are sufficiently different states, as required for double-trigger behavior. If this is not the case, then a sizeable portion of employed agents default in all cohorts, which brings the model too close to a frictionless option model and the overshooting problems already observed in section 3.

²⁴The estimated direct utility benefit parameter θ may also capture deviations of house price expectations from the empirically observed long-run house price growth rates to which the model is calibrated. For instance, if expected real house price growth μ is increased by one percentage point on an annual basis, then θ is instead estimated at a lower value of about 0.26. This reflects the general importance of house price expectations in this literature. The empirical performance of the model with such modestly overoptimistic house price expectations is very similar to the benchmark calibration. However, a model with very overoptimistic house price expectations, and accordingly a direct utility benefit θ close to zero, does not perform well empirically because in such a model, a strong fall in house prices will also induce employed and non-liquidity-constrained households to default, cf. the detailed explanation in section 6.3.

²⁵Laufer (2018) also includes a default penalty to match data on default rates among underwater borrowers, which in his model takes the form of a higher rent-price ratio for defaulters relative to non-defaulters. However, he estimates an extremely large default penalty that corresponds to 29% of permanent income. Among the many differences in the model structure, calibration and analyzed data between the two papers, two features seem especially likely to explain these different findings. The first is that Laufer includes time-varying house price expectations in his model, which imply that during the crisis, homeowners expect much lower house price growth rates than during the boom, and in fact, they expect strongly falling house prices for an extended period of time. In addition, he assumes the required rent after default to be proportional to the homeowner's pre-default house price such that a fall in house prices also lowers the required rent payments after default. Both of these features are likely to strongly increase default incentives during the crisis, such that Laufer's model requires a very high default penalty to match actual default rates. By contrast, these features are absent in my model, such that it requires much weaker default penalties.

6 Results

This section first explains the mechanism that generate default in the model. I then present the main results of how well the model fits the rise in default rates across loan cohorts. Finally, I explain some insights into the modeling of double-trigger behavior with empirically accurate house price sensitivity.

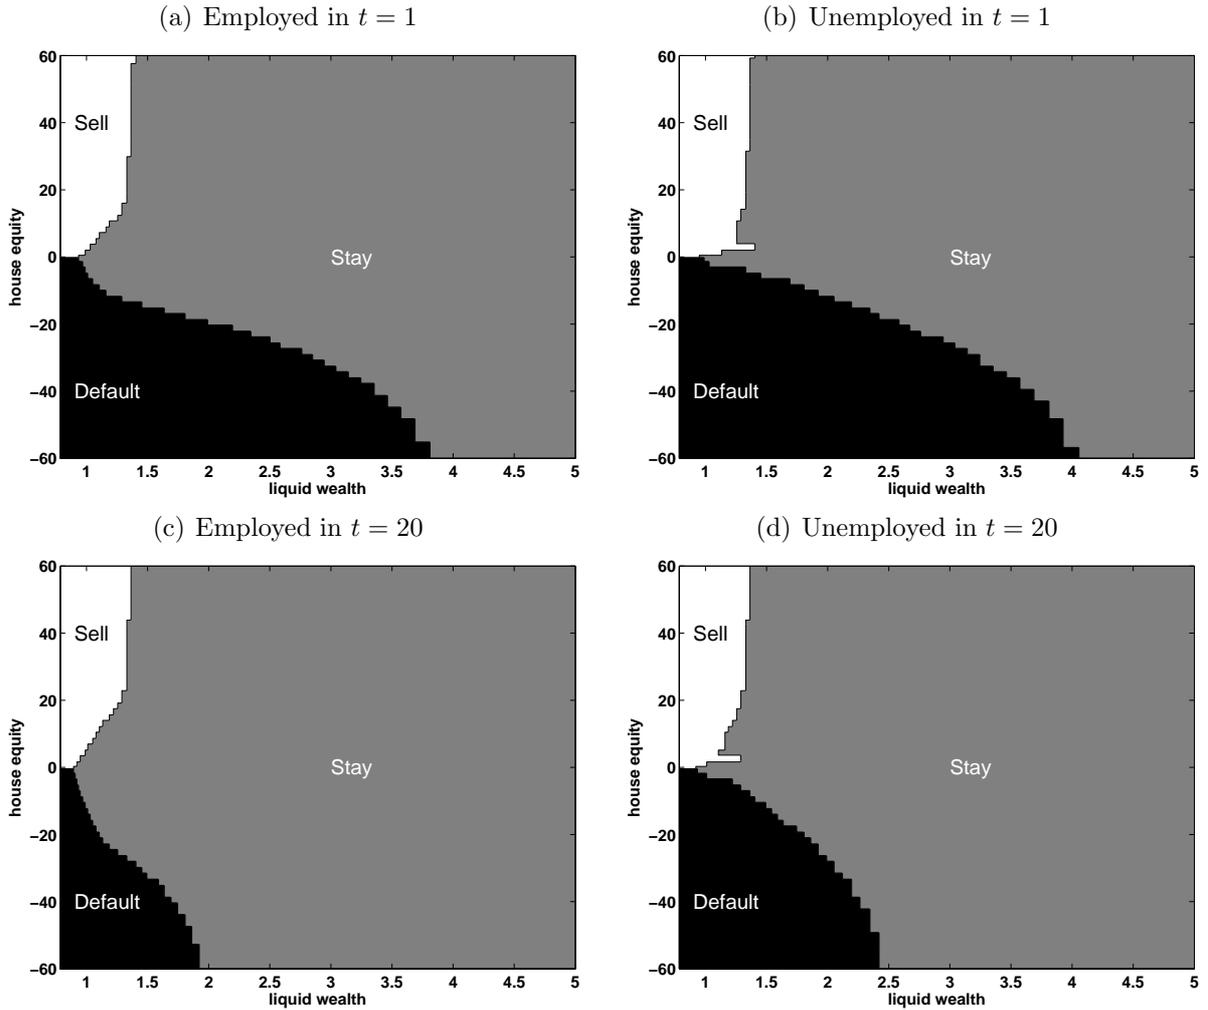
6.1 The Default Mechanism

The repayment policy function of a borrower in the model is presented in figure 5 as a function of house equity, liquid wealth, employment status and time. Several features are noteworthy. First, negative equity is a necessary condition for default. By contrast, with positive equity, selling is strictly preferred to defaulting because the borrower is the residual claimant of the house's value after the mortgage balance has been repaid.

Second, negative equity is not sufficient for default. There are many combinations of state variables where a borrower with negative equity prefers to stay in the house and service the mortgage. In a negative equity situation, the basic trade-off of the borrower is the following (postponing the role of the borrowing constraint until the next paragraph). The cost of staying in the house is that the borrower needs to make the mortgage payment, which is higher than the rent for an equivalent property. The benefit of staying is that the borrower receives the utility benefit of owning a house and keeps the option to default, sell or stay in a later period. Specifically, there are possible future states of the world with positive equity. However, the probability of reaching those states depends on the current house price. This establishes a default threshold level of the house price. Of course, when making this decision, the rational borrower will also need to discount these future gains and take risk aversion into account.

Third, the level of negative equity at which the borrower exercises the default option depends on non-housing state variables: liquid wealth and employment status. Specifically, a borrower who is unemployed and/or has low liquid wealth will default at lower levels of negative equity. There are two reasons for terminating the mortgage in these states. One is that current borrowing constraints may bind, and the borrower terminates

Figure 5: Repayment Policy Function



Notes: Repayment choice as a function of the state variables liquid wealth, house equity, employment status and time. Blue region: Default. Green region: Sell. Red region: Stay.

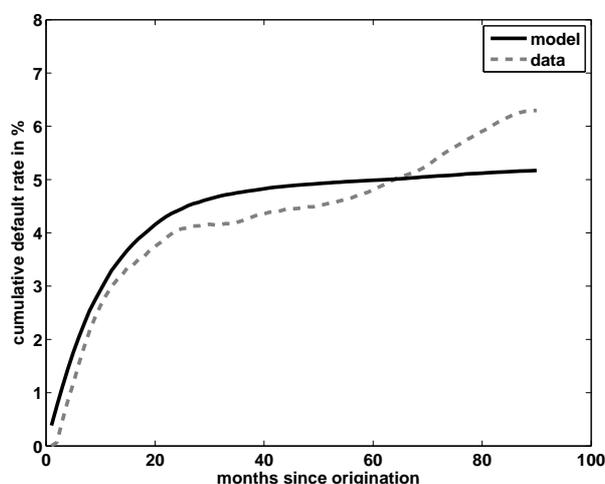
the mortgage to increase current consumption. The other reason is that in these states, it becomes very likely that borrowing constraints will bind in the future and the agent will be forced to terminate the mortgage then. However, an anticipated future mortgage default creates an incentive to default today to save the difference between the mortgage payment and the rent in the meantime. This also explains why unemployment, which is persistent, shifts the default frontier to the right.

Fourth, over time, the default region shrinks. This is mainly due to the effect of inflation, which diminishes the real difference between the effective mortgage payments and rents. This has two implications. First, a liquidity-constrained borrower cannot increase current consumption much by defaulting on the mortgage. Second, staying in

the home eventually dominates renting in all states because the real value of the mortgage payment falls below the real rent.

In order to better understand default behavior over the life-cycle of a loan, figure 6 presents the cumulative default rate for loans originated in 2002. This is the cohort for which I have the longest time dimension and on which the model is estimated. Accordingly, the dynamics of default over the life-cycle of this cohort are captured relatively well by the model.

Figure 6: Cumulative Default Rates of 2002 Cohort: Model vs. Data



Though this cohort faces rising average house prices during the immediate time after origination, as seen in figure 1(b), some individuals experience falling house prices and negative equity as a consequence of individual house price shocks. Households with negative equity default when prolonged stretches of unemployment have exhausted their buffer-stock savings, cf. the default region of the state space in figure 5. In fact, more than 99% of all borrowers in this cohort who default are unemployed when they default. This number is similar in the other loan cohorts and never falls below 93%. Thus, the presented model does indeed micro-found the double-trigger hypothesis, cf. the discussion in section 6.3 below.

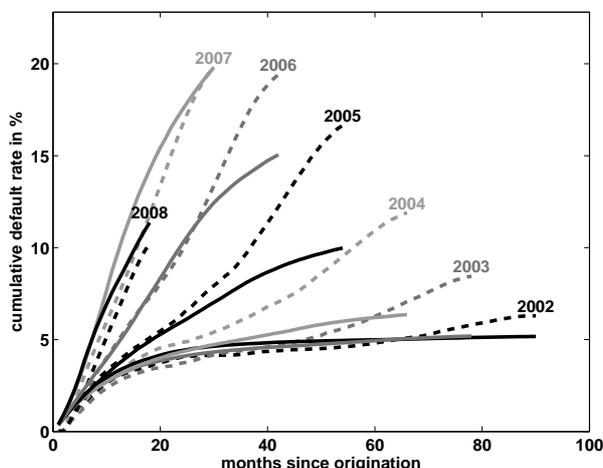
Eventually, the cumulative default rate levels off due to two reasons. First, borrowers who are still active have amortized their mortgages sufficiently, such that most have positive equity. Second, due to the mortgage tilt effect the difference between the real

mortgage payment and real rents shrinks over time, such that a default becomes less appealing.

6.2 The Rise in Cumulative Default Rates

The next step is to compare the default behavior of different cohorts during the time period of the U.S. mortgage crisis. Figure 7 and table 3 compare model predictions and empirical observations on cumulative default rates for cohorts of loans originated each year between 2002 and 2008. The model can explain the broad pattern in the data. The more adverse house price paths of later cohorts cause more borrowers to have negative equity. In the model, the borrowers with negative equity who also experience liquidity problems due to unemployment default on their mortgage. This means the model attributes the rise in cumulative default rates across cohorts to the different aggregate house price paths observed in figure 1(b).

Figure 7: Cumulative Default Rates of all Cohorts: Model (solid lines) vs. Data (dashed lines)



During roughly the first two years after loan origination, the model accurately predicts the observed increase in default rates across loan cohorts. In contrast, during the later period after loan origination, the increase in default rates across cohorts is too weak in the model relative to the data. In the model, this is due to a strong effect of inflation, the mortgage tilt effect. This effect diminishes the difference between real mortgage payments and rents over time. The model is sensitive to this difference. As a consequence, default

Table 3: Cumulative Default Rates (in %) of all Cohorts at certain Points since Origination: Model vs. Data

Cohort		Months since Origination							
		6	18	30	42	54	66	78	90
2002	Data	1.5	3.6	4.2	4.4	4.6	5.1	5.8	6.3
	Model	2.0	4.0	4.6	4.9	5.0	5.1	5.2	5.2
2003	Data	1.3	3.4	4.2	4.7	5.6	7.1	8.4	
	Model	1.9	3.7	4.3	4.6	4.8	5.0	5.1	
2004	Data	1.7	4.3	5.4	7.0	9.5	11.9		
	Model	1.9	3.8	4.7	5.4	6.0	6.4		
2005	Data	1.9	5.1	7.9	12.2	16.6			
	Model	2.0	4.9	7.0	8.8	9.9			
2006	Data	2.1	7.3	13.4	19.4				
	Model	2.5	7.5	12.5	15.2				
2007	Data	3.2	11.8	19.8					
	Model	3.9	14.4	20.1					
2008	Data	2.3	10.2						
	Model	4.0	11.3						

rates do not react strongly enough to falling house prices in periods long after loan origination.

Indeed, the model provides a much better fit to the data when it is calibrated to a lower inflation rate, cf. online appendix C.4. One possible interpretation is that in the real world, borrowers do not fully understand the mortgage tilt effect, and they underestimate it relative to the rational agent in the model. Another fact is that during the 2008 to 2010 period, the inflation rate was, on average, only about 1.4%, which is much lower than the calibrated value of 2.4%. This could also explain the discrepancy between the model and the data, particularly if during the crisis, borrowers expected inflation to be low for an extended period of time in the future.

Another plausible explanation is that other life events, such as marital break-up or other income or expenditure shocks that were excluded from the model, could be responsible for default in later periods. The structural model only analyzes whether and how unemployment shocks can act as the trigger event and finds that they can definitely

play an important role, especially during the early months after origination. However, assessing the role of other life events and decomposing actual default rates into the different causes within the double-trigger paradigm is an interesting area for future research. With respect to divorce, I have also analyzed one alternative specification that included this risk in the model and calibrated it to observed divorce rates. When repeating the analysis, the fit to the data was very similar to that observed for the benchmark model, presumably because observed divorce rates are relatively small.

6.3 A Lesson for Building Empirically Accurate Micro-Founded Models

The above results show how one can broadly explain the rise of default rates during the mortgage crisis within a structural double-trigger model. Important elements of this model are obviously the unemployment shocks and a liquidity constraint, which are also included in a mortgage context in a few other papers cited in the introduction. However, during the investigation, it became clear that these assumptions are not sufficient to yield an empirically accurate double-trigger model. This insight is important for future work and is not contained in the prior literature. The problem is that a sharp decline in house prices may also induce employed and not currently liquidity-constrained homeowners to default. This will then lead to a too strong increase in default rates across cohorts in the model, which I document in detail in the online appendix C.3. Accordingly, there needs to be a reason that prevents employed agents from defaulting after a strong fall of house prices. In the presented model the direct utility benefit from living in the bought house plays this role. Such a feature relates to what are often called the transaction costs of default, which could, for example, reflect factors such as moving costs, reduced credit scores and the psychological stigma of default. Mortgage researchers have long been aware that such costs may be important for bringing frictionless default models closer to the data, cf. the survey by Foote and Willen (2018). However, this paper shows that such a feature is also enormously important for truly micro-founding double-trigger behavior and ensuring an accurate sensitivity of default rates to changes in aggregate house prices. This is of great relevance for the use of such structural models in any policy

or macroeconomic risk analysis.

7 Analysis of Two Crisis Mitigation Policies

During mortgage crises, governments frequently implement mitigation policies. Presumably, the reason is that governments are concerned about a destabilization of the financial system due to the massive losses that mortgage lenders incur from mortgage default in such situations. This section applies the presented structural model to policy analysis in this context. For simplicity, I study a situation where the government decides to neutralize all losses of lenders using a suitable bailout policy, and I ask the question: Should the government bail out lenders or subsidize homeowners for not defaulting? Thus, I only analyze the relative choice between these two options and not whether stabilization policies should be implemented at all. I also do not attempt to provide a full welfare analysis and instead focus on which of these policies has a lower cost for taxpayers. Alternatively, one can view this comparison as an estimate of how much of a subsidy would the government or lenders themselves need to make to potentially defaulting homeowners in order to prevent a given amount of losses to the financial system?

The two analyzed policy measures do not directly resemble specific policies enacted during the recent crisis. However, they are arguably related, as actual stabilization policy has encompassed both types of measures. The Troubled Asset Relief Program (TARP) may, for instance, be viewed as a bailout to lenders because it allowed the government to buy distressed mortgage-backed securities, and to provide guarantees and loans under conditions favorable to banks. In contrast, the Home Affordable Modification Program (HAMP) helped borrowers to receive loan modifications that lowered their current mortgage payments because the government shared the costs this created for lenders. Thus, this program may be viewed as a subsidy to homeowners to prevent default. The assumption that the government wants to neutralize all losses of lenders is of course stark and may seem unrealistic. Here this assumption is made on purpose to facilitate the analysis. The reason is that this allows me to compare the costs of the two policies in a situation where both achieve exactly the same goal. In reality, governments could of course scale

down these policies in a suitable way.

In the event that lenders are bailed out, the government needs to cover the negative equity of defaulters, i.e. the amount by which the outstanding mortgage balance exceeds the value of the collateral. In contrast, the government could also give subsidies to homeowners who would otherwise default, such that they continue to service the mortgage. An alternative interpretation of such a subsidy policy is a temporary mortgage payment reduction for the borrower. This policy might well be cheaper because homeowners are willing to accept some negative equity and thus bear some of the losses on the house's value unless they face severe liquidity problems. The subsidies then only have to overcome the temporary liquidity shortage to neutralize the losses for lenders. However, it is also possible that subsidizing homeowners simply delays default to a later period, and thus the subsidy policy ends up being more expensive in the long run. These opposing effects make a quantitative analysis desirable.

The two policies are compared by calculating the average cost per borrower who would default in the absence of an intervention. For the bailout of lenders, this simply amounts to the average negative equity of a defaulter, which can be readily computed during the simulation. For the subsidy policy, I calculate for each potential defaulter the minimum subsidy amount required to make the borrower stay in the house.²⁶ When doing this, the standard policy functions are used. This means borrowers will consume out of the subsidy, but further negative incentive effects are ruled out. The total sum of all subsidies to a cohort is divided by the number of defaulters without any intervention to make it comparable to the other bailout policy. The required real payment streams of both policies are compared by calculating present discounted values using the real interest rate r . Of course, the accuracy of these calculations depends on the model capturing actual default

²⁶In order to define the subsidy more formally, denote the optimal default policy of a borrower by $d(X_{it}, L_{it}, P_{it}, t)$, which is a function of his state variables liquid wealth X_{it} , employment status L_{it} , house price P_{it} and time t . Let $d(X_{it}, L_{it}, P_{it}, t) = 0$ denote the decision to service the mortgage and $d(X_{it}, L_{it}, P_{it}, t) = 1$ to default on the mortgage. When a borrower receives a subsidy, his liquid wealth increases and this affects his willingness to default. The minimum subsidy S_{it}^* for an individual borrower i in period t is calculated as follows. If borrower i defaults given his state variables ($d(X_{it}, L_{it}, P_{it}, t) = 1$), then the minimum subsidy S_{it}^* is calculated as the smallest possible subsidy amount S_{it} such that this borrower does not default when receiving this amount ($d(X_{it} + S_{it}, L_{it}, P_{it}, t) = 0$). For a non-defaulting borrower ($d(X_{it}, L_{it}, P_{it}, t) = 0$) the subsidy S_{it}^* is zero. For each cohort and time period t , the required individual subsidy amounts S_{it}^* are then summed over all borrowers.

behavior. Thus, it is an important advantage that the model is broadly consistent with empirical evidence. There is also an argument to be made that one should focus more on the earlier cohorts, which are observed for more time periods, in order to accurately account for the delayed default effect of the subsidy.

Table 4 presents the results of this analysis for the different cohorts. Bailing out lenders implies average real present discounted costs between 4.5% and 9.8% of the initial house price per borrower who defaults. By contrast, subsidizing homeowners costs, on average, only between 0.6% and 1.0% of the initial house price in real present discounted value terms. Depending on the cohort, bailing out lenders is thus between 7.1 and 9.9 times more expensive than subsidizing homeowners. In other words, for each dollar spent on subsidies to potentially defaulting homeowners, one can prevent about 7 – 10 dollars of losses to the financial system. These are large differences.

Table 4: Costs of Different Mitigation Policies

Cohort	2002	2003	2004	2005	2006	2007	2008
Bailout to Lenders	4.5	4.6	5.3	6.7	8.3	9.8	7.7
Subsidy to Borrowers	0.6	0.6	0.7	0.7	0.9	1.0	0.9
Ratio Bailout / Subsidy	7.1	7.4	8.1	9.0	9.5	9.9	9.0

Notes: Rows 1 and 2 present the average real discounted cost of the respective policy per borrower who would default without an intervention, expressed in percent of the initial house price. Row 3 reports the ratio between row 1 and row 2.

A couple of caveats apply to this policy analysis. First, these are partial equilibrium results. Still, it seems that general equilibrium effects of subsidizing homeowners would also be more favorable because keeping borrowers in their houses avoids downward pressure on house prices due to foreclosure sales. Second, the subsidy would also help lenders to avoid further administrative costs related to foreclosures and housing sales. Both of these points further strengthen the case for the subsidy.

However, there are also reasons to believe that the costs of the subsidy might be underestimated in these calculations, or at least that a real-world implementation of this policy would need to pay attention to further details. One is that there may be practical problems and high informational requirements associated with implementing such an

individually targeted minimum subsidy to homeowners. Other concerns are related to moral hazard issues. The subsidy could make unemployed borrowers more reluctant to accept new job offers and thus prolong their unemployment spells, cf. Mulligan (2012) for arguments on adverse labor market effects from mortgage modification. However, this problem could potentially be addressed by making the subsidy policy conditional on the borrower exerting a reasonable job search effort and accepting job offers he receives. In the long run both policies may also have negative incentive effects on the screening efforts of lenders and may lead to more risky loans. Thus, these calculations are probably most accurate for a situation where the government surprises private agents with such policies, which are then implemented temporarily during a crisis and only applied to old loans and not new ones.

Several of the points raised above would merit further investigation in future work. Nevertheless the magnitude of the differences between policies suggest that there is indeed considerable potential to improve on the practice of simply bailing out lenders, and such improvements may help to reduce the taxpayer cost of mitigating a mortgage crisis.

8 Conclusions

In its analysis, this paper has explicitly included the economic structure of different default theories and has presented simulations of these models for the observed path of aggregate house prices and a realistic microeconomic distribution. Theoretical predictions were then compared to data on default rates on prime fixed-rate mortgages to assess the explanatory power of the theories during the U.S. mortgage crisis. This comparison revealed that the frictionless default theory is too sensitive to the mean shifts in the house price distribution observed in recent years. By contrast, the double-trigger hypothesis, attributing default to the joint occurrence of negative equity and a life event, is consistent with the dynamics of default rates during the U.S. mortgage crisis.

Based on this finding, a structural dynamic stochastic model with liquidity constraints, unemployment shocks and a direct utility benefit of owning the bought house was presented to provide micro-foundations for the double-trigger hypothesis. In this model,

the liquidity problems associated with unemployment act as a trigger event for default in negative equity situations. The direct utility flow from living in the bought house prevents employed non-liquidity-constrained borrowers from defaulting after a strong fall in house prices. Both features are important for micro-founding double-trigger behavior. The model is broadly consistent with the data and explains most of the rise in mortgage default rates as a consequence of aggregate house price dynamics. This accurate sensitivity of default rates to changes in aggregate house prices is a key requirement of making such models useful for policy or macroeconomic risk analysis.

The structural model was used to formally analyze two important ways of implementing mitigation policies in a mortgage crisis, namely a bailout to lenders or a subsidy to homeowners to prevent default. These abstract policy options are broadly related to actual policies such as the Troubled Asset Relief Program (TARP) and the Home Affordable Modification Program (HAMP). The analysis shows that if the government desires to neutralize losses for lenders, then subsidizing homeowners is about 7–10 times cheaper than bailing out lenders when liquidity problems are a key determinant of mortgage default. These results suggest that there is great potential to reduce the taxpayer cost of mitigating a mortgage crisis for taxpayers. In future research, extended versions of the model could also be used to analyze a range of related policy measures meant to prevent or mitigate mortgage crises. One example is how the design of unemployment insurance may help to prevent mortgage default. In such analyses, it would also be useful to pay more attention to incentive and information problems associated with these policies.

The paper provides evidence that the observed aggregate house price dynamics play a very important role in the rise of mortgage default during the U.S. mortgage crisis. Together with the presented evidence on stable loan characteristics, these findings caution against attributing too much of the increase in default just to lax lending standards. Though the extreme movements in house prices were a rare historical event, the reaction of borrowers can be explained quantitatively by the double-trigger theory of mortgage default. This finding may also help policy makers to draw lessons from the recent crisis for the prevention of future mortgage crises.

References

- Adelino, M., K. Gerardi, and P. S. Willen (2013). Why don't Lenders renegotiate more Home Mortgages? Defaults, Self-Cures and Securitization. *Journal of Monetary Economics* 60(7), 835–853.
- Amromin, G. and A. L. Paulson (2009). Comparing Patterns of Default among Prime and Subprime Mortgages. *Federal Reserve Bank of Chicago Economic Perspectives Q2*, 18–37.
- Bajari, P., S. Chu, and M. Park (2010). An Empirical Model of Subprime Mortgage Default from 2000 to 2007. *Working Paper*.
- Bhutta, N., J. Dokko, and H. Shan (2017). Consumer Ruthlessness and Mortgage Default during the 2007 to 2009 Housing Bust. *Journal of Finance* 72(6), 2433–2466.
- Calhoun, C. A. (1996). *OFHEO House Price Indexes: HPI Technical Description*. Washington, D.C.: Office of Federal Housing Enterprise Oversight.
- Campbell, J. Y. and J. F. Cocco (2003). Household Risk Management and Optimal Mortgage Choice. *Quarterly Journal of Economics* 118, 1449–1494.
- Campbell, J. Y. and J. F. Cocco (2015). A Model of Mortgage Default. *Journal of Finance* 70(4), 1495–1554.
- Carroll, C. D. (1997). Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis. *Quarterly Journal of Economics* 112(1), 1–55.
- Case, K. E., R. J. Shiller, and A. K. Thompson (2012). What have They been thinking? Homebuyer Behavior in Hot and Cold Markets. *Brookings Papers on Economic Activity Fall 2012*, 265–298.
- Chatterjee, S. and B. Eyigungor (2015). A Quantitative Analysis of the U.S. Housing and Mortgage Markets and the Foreclosure Crisis. *Review of Economic Dynamics* 18(2), 165–184.

- Corbae, D. and E. Quintin (2015). Leverage and the Foreclosure Crisis. *Journal of Political Economy* 123(1), 1–65.
- Corradin, S. (2014). Household Leverage. *Journal of Money, Credit and Banking* 46(6), 567–613.
- Davis, M. A., A. Lehnert, and R. F. Martin (2008). The Rent-Price Ratio for the Aggregate Stock of Owner-Occupied Housing. *Review of Income and Wealth* 54(2), 279–284.
- Deaton, A. (1991). Saving and Liquidity Constraints. *Econometrica* 59(5), 1221–1248.
- Demyanyk, Y. and O. Van Hemert (2011). Understanding the Subprime Mortgage Crisis. *Review of Financial Studies* 24(6), 1848–1880.
- Deng, Y., J. M. Quigley, and R. Van Order (2000). Mortgage Terminations, Heterogeneity and the Exercise of Mortgage Options. *Econometrica* 68(2), 275–307.
- Duffie, D. and K. J. Singleton (1993). Simulated Moments Estimation of Markov Models of Asset Prices. *Econometrica* 61(4), 929–952.
- Eberly, J. and A. Krishnamurthy (2014). Efficient Credit Policies in a Housing Debt Crisis. *Brookings Papers on Economic Activity Fall 2014*, 73–136.
- Elul, R., N. S. Souleles, S. Chomsisengphet, D. Glennon, and R. Hunt (2010). What ‘Triggers’ Mortgage Default? *American Economic Review Papers and Proceedings* 100(2), 490–494.
- Ferreira, F. and J. Gyourko (2015). A New Look at the U.S. Foreclosure Crisis: Panel Data Evidence of Prime and Subprime Borrowers from 1997 to 2012. *Working Paper*.
- Foote, C. L., K. S. Gerardi, L. Goette, and P. S. Willen (2008). Subprime Facts: What (We Think) We Know about the Subprime Crisis and What We Don’t. *Federal Reserve Bank of Boston Public Policy Discussion Paper No. 08-2*.
- Foote, C. L., K. S. Gerardi, L. Goette, and P. S. Willen (2009). Reducing Foreclosures: No Easy Answers. *NBER Macroeconomics Annual* 24, 89–138.

- Foote, C. L., K. S. Gerardi, and P. S. Willen (2008). Negative Equity and Foreclosure: Theory and Evidence. *Journal of Urban Economics* 64(2), 234–245.
- Foote, C. L. and P. S. Willen (2018). Mortgage-Default Research and the Recent Foreclosure Crisis. *Annual Review of Financial Economics*. forthcoming.
- Fuster, A. and P. S. Willen (2013). Payment Size, Negative Equity, and Mortgage Default. *Working Paper*.
- Garriga, C. and D. E. Schlagenhauf (2009). Home Equity, Foreclosures, and Bail-outs. *working paper*.
- Gerardi, K. S., K. F. Herkenhoff, L. E. Ohanian, and P. S. Willen (2013). Unemployment, Negative Equity and Strategic Default. *Working Paper*.
- Gerardi, K. S., A. Lehnert, S. M. Sherlund, and P. S. Willen (2008). Making Sense of the Subprime Crisis. *Brookings Papers on Economic Activity* (Fall), 69–159.
- Gerardi, K. S., A. H. Shapiro, and P. S. Willen (2007). Subprime Outcomes: Risky Mortgages, Homeownership Experiences, and Foreclosures. *Federal Reserve Bank of Boston Working Paper No. 07-15*.
- Ghent, A. C. and M. Kudlyak (2011). Recourse and Residential Mortgage Default: Theory and Evidence from U.S. States. *Review of Financial Studies* 24(9), 3139–3186.
- Guiso, L., P. Sapienza, and L. Zingales (2013). The Determinants of Attitudes towards Strategic Default on Mortgages. *Journal of Finance* 68(4), 1473–1515.
- Guvenen, F. (2006). Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective. *Journal of Monetary Economics* 53(7), 1451–1472.
- Gyourko, J. and J. Tracy (2014). Reconciling Theory and Empirics on the Role of Unemployment in Mortgage Default. *Journal of Urban Economics* 80, 87–96.
- Hatchondo, J. C., L. Martinez, and J. M. Sánchez (2015). Mortgage Defaults. *Journal of Monetary Economics* 76, 173–190.

- Herkenhoff, K. F. and L. E. Ohanian (2013). Foreclosure Delay and U.S. Unemployment. *Working Paper*.
- Jagtiani, J. and W. W. Lang (2011). Strategic Default on First and Second Lien Mortgages During the Financial Crisis. *Journal of Fixed Income* 20(4), 7–23.
- Jeske, K., D. Krueger, and K. Mitman (2013). Housing, Mortgage Bailout Guarantees and the Macro Economy. *Journal of Monetary Economics* 60(8), 917–935.
- Joint Center for Housing Studies at Harvard University (2008). *The State of the Nation's Housing*.
- Kau, J. B. and D. C. Keenan (1995). An Overview of the Option-Theoretic Pricing of Mortgages. *Journal of Housing Research* 6(2), 217–244.
- Kau, J. B., D. C. Keenan, and T. Kim (1994). Default Probabilities for Mortgages. *Journal of Urban Economics* 35, 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.
- Kau, J. B., D. C. Keenan, W. J. Muller, and J. F. Epperson (1995). The Valuation at Origination of Fixed-Rate Mortgages with Default and Prepayment. *Journal of Real Estate Finance and Economics* 11(1), 5–36.
- Laufer, S. (2018). Equity Extraction and Mortgage Default. *Review of Economic Dynamics* 28, 1–33.
- Mayer, C., K. Pence, and S. M. Sherlund (2009). The Rise in Mortgage Defaults. *Journal of Economic Perspectives* 23(1), 27–50.
- McFadden, D. (1989). A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration. *Econometrica* 57(5), 995–1026.
- Mian, A. and A. Sufi (2009). The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis. *Quarterly Journal of Economics* 124(4), 1449–1496.

- Mulligan, C. B. (2012). *The Redistribution Recession: How Labor Market Distortions Contracted the Economy*. Oxford University Press.
- Pakes, A. and D. Pollard (1989). Simulation and the Asymptotics of Optimization Estimators. *Econometrica* 57(5), 1027–1057.
- Palmer, C. (2013). Why did so many Subprime Borrowers Default during the Crisis: Loose Credit or Plummeting Prices? *Working Paper*.
- Quercia, R. G. and M. A. Stegman (1992). Residential Mortgage Default: A Review of the Literature. *Journal of Housing Research* 3(2), 341–379.
- Stanton, R. and N. Wallace (2011). The Bear’s Lair: Index Credit Default Swaps and the Subprime Mortgage Crisis. *Review of Financial Studies* 24(11), 3250–3280.
- Vandell, K. D. (1995). How Ruthless is Mortgage Default: A Review and Synthesis of the Evidence. *Journal of Housing Research* 6(2), 245–264.
- Zeldes, S. (1989). Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence. *Quarterly Journal of Economics* 104(2), 275–298.

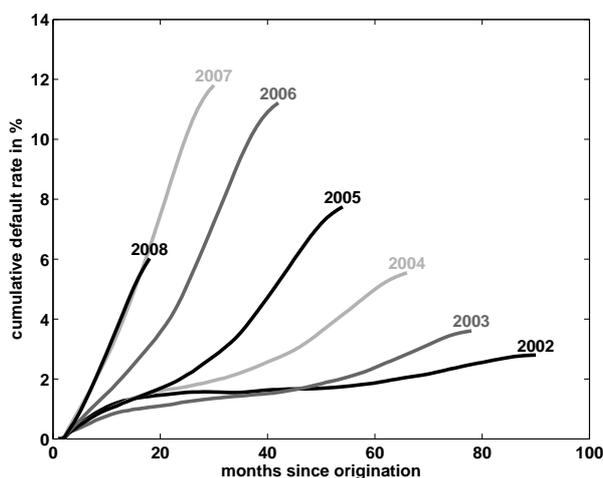
Appendices (Only For Online Publication)

A Appendix to Data and Empirical Facts

The main analysis of the paper is focussed on loans with an initial loan-to-value (LTV) ratio above 95% for the substantive reasons explained in section 2.1. This section shows that the general empirical patterns concerning the rise in default rates and relatively stable loan characteristics at origination across cohorts extend to other loans with lower LTVs as well. In other words the full sample exhibits broadly the same dynamics across cohorts than the sample of highly leveraged borrowers analysed in the paper, though there are of course various level differences.

Figure A.1 presents the observed cumulative default rates for different loan cohorts for the full sample of loans with all possible initial LTVs. Of course the data still refers to prime, fixed-rate, 30-years mortgages as in the main text. Though the level of default rates is in general somewhat lower than for the sample of highly leveraged borrowers, the rise in default rates follows a very similar pattern.

Figure A.1: Cumulative Default Rates in the Full Sample



Evidence on average loan characteristics at origination for different cohorts for the full sample are presented in table A.1. The observed pattern is similar to the one for highly leveraged borrowers in the main text. Average FICO credit scores are almost constant across cohorts. The average loan-to-value ratio changes a bit more across cohorts than

for the above 95% LTV sample. But surprisingly in the full sample many of the later cohorts even have a lower average LTV than the first cohort. The dynamic pattern of mortgage rates and debt-to-income (DTI) ratios is very similar to the sample in the main text. Again changes in average mortgage rates are not closely correlated with the increase in default rates, but average DTIs at origination increase somewhat over time.

Table A.1: Average Loan Characteristics at Origination by Loan Cohort for the Full Sample

Cohort	2002	2003	2004	2005	2006	2007	2008	All
LTV in %	77.9	75.1	75.7	73.8	74.8	76.4	78.4	76.0
FICO score	714	714	712	716	712	710	716	714
Mortg. rate in %	6.7	5.9	6.0	6.0	6.6	6.5	6.2	6.2
DTI in %	34	37	36	37	38	39	38	37

These empirical facts show that the dynamics of default rates and loan characteristics across cohorts are very similar in the full and the more restricted sample. This is evidence against the considered sample being somehow special and suggests that several of the drawn conclusions may well extend to prime, fixed-rate, 30-years mortgages more generally. Furthermore, in section B.4 below I show explicitly that under plausible assumptions on second mortgages the main conclusions of the reduced-form exercise generalize to loans with an initial LTV of the first mortgage between 75% and 84%.

At least for the prime mortgage market the facts documented in this section are prima facie evidence against explanations that emphasize the role of a deterioration in lending standards or increased borrower leverage for the rise in default rates like Corbae and Quintin (2015). For the full sample I have also conducted an additional analysis of compositional changes with respect to initial LTVs. For this purpose I grouped all loans with respect to their initial LTV into 10 bins with a width of 10 percentage points (the two exceptions are the first bin with all LTVs between 1 – 14% and the last bin with all LTVs above 95%). Using the observed default rates of each of these bins for each cohort and the actual LTV composition of each cohort I then computed two counterfactual aggregate default rates for all cohorts, i.e. the equivalent to figure A.1. In the first counterfactual

scenario I keep the default rates of each LTV bin fixed at their observed level of the 2002 cohort and only vary the LTV composition across cohorts according to the actual data. In the second counterfactual I only vary the default rates of each bin across cohorts as observed in the actual data, but counterfactually keep the LTV composition fixed as it is observed for the 2002 cohort. This decomposition analysis shows that for all practical purposes the aggregate default dynamics across cohorts witnessed in figure A.1 are entirely driven by increases in default profiles for each initial LTV bin and not by changes in the composition of initial LTVs. The same conclusion applies if one uses a cohort other than the 2002 one for fixing one of the two margins in the counterfactual calculations. These results reinforce the conclusion in the main paper that the fall in house prices and not compositional effects are key for understanding the rise in default rates at least for prime fixed-rate mortgages during the crisis.

B Appendix to Reduced Form Models

B.1 Estimation Procedure

The model parameters are estimated by a simulated method of moments procedure. Let θ stand in for the parameter to be estimated in the respective model. The idea of the estimation is to choose θ such that the cumulative default rates for the 2002 cohort simulated from the model match as well as possible those observed in the data. Collect the variables d_{it} in one vector $D_i = [d_{i1}, \dots, d_{iT}]'$ for each individual. The mean of this vector $\bar{D} = \frac{1}{N} \sum_{i=1}^N D_i$ represents the empirically observed cumulative default rate. The expected value of D_i is $E[D_i] = D(\theta)$ and denote the expected value evaluated by simulation of S individuals from the model by $\tilde{D}(\theta)$. The deviation of the model from the data is then given by $G(\theta) = \bar{D} - \tilde{D}(\theta)$. The simulated method of moment estimator of θ minimizes $G(\theta)'WG(\theta)$ where W is a weighting matrix. I weight all moments equally by using an identity matrix as the weighting matrix. θ is then estimated by minimizing a least squares criterion function given by

$$\sum_{t=1}^T \left(\bar{d}_t - \tilde{d}_t(\theta) \right)^2 \tag{15}$$

where \bar{d}_t and $\tilde{d}_t(\theta)$ are the t -th element in the vectors \bar{D} and $\tilde{D}(\theta)$, respectively. Here, $\tilde{d}_t(\theta)$ is evaluated using a frequency simulator such that $\tilde{d}_t(\theta) = \frac{1}{S} \sum_{j=1}^S \tilde{d}_{jt}(\theta)$ and $\tilde{d}_{jt}(\theta)$ represents the outcome for period t of applying the decision rules to the drawn history j of the underlying shocks. The minimization problem is solved by a grid search algorithm.

B.2 Robustness Checks

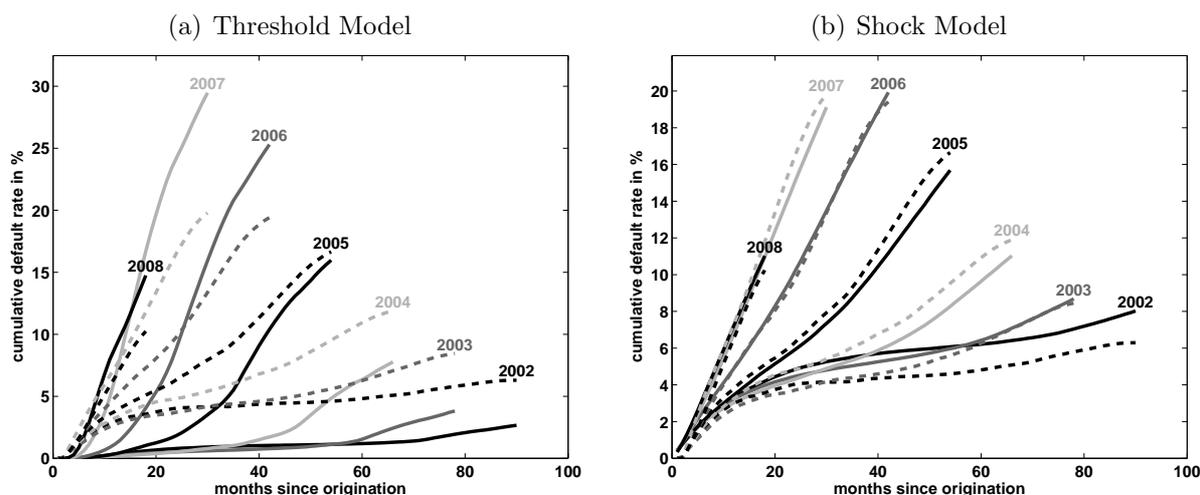
This section reports a battery of robustness checks that were performed to scrutinize the reduced form results. I find that the results are robust across all the modifications considered here. For brevity I do not report the graphs corresponding to figure 3 for all these checks. But these are available upon request.

Instead of estimating the models on the 2002 cohort with low default rates, I also estimate them on the 2008 cohort with very high default rates. This does not affect the good fit of the shock model. However, now the threshold model greatly undershoots the default rates of early cohorts and also still overshoots the 2006 and 2007 cohort. Thus, the comparison across models is unaffected. I have also used the 2003 and 2005 cohorts to estimate the models and always found the same results across the two models.

Another robustness check replaces the out-of-sample test with an in-sample test. Here, I estimate the two models on all cohorts and then examine the fit within that sample. This exercise is informative on the best possible fit both models can give to the data. These results are thus worth to report in more detail. The estimated parameters are then -15.9% for ϕ and 1.36% for ψ , and the results are shown in figure B.1. The threshold model still has considerable problems to match the data even under these most favorable circumstances. It generally undershoots earlier cohorts and the early months after origination for all cohorts and at the same time still overshoots the late months of the 2006 and 2007 cohorts. In contrast, the shock model gives a very good fit to the data. The conclusions across models are essentially unchanged.

I also examine the role of the variation in mortgage rates and the distribution of loan-to-value ratios across cohorts in three alternative specifications. In the first specification, I keep the within cohort LTV distribution fixed across cohorts according to the average

Figure B.1: In-Sample Fit of the Two Models



frequency. The second specification abstracts from within cohort heterogeneity such that everyone has the same LTV according to the respective within cohort average. The third specification is the same as the second except that the LTV and mortgage rate are not varied across cohorts. All these changes have very modest effects on both models and leave the conclusions across models unaffected. This implies that the double-trigger model attributes the rise in default rates to the variation in aggregate house prices and not the changes in contract characteristics across cohorts. It also suggests that abstracting from this heterogeneity across cohorts in the structural model is not too restrictive.

In section 3.4, it was assumed that the individual house price shocks are normally distributed. The major argument supporting this choice is that by the central limit theorem the mean of individual shocks converges asymptotically to a normal distribution anyway. However, since the analysis also covers periods where t is still small, I perform an additional check here. Instead of using a normal distribution for the individual shocks I specify them as being uniformly distributed on the interval $[-b_t, b_t]$. The parameter b_t is then chosen such that the variance of the uniformly distributed shock in period t in the respective census division is identical to the one used in the standard framework. I find that the results are almost identical.

Another potential concern is that the simplicity of the presented reduced-form models with only one constant parameter somehow biased the results against the frictionless

option model. There is also no strong reason why the default threshold parameter ϕ and default shock probability ψ should be constant over the course of a loan. It turns out that the results are robust to changing this assumption. As a check I have performed a scenario where the respective default parameter depends fully on the month since origination t . The constant parameters in the model are then replaced with ϕ_t and ψ_t that are allowed to differ each period from $t = 1, \dots, T$ when fitting the models to the 2002 cohort. Under these circumstances, both models almost perfectly match the 2002 cohort. The cumulative default rates simulated for the other cohorts then inherit the non-smoothness of the first differences of the cumulative default rate of the 2002 cohort. However, subject to that qualification the conclusions on the out-of-sample fit remain essentially unchanged. The threshold model still greatly overshoots the later cohorts. The shock model generates default rates comparable in magnitude to the benchmark.

B.3 Using an alternative definition of default

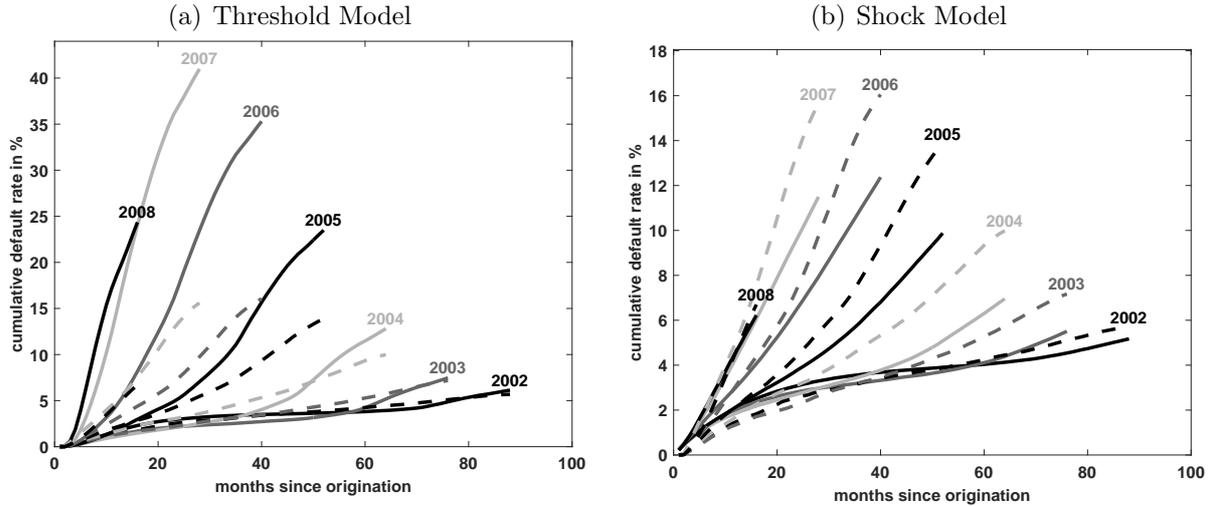
In this section, I use a different definition of default. Instead of using a 60 or more days definition of default as in the main text, I now consider a loan to be in default once it is at least 120 days past due.²⁷ This is a more demanding definition and by all accounts being 120 days past due is considered as a very serious delinquency. This change of definition affects the levels of the data on cumulative default rates which is used to estimate and test the models. The level of default rates is lower now. However, the broad dynamics across cohorts are similar to the ones analyzed in the main text.

Again, I estimate both reduced form models on the 2002 cohort and use the remaining cohorts to test the estimated models. For the threshold model this yields an estimate of ϕ of -11.9% and for the shock model ψ is estimated as 0.83% . The results are reported in figure B.2. These are qualitatively very similar to the ones of the main text and the conclusions across models are unchanged. Though one can debate what is the appropriate definition of default, I conclude from this exercise that this issue is not key for my results.

Furthermore, I have also investigated the effect of using a definition of default that

²⁷Now I backdate the period of default by 3 months to capture the time when the first payment has been missed.

Figure B.2: Using an alternative definition of default (120+ days)



requires a loan to be in foreclosure. This also generates similar results (which are available upon request) and does not resolve the empirical problems of a frictionless option-model documented in section 3.

B.4 Extension to lower Loan-to-Value Ratios

The paper is focussed on loans with an LTV above 95% because these borrowers should be least likely to have a second mortgage on their home, cf. the discussion in section 2.1. The question arises whether the results of the paper also generalize to loans with a lower LTV. This section provides some evidence on this by repeating the reduced-form analysis of section 3 for loans with an LTV of the first mortgage between 75% and 84%. Due to the discussed data problems this section is necessarily somewhat tentative. Nevertheless, some very interesting results emerge.

First, I take the data for the loans with a LTV of the first mortgage between 75% and 84% at face value and assume that no one has a second mortgage. Accordingly, the LTV varies within cohorts in steps of one percentage point between 75% and 84%. Changes to the distribution of loans over this support across cohorts observed in the mortgage data are again taken into account. The mortgage rate is again kept constant within a cohort and set equal to the respective cohort average. When estimating the models on the 2002 cohort, I find that neither of the two models can capture this data well. Even

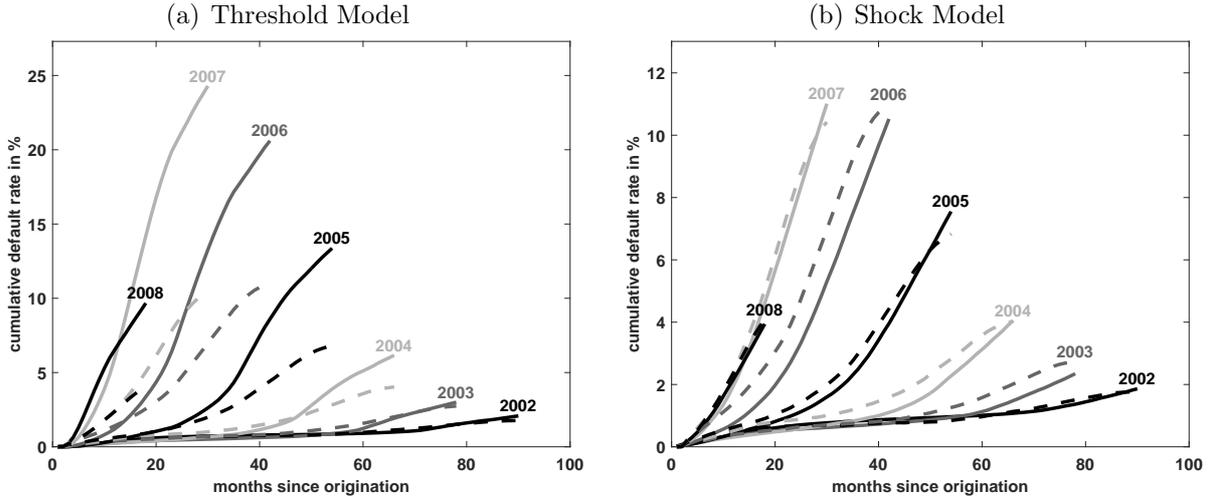
for the most extreme parameter values of $\phi = 0$ and $\psi = 1$, both models undershoot the cumulative default rate of the 2002 cohort substantially for at least the first 60 months after origination. The reason is that the equity buffer generated by the down-payment is substantial for these borrowers. Because the 2002 cohort faced strongly increasing average house prices immediately after origination, too few borrowers in the simulation experience negative equity compared to observed default rates. It is important that both models fail if we take this data at face value. One can draw two possible conclusions from these results. Either we need a completely new theory of default for these loans or it is crucial to take second mortgages into account. I present evidence on the second explanation next.

Elul, Souleles, Chomsisengphet, Glennon, and Hunt (2010) report that 26% of all borrowers have a second mortgage and this adds on average 15% to the combined LTV. However, they neither report a breakdown of these statistics by the LTV of the first mortgage nor when borrowers take out the second mortgage. Faced with this situation, I model a very simple form of intra-cohort heterogeneity taking these estimates of the frequency and size of second mortgages into account. I assume that 74% of borrowers have only one mortgage with a distribution of LTVs as in the mortgage data. However, 26% of borrowers in each cohort independently of the LTV of the first mortgage also have a second mortgage adding 15% to the combined LTV. This implies that the support of the LTV distribution is expanded and now also includes values between 90% and 99%. It is assumed that borrowers got the second mortgage at the same time as the first one and pay the same mortgage rate on both. Admittedly, these are very crude assumptions. This exercise can only provide preliminary evidence until better data is available and should be regarded with considerable caution.

For this setup the reduced-form models are estimated again on the 2002 cohort. This yields estimates of $\phi = -7.8\%$ and $\psi = 2.25\%$. The estimated models are again tested on their ability to predict out-of-sample. Figure B.3 presents the results for all cohorts. The threshold model overshoots the data again. In contrast, the shock model provides an excellent fit to the data. Thus, the double-trigger theory also provides a better expla-

nation for this data under the maintained assumptions on second mortgages. Due to the discussed data problems I would personally put a lower weight on these results compared to the benchmark results. However, these results are at least suggestive that the main conclusions on the relative merit of the two theories may well extend to loans with a lower LTV.

Figure B.3: Results for borrowers with a first mortgage LTV of 75 – 84% taking second mortgages into account



Notes: Solid lines: model. Dashed lines: data.

C Appendix to Structural Model

C.1 Value Functions

The state variables of the optimization problem for an owner are liquid wealth $X_t = A_t + Y_t$, employment status L_t , house price P_t and time t . The choice variables are consumption C_t and the mortgage termination choice. The value function of an owner $V^o(\cdot)$ can then be written as

$$V^o(X_t, L_t, P_t, t) = \max \left\{ V^s(X_t, L_t, P_t, t), V^r(X_t + P_t - \frac{M_t}{\Pi_t}, L_t, t), V^r(X_t, L_t, t) \right\}$$

which reflects the optimal choice between staying in the house with value $V^s(X_t, L_t, P_t, t)$, selling with value $V^r(X_t + P_t - \frac{M_t}{\Pi_t}, L_t, t)$ and defaulting with value $V^r(X_t, L_t, t)$. Selling and defaulting involve a permanent transition to the rental market. In case of staying

the value $V^s(X_t, L_t, P_t, t)$ is given by

$$\begin{aligned}
V^s(X_t, L_t, P_t, t) &= \max_{C_t} \left\{ \frac{C_t^{1-\gamma}}{1-\gamma} + \theta + \beta \mathbb{E}_t [V^o(X_{t+1}, L_{t+1}, P_{t+1}, t+1)] \right\} \\
\text{s.t. } X_{t+1} &= (1+r) \left(X_t - \frac{m}{\Pi_t} + \tau r^m \frac{M_t}{\Pi_t} - C_t \right) + Y_{t+1} \\
C_t &\leq X_t - \frac{m}{\Pi_t} + \tau r^m \frac{M_t}{\Pi_t}.
\end{aligned}$$

The value function of a renter $V^r(X_t, L_t, t)$ is given by

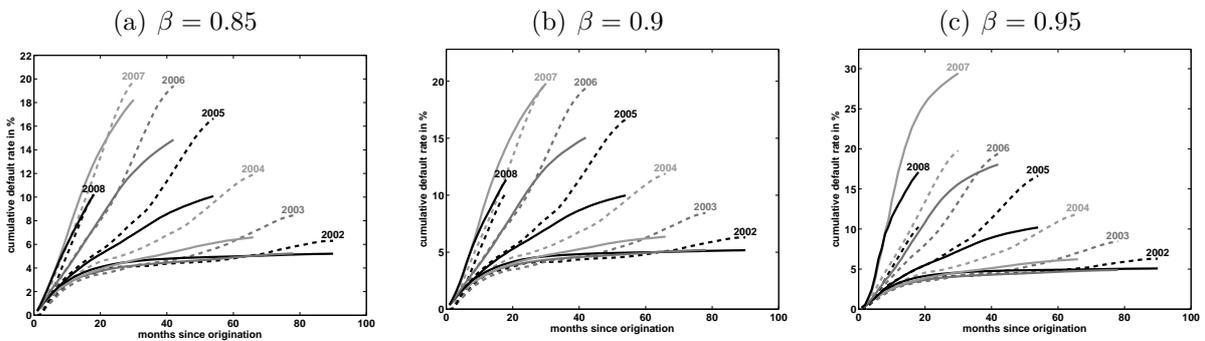
$$\begin{aligned}
V^r(X_t, L_t, t) &= \max_{C_t} \left\{ \frac{C_t^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}_t [V^r(X_{t+1}, L_{t+1}, t+1)] \right\} \\
\text{s.t. } X_{t+1} &= (1+r) (X_t - R - C_t) + Y_{t+1} \\
C_t &\leq X_t - R.
\end{aligned}$$

C.2 Dependence on Preference Parameters

This section explores how the model depends on the predetermined preference parameters. Specifically, I compute results for alternative parameter values for β and γ in order to get an idea how the model behaves in different parts of the parameter space. The benchmark preference parameter values are $\beta = 0.9$ and $\gamma = 5$.

First, I consider alternative values of β equal to 0.85 and 0.95. For these value of β , the parameter θ is then reestimated in order to fit the 2002 cohort. This yields values of θ of 0.39 and 0.16 respectively. The results of these experiments are compared to the benchmark results in figure C.1.

Figure C.1: Sensitivity to Preference Parameter β

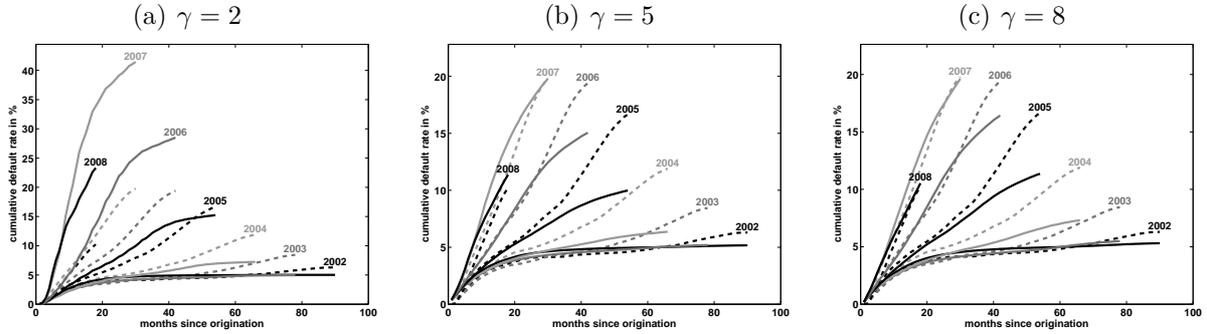


Notes: Solid lines: model. Dashed lines: data.

Next, I consider alternative values of γ equal to 2 and 8. θ is then estimated as

0.06 and 0.64 respectively. Figure C.2 compares these alternative calibrations to the benchmark.

Figure C.2: Sensitivity to Preference Parameter γ



Notes: Solid lines: model. Dashed lines: data.

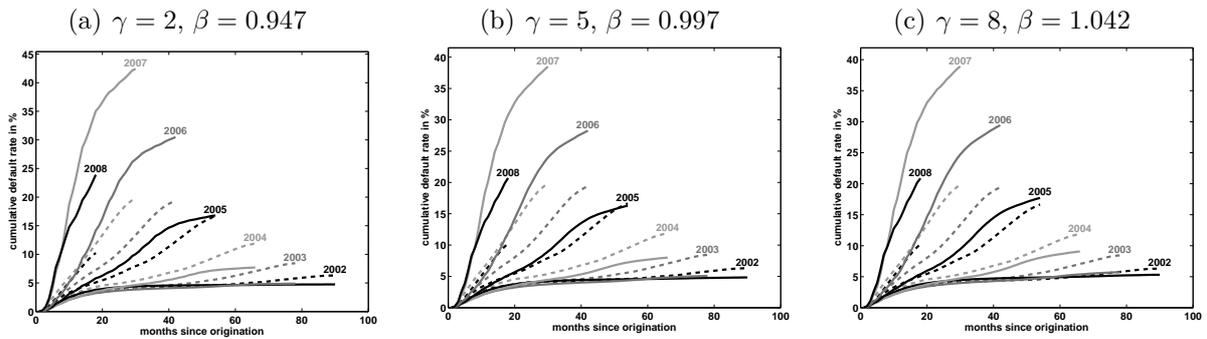
These results show that the model works as well or better than in the benchmark calibration for higher values of γ and/or lower values of β . These parameter changes make the agent less willing to substitute intertemporally and/or more impatient to consume today. This worsens the liquidity problem caused by unemployment. The model can only feature double-trigger behavior when being employed and being unemployed are sufficiently different. In contrast, for lower values of γ and higher values of β temporary income reductions can more easily be smoothed out. The model then implies that a sizeable portion of employed agents default in all cohorts. This brings the model too close to a frictionless option model and the model then inherits all the problems of such a specification witnessed already in section 3.

C.3 Details on the Role of a Direct Utility Benefit

In this section, I show that versions of the model which abstract from a direct utility benefit of living in the bought house are excessively sensitive to changes in aggregate house prices. As an illustration I conduct simulations of the model where the utility flow parameter θ is set to 0. I then adjust the discount factor β to again match the default rates of the 2002 cohort for $\theta = 0$. Figure C.3 presents these results for the benchmark value of the CRRA coefficient $\gamma = 5$, but also for values of γ of 2 and 8. Each of those experiments is associated with a different value of β . These results confirm the key role of including a direct utility benefit of living in the bought house in order to prevent employed

and not currently liquidity-constrained agents from defaulting when house prices decline strongly. Without this feature the model cannot accurately capture the sensitivity of default rates to aggregate house prices.

Figure C.3: Performance of the model without a direct utility benefit of living in the bought house ($\theta = 0$) for different values of γ



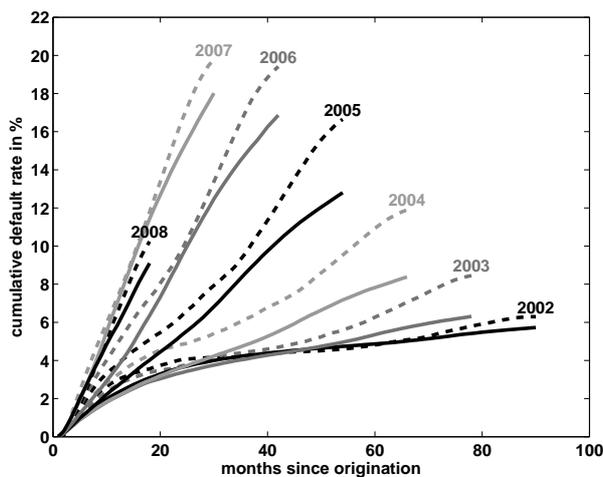
Notes: Solid lines: model. Dashed lines: data.

C.4 Role of Inflation

In this section, I show that the mortgage tilt effect caused by inflation plays an important role for the performance of the model in later periods after origination. In the benchmark calibration, the inflation rate is set to 2.4%, which is the average value during the simulation period (2002-2010). In order to document the sensitivity of default rates to the inflation rate I investigate an alternative calibration where π is set ad-hoc to 1%. This alternative calibration is only meant to be an illustration how default decisions are affected when borrowers either expect the inflation rate to be a bit lower (1.4% lower) during the simulation period and the future than it was during these years, or in some other way underestimate the mortgage tilt effect. Unfortunately I have no data on how large such a deviation of borrower expectations from observed inflation may have been in reality, which implies that I can only look at an ad-hoc scenario. All other parameters are unchanged, but θ is reestimated at a value of 0.44 to fit the 2002 cohort.

Figure C.4 presents the results of this exercise. The fit of the model improves in the later period after origination relative to the benchmark results. This shows that the strength of the mortgage tilt effect is responsible for the problems of the model in the main text to capture default decisions in later periods after origination.

Figure C.4: Performance of the Model for a Lower Inflation Rate



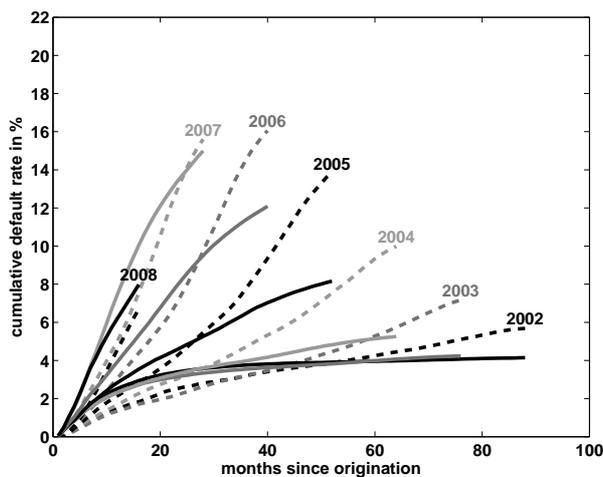
Notes: Solid lines: model. Dashed lines: data.

Using this alternative calibration I have also repeated the policy analysis from section 7. This allows to check how the policy results change in a model that captures the data even better than the benchmark (but admittedly makes an ad-hoc assumption on the inflation rate). The absolute costs of both policies tend to increase a bit relative to the benchmark and the relative cost of the bailout to lenders also increases. Bailing out lenders is then between 8.6 and 11.2 times more expensive than subsidizing homeowners using this alternative calibration. Thus, the conclusions across policies are robust or even strengthened relative to the benchmark calibration.

C.5 Using an alternative definition of default

This section reports the results for the structural model when the 120 days definition instead of the 60 days definition is used to measure default empirically as in section B.3. All other procedures are as in the main text. The estimate of θ is then 0.32. Figure C.5 reports the results for all cohorts. The fit of the model to this data is qualitatively similar to the one of the main text that uses a 60 days default definition. The results of the policy analysis are also essentially unchanged and bailing out lenders is found to be between 8.7 and 11.2 times more expensive than subsidizing homeowners. This analysis shows that the main results of the structural model are robust to using an alternative reasonable definition of default.

Figure C.5: Model Results for the 120+ Days Default Definition



Notes: Solid lines: model. Dashed lines: data.

C.6 A Structural Single-Trigger Model

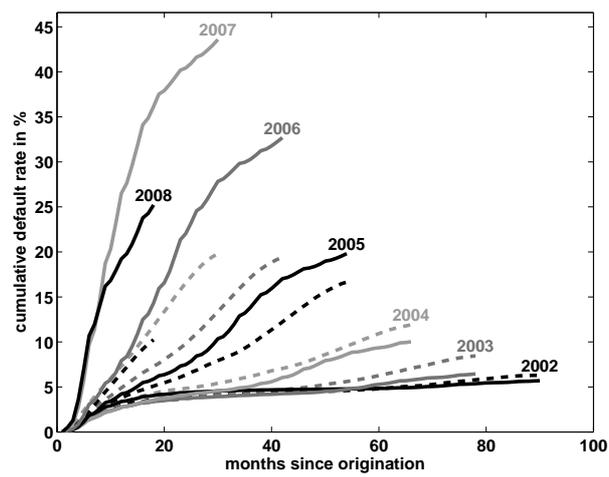
This section investigates a structural single-trigger model in the spirit of the frictionless option theoretic literature. The aim is to confirm that the conclusions on the relative merit of the two theories drawn in the reduced form section 3 also carry over to a comparison of structural models of these theories.

For this purpose, it is convenient that the structural double trigger model of sections 4 and 5 in fact nests a single-trigger model if one suitably modifies some model parameters. The single-trigger paradigm relies on the assumption that borrowers have access to perfect and complete financial markets such that they can perfectly insure against income fluctuations. Accordingly, mortgage default decisions are unaffected by income fluctuations and liquidity problems. The structural model nests such a case when the following changes to model parameters are made. The net replacement rate in case of unemployment ρ is set to 1 such that income is identical in the employed and unemployed state. Given the actual replacement rate of unemployment insurance and the unemployment rate implied by the job separation and finding rates from the benchmark calibration, this insurance costs about 2% of net income when employed. Accordingly, employed net income $(1 - \tau)Y_0$ is lowered by 2% such that it is net of insurance fees, which is achieved by resetting the debt to income ratio to a slightly higher value of $DTI = 0.4085$. These

modifications are equivalent to the borrower buying insurance against the part of income risk not covered by the public unemployment insurance system. In the resulting model, default is no longer driven by employment status because income does not vary across employment states. The second modification is to delete the borrowing constraint of equation (13) from the model such that it features a perfect credit market for unsecured credit. In practice this is achieved by changing equation (13) to $A_{t+1} \geq -\xi$ and then setting ξ to a large number instead of zero as in the benchmark calibration. Specifically, I set ξ equal to five times annual net income and confirm that during the simulation the optimal asset path chosen by borrowers remains far away from this constraint. The other features, parameters and initial conditions of the model remain the same as in the benchmark model and calibration. Thus, the structural single trigger model is put as much as possible on an equal footing with the structural double trigger model. The direct utility benefit parameter θ is then estimated at a value of -0.035 to match the cumulative default rate of the 2002 cohort, which this model specification also matches very well.

Figure C.6 presents the resulting fit of the structural single-trigger model to all cohorts. The model strongly overpredicts the rise in default rates. Its performance is broadly comparable to the one of the reduced-form single-trigger model. When comparing the structural single and double-trigger models to each other, one again finds that the double-trigger model explains the data much better than the single trigger model. The impression from visually inspecting figure C.6 and figure 7 of the main text is confirmed by simple measures of goodness of fit of the model predictions to the data. For instance the root mean squared error (mean absolute error) of the out-of-sample forecast is 3.5 (2.8) times higher for the structural single-trigger than for the structural double-trigger model. Thus, this exercise using structural models confirms the reduced form results on the relative merit of the two theories.

Figure C.6: Performance of a Structural Single-Trigger Model



Notes: Solid lines: model. Dashed lines: data.