No. 72

UNIVERSITY OF COLOGNE

WORKING PAPER SERIES IN ECONOMICS

MORTGAGE DEFAULT DURING THE U.S. MORTGAGE CRISIS

THOMAS SCHELKLE

Department of Economics
University of Cologne
Albertus-Magnus-Platz
D-50923 Köln
Germany

http://www.wiso.uni-koeln.de

Mortgage Default during the U.S. Mortgage Crisis*

Thomas Schelkle†

This Version: May 16, 2014
First Version: October 2011

Abstract

Which of the main competing theories of mortgage default can quantitatively explain the rise in default rates during the U.S. mortgage crisis? This paper finds that the double-trigger hypothesis attributing mortgage default to the joint occurrence of negative equity and a life event like unemployment is consistent with the evidence. In contrast a traditional frictionless default model predicts a too strong increase in default rates. The paper also provides microfoundations for double-trigger behavior in a model where unemployment may cause liquidity problems for the borrower. Using this framework for policy analysis reveals that a mortgage crisis may be mitigated at a lower cost by relieving the liquidity problems of borrowers instead of bailing out 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 PhD thesis at the London School of Economics. I am grateful to Francesco Caselli and Rachel Ngai for many helpful discussions and advice. I also thank Wouter den Haan, Alexander Ludwig, Albert Marcet, Xavier Mateos-Planas, Stephan Seiler, Jonathan Temple and seminar participants at Aachen, Cologne, LSE, 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

After 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 then contributed to the start of the recent financial and economic crisis. These events highlight the importance of understanding the economic mechanisms triggering mortgage default and the rise in default rates. Insights into these issues may then inform political debates on how to prevent future mortgage crises or mitigate ones 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 Unites States between 2002 and 2010.

The paper considers the two major mortgage default theories - the frictionless option-theoretic model and the double-trigger hypothesis. The traditional frictionless literature assumes that borrowers "ruthlessly" default on their mortgage to maximize their financial wealth.[1] 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 of the borrower.

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 it attributes default to the joint occurrence of negative equity and a life event like unemployment or divorce. The double-trigger hypothesis is well-known among mortgage researchers.[2] But it is usually only discussed verbally or with stylized models.

These two microeconomic theories are tested using the observed variation in aggregate house prices and default rates across cohorts of loans of

[1] An example is the paper by Kau, Keenan, and Kim (1994). The surveys of Quercia and Stegman (1992) and Vandell (1995) provide further discussion and references.

[2] Discussions can for example be found in Gerardi, Shapiro, and Willen (2007), Foote, Gerardi, and Willen (2008) and Foote, Gerardi, Goette, and Willen (2009).

prime fixed-rate mortgages with high initial loan to value ratios.[3] In the data mortgage borrowers who experienced a more adverse path of average house price growth rates defaulted much more frequently.[4] *Qualitatively* both theories are consistent with this observation. However when I simulate and estimate reduced form models of the two theories and assess their ability to predict out-of-sample, this reveals important *quantitative* differences between the theories. I find that the frictionless theory is excessively sensitive to changes in aggregate house prices and predicts a far too strong rise in default rates. In contrast, the double trigger hypothesis is consistent with the evidence. The economic reason is that default rates have increased roughly in proportion to the number of borrowers who experience any level of negative equity as predicted by the double-trigger theory. In contrast, the predictions of the frictionless theory are based on the number of homeowners experiencing extreme levels of negative equity. During the crisis this number has increased much more strongly across loan cohorts than actual default rates.

Based on this finding the second part of the paper micro-founds the double-trigger hypothesis in a structural dynamic stochastic partial equilibrium model of mortgage default. Borrowers in that 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 for formally analyzing two possible mitigation policies in a mortgage crises that may help to stabilize the finan-

[3]Specifically, I focus on loans with an initial loan to value ratio above 95%. The reasons for this are explained in detail in section 2. However I also briefly show in an extension that under plausible assumptions the results generalize to loans with lower initial loan to value ratios.

[4]The variation in default rates across cohorts does not seem to be driven by observed differences in loan and borrower characteristics as I document empirically in section 2.

cial system. If the government desires to neutralize the losses of mortgage lenders from default, it could either bail out the lenders or mitigate the liquidity problems of homeowners who would otherwise default. The analysis shows that a subsidy policy to homeowners is the cheaper option in this model where default is partly driven by liquidity problems.

The paper relates to different strands of prior theoretical and empirical work. The structural model of the paper builds on previous theoretical work by Campbell and Cocco (2003, 2011) and Corradin (2014) who also model liquidity constraints in a mortgage framework. Other papers use equilibrium models to examine the role various institutional features like bail-out guarantees or mortgage product innovation and falling house prices played for the mortgage crisis including Chatterjee and Eyigungor (2011), Corbae and Quintin (2011), Jeske, Krueger, and Mitman (2013) and Garriga and Schlagenhauf (2009). In contrast to this related work, my paper focusses on comparing theoretical models on the decision to default in detail to empirical observations for different loan cohorts. Though these papers also include income shocks and liquidity constraints in a mortgage framework, my analysis reveals that this 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 feature double-trigger behavior and are consistent with the evidence. Furthermore I use the micro-founded double-trigger model to formally evaluate different mitigation policies in a mortgage crisis.

The paper is also related to a vast empirical literature that studies the determinants of mortgage default.[5] This literature provides a wealth of evi-

[5]Studies 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 then caused an enormous increase in empirical work on mortgage default. Examples of this empirical research include Amromin and Paulson (2009), Bajari, Chu, and Park (2010), Bhutta, Dokko, and Shan (2010), Demyanyk and Van Hemert (2011), Elul, Souleles, Chomsisengphet, Glennon, and Hunt (2010), 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) and Palmer (2013),

dence that negative equity or falling house prices are strong determinants of default. Some studies have also investigated the role of life events as triggers for default. Many studies 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 at the individual level that unemployment and income shocks increase the probability of default. My paper is motivated by these prior empirical results.[6] But it uses a very different methodology. Specifically, I simulate theoretical models for the observed aggregate house price paths and realistic microeconomic house price distributions and compare these predictions to empirical observations. This reveals the excess sensitivity of a purely negative equity threshold based default theory to changes in aggregate house prices and the empirically 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 is therefore complementary to the prior empirical literature.

The paper is structured as follows. Section 2 describes the data and empirical facts on 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 for 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 and the key facts the paper attempts to explain. Information on mortgage contract

among others.

[6]Another interesting empirical fact is the great heterogeneity in default behavior for borrowers with the same level of negative equity (Quercia and Stegman 1992). The structural double-trigger model of this paper can rationalize this fact because in the model the default threshold of negative equity depends on liquid wealth and employment status. Individual heterogeneity in these variables, which are unobserved in all standard mortgage data sets, may then account for the heterogeneity in default behavior of borrowers with the same level of negative equity.

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 focussed on loans originated between 2002 and 2008.

I restrict the sample to prime, first, fixed-rate, 30-years mortgages that have a standard amortization schedule (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 these are the most common mortgage contracts. The data base contains around 23 million loans with these characteristics in 2010.[7] I further focus the analysis on loans with a loan-to-value ratio (LTV) above 95%, which depending on the year represents about $20 - 30\%$ of all 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 the model. This is important due to the highly non-linear relationship between default decisions and negative equity in the theoretical models. Furthermore, the loans with a high LTV default most frequently, so it makes sense to focus an analysis of mortgage default on them. But 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.[8] 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. However

[7] Amromin and Paulson (2009) estimate that the LPS data cover about 60% of the prime market between 2004 and 2007.

[8] Elul, 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 and this adds on average 15% to the combined LTV. Unfortunately, they do not report a break-down of these statistics by the LTV of the first mortgage.

in online appendix A.4 I show 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%.

The data set contains for each loan cohort (defined by origination month) over time how many active loans 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.[9] However I also show in online appendices A.3 and B.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.[10]

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 evidence discussed in section 3.4.

The key empirical facts on mortgage default rates and house prices

[9]The period of default is backdated by one month to capture the time when the first payment has been missed.

[10]The reason for also looking at a 120 days definition is that a considerable fraction of loans which 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 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. This means that there exists a certain tension between the theoretical and empirical literature with respect to the used concept of default. The rationale for also looking at a 120 days definition is then 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 on the correspondence between theoretical and empirical concepts of default. In this check I find that my results are robust to using this alternative reasonable 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 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.

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.[11] 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 is the main aim of the paper.

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

(a) Cumulative Default Rates (b) Mean Real House Price Paths

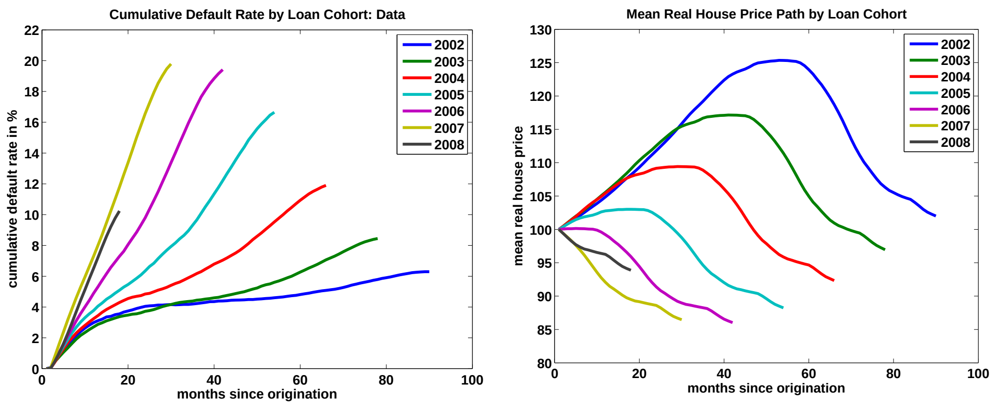

Before proceeding to the analysis I briefly discuss one alternative explanation for the rise in default rates observed in figure 1(a). This explanation is that lending standards and loan quality deteriorated sharply before the mortgage crisis. Thus, I first present evidence that average loan quality is fairly stable across cohorts in my data set.[12]

One concern is that the loan-to-value ratio (LTV) might have increased

[11]In 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 shown curves for an origination year are averages of the underlying twelve cohorts defined by origination month.

[12]It is also important to remember that I only look at data on prime fixed-rate mortgages. Therefore any compositional shifts that occurred in the mortgage market towards subprime lending or variable rate mortgages do by construction not affect my analysis. We see clearly from figure 1(a) that mortgage default rates have increased substantially even without such compositional effects.

over time leaving a smaller buffer before borrowers experience negative equity. I only consider loans that have a 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.[13]

The second row of table 1 reports the average FICO credit score at origination of the different loan cohorts. 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 as such less attractive to the borrower. There is some variation in this variable across cohorts. But the mortgage rate and default rates seem to be fairly uncorrelated across cohorts.

The average debt-to-income (DTI) ratio representing the share of the required mortgage payment in gross income is presented in the last row of table 1.[14] This has increased over time indicating that borrowers in later cohorts need to devote more of their gross income to service the mortgage. But the increase was quite modest.

These statistics show that there is no evidence in favor of a strong deterioration of lending standards over time in my data set of prime fixed-rate mortgages with a LTV above 95%.[15] This casts doubts on explanations

[13]When I analyse the reduced-form models in section 3 I even control for changes between cohorts in the within-cohort distribution of LTVs and find that the observed changes are irrelevant for the models considered here.

[14]The 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 only cover prime, fixed-rate, 30-years 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 to the actual loan pool I consider.

[15]This conclusion might be specific to the prime market. For example Demyanyk and Van Hemert (2011) present evidence that loan quality deteriorated in the subprime market. But Amromin and Paulson (2009) also note that it is less obvious that a similar deterioration was present in the prime market. A particular advantage of my descriptive statistics is that they are based on all loans in the LPS data base satisfying 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.

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	40	40	40	40	42	42	40

of the mortgage crisis that rely solely on lax lending standards. Instead this paper shows that the fall in house prices can explain the rise in default rates within a formal model.

3 Reduced Form Models

This section presents evidence on mortgage default from 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). The aim is to discriminate between these theories and their key mechanism in a relatively general way that is independent of the exact specification of the respective structural model. Building on these results the following section then develops a structural economic model.

3.1 Model Setup

The paper considers individual borrowers who took out a fixed-rate 30-years mortgage. Each loan cohort defined by origination date consists of many borrowers who are indexed by $i = 1, \ldots, N$ and observed in periods $t = 1, \ldots, T$ after loan origination. Borrowers take 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 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 defaulted yet.

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. m_i are fixed nominal monthly payments covering mortgage interest and principal. These payments are determined at the beginning of the contract and satisfy

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

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 heterogenous with respect to the LTV. It is assumed that agents take 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 an individual homeowner evolves stochastically over time as described in section 3.4 below. P_{i0} is normalized to 100.

3.2 The Threshold Model

The first model assumes that borrowers with negative equity default on their mortgage at 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 $\phi < 0$. If in period t the borrower has not defaulted yet then the default decision in that period is described by

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

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. 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 to receive a default shock ψ is constant and satisfies $0 \leq \psi \leq 1$ and default shocks are independently and identically distributed over time. If the borrower has not defaulted yet, the default decision in period t is determined by

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

This is a reduced-form of a double-trigger model. Here the default shock represents a life event that combined with negative equity triggers default. The parameter ψ represents the probability that a life event occurs and is estimated from the data. Possible examples for 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 modelled and simulated. The information in this subsection applies also to the structural model in the following section. The general aim is to base the simulation framework for house prices as closely as possible 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

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

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 because otherwise theoretical models cannot explain any default during times of positive aggregate house price growth. The formulation is consistent with the approach used by the FHFA to estimate the house price index, cf. the description in Calhoun (1996).[16]

In equation (5) a census division index was suppressed for convenience. But 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. But the FHFA provides estimates 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 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 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

[16] I 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 $\ln(P_{it}) = \ln(P_{i,0}) + \sum_{\tau=1}^{t} g_\tau^{agg} + \sum_{\tau=1}^{t} g_{i\tau}^{ind}$ where $\ln(P_{i,0}) + \sum_{\tau=1}^{t} g_\tau^{agg} = \beta_t + N_i$ and $\sum_{\tau=1}^{t} g_{i\tau}^{ind} = H_{it}$ give equation (1) in Calhoun (1996).

a quadratic formula in time for the variance of the total individual part of the house price change since purchase given by

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

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 then given by

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

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 the variance of a single g_{it}^{ind} is decreasing over time.[17]

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. For the shock model I also draw default shock histories.[18] When computing the mortgage balance the mortgage rate is kept constant within a cohort and set equal to the respective cohort average. But borrowers within a cohort are heterogenous with respect to the initial LTV which varies in steps of one percentage point between 95% and 104%.[19] The frequency of these different loan-to-value ratios at origination is varied

[17]On 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.

[18]Specifically, I draw histories from an i.i.d. uniform distribution on the interval $[0, 1]$. For a given parameter ψ the default shock occurs for the respective individual and month if the uniform draw is smaller or equal to ψ.

[19]The few loans with a LTV above 104% are subsumed under the 104% LTV group.

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 is used to compute the real mortgage balance. The decision rules are then applied to these shock histories and paths of the real mortgage balance.

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 where 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.[20] 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. It consists of informally comparing simulated and empirically observed default rates for these remaining cohorts and checking which estimated model gives 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 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.

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 over-predicts default

[20]The online appendix A.1 provides details. The simulated method of moments was developed by McFadden (1989), Pakes and Pollard (1989) and Duffie and Singleton (1993).

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

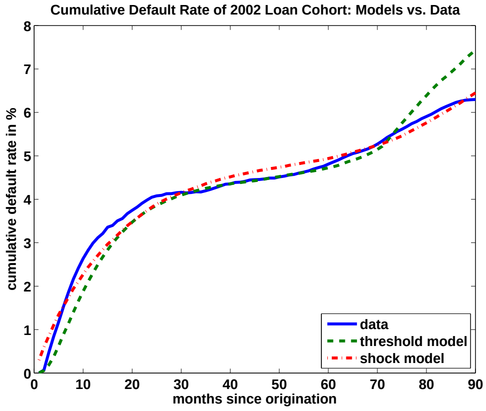

rates for the later cohorts in the simulation period by at least one order of magnitude. The threshold model is excessively sensitive to the shifts in the mean of the house price distribution observed in the data. In contrast, the shock model gives a good fit to the broad dynamics in the data.[21]

Figure 3: Cumulative Default Rates for all Cohorts: Models vs. Data

(a) Threshold Model (b) Shock Model

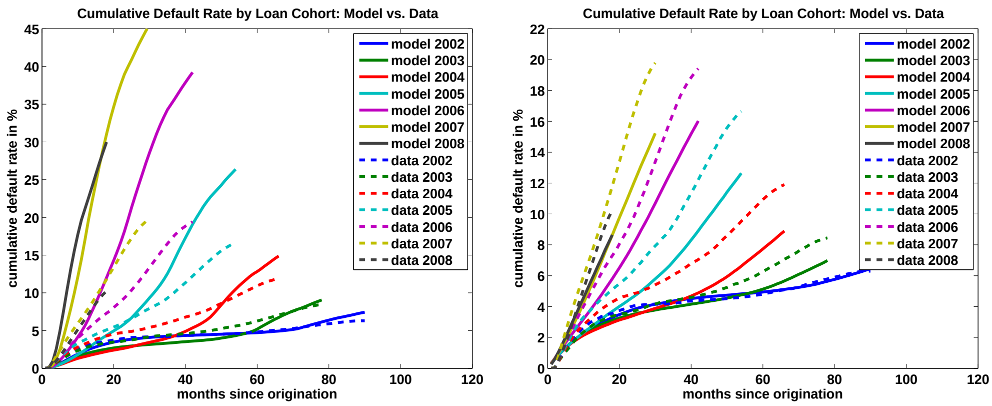

The explanation for the difference between models is the following. The

[21]Admittedly the shock model generates a slightly too low increase in default rates. It is interesting in this context to examine the empirical fit of the model when it is estimated on all cohorts, which is provided in appendix A.2. This shows that for the resulting parameter estimate the model is able to almost perfectly explain the 2003 to 2008 cohorts, while overshooting only the 2002 cohort. This is a hint that there may indeed be a small difference between the borrowers of the 2002 cohort and the other cohorts though they appear to be similar based on observed characteristics.

shock model predicts that a fraction ψ of borrowers with negative equity default each period. When the whole equity distribution shifts left due to the fall in aggregate house prices, the shock model predicts that the default rate should increase in proportion to the increase in the number of borrowers who experience negative equity. It turns out that observed default rates approximately exhibit this pattern. But the threshold model is concerned with the far left tail of the equity distribution. It predicts that all borrowers with an extreme level of negative equity below ϕ times the initial house price default. When the equity distribution shifts left the number of borrowers with such an extreme level of negative equity increases faster than the observed default rate. This generates the inconsistency with the data.

I have conducted a large number of robustness checks to scrutinize these results, which are reported in detail in online appendices A.2, A.3 and A.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.

Two conclusions can be drawn from the results of this section. First, an empirically successful structural model cannot rely on a single-trigger or negative equity threshold mechanism alone. Instead some feature other than house price shocks must play a role. Second, in 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 in aggregate house prices is key for understanding the observed rise in default rates.

4 Structural Model

This section introduces a micro-founded model of double-trigger behavior where unemployment acts as the life event that may trigger default together with negative equity. The aim of the structural model is to analyze whether and how unemployment may play this role, 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 with a fixed-rate mortgage each period chooses non-housing consumption and 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 defaulting on the loan by "walking away". The homeowner faces uncertainty on the future price of the house, unemployment shocks and a borrowing constraint for unsecured credit. 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 instalments 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 for a direct utility benefit of owning a house. Household decisions over the length of the mortgage contract are determined by maximizing expected utility given by

$$U = \mathrm{E}_0 \sum_{t=1}^{T} \beta^{t-1} \left(\frac{C_t^{1-\gamma}}{1-\gamma} + \theta \mathcal{I}(own_t) \right) + \beta^T \frac{W_{T+1}^{1-\gamma}}{1-\gamma} \tag{8}$$

which is derived from consumption C_t in periods 1 to T and remaining wealth W_{T+1} at the end of the contract.[22] 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. $\mathcal{I}(own_t)$ 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 house or the benefit that an owner cannot be asked to move out by a landlord as may happen to a renter.

In each period the homeowner has to decide how much to consume and on staying or leaving the house. If the agent wants to leave this can be done by either selling the house (and repaying the current mortgage balance) or defaulting on the loan by "walking away".[23]

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

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

where A_t denotes real asset holdings and Y_t 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. But the nominal mortgage interest $r^m M_t$ is tax deductable and τ is the tax rate. All nominal variables

[22] Following 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 including a retirement period. This had no significant effect on the results.

[23] The model abstracts from mortgage termination through refinancing for computational reasons. 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.

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 case 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

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

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 life.[24] 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 extended period of the time and experiences a strong fall in his credit rating. This is one of the costs of defaulting from the borrower's point of view. 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

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

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 at least since 1975. In this paper I take these observations as given and specify the exogenous variables of the model accordingly. But explaining this pattern is an important area for future research. However a fully realistic specification would also require to make α cohort-specific. But I use an approximation for computational reasons such that α is constant across cohorts and calibrated to a suitable average.

In contrast, if the agent decides to default on the mortgage by "walking away"[25] 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 borrowing constraint for unsecured credit[26] 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 period depend on current wealth and the house tenure choice.

Remaining wealth at the end of the contract for a homeowner 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 exogenously given by

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

[25]The 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 borrowers from defaulting who own relatively high valued properties (above $\$200,000$ in real 2005 terms). Since my data set contains borrowers with much lower house values the neglect of recourse does not seem to be a major concern.

[26]In modelling borrowing constraints the model builds on the buffer-stock saving framework of Zeldes (1989), Deaton (1991) and Carroll (1997).

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 $(1 - f)$. I focus on income fluctuations due to unemployment risk here because unemployment involves a severe fall in labor income from one month to another. This makes it a very plausible cause for 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.[27]

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. But the house price realizations in the simulation of the model are generated from the division specific data and distributions.

[27] My formulation abstracts from deterministic changes to labor income like a life-cycle profile 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 period until buffer stock savings are exhausted and default occurs.

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 $DTI = \frac{m}{Y_0}$.[28] I assume that borrowers were employed when they got their loan. 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[29] and then 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 in informally comparing the out-of-sample model predictions on 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 a set

[28]The name debt to income ratio is standard mortgage terminology, but can be easily misunderstood. It means the ratio of the monthly mortgage payment to gross income.

[29]Computational details are relegated to appendix B.2

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 attention to 30-years ($T = 360$ 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 before, 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 0.8% and a yearly standard deviation of 1.9%. This means that agents in the model have expectations on 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 the ones 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 calculator[30] 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 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[31] 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.[32] I then set $s = 1.8\%$ and $f = 31\%$ which are the average values of the computed monthly finding and separation probabilities during 1990 to 2010. These values imply a steady state unemployment rate around 5.7%.

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 this 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).

[30] http://www.oecd.org/social/soc/benefitsandwagestax-benefitcalculator.htm

[31] Since the data on median unemployment duration is reported in weeks, I first transform it to months by multiplying the weekly value by 12/52.

[32] As 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 gives 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. But adapting Guvenen's argument to β, the reason is that I am analyzing borrowers who were only able to make a very small down-payment. This could be due to the fact that they are very impatient. The agents who are net savers could then have a higher discount factor. In the online appendix B.3 I investigate how the results depend on the specific choice of β and γ.[33]

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. As in the section on reduced-form models the parameter is chosen such that cumulative default rates simulated from the model match those observed in the data using only information from the 2002 cohort. This yields an estimate for θ of 0.28. The remaining data is again used to test the ability of the estimated model to predict out of sample.

6 Results

This section first explains the mechanism generating default in the model. Then the main results how well the model fits the rise in default rates across loan cohorts are presented.

6.1 The Default Mechanism

The repayment policy function of a borrower in the model is presented in figure 4 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. Instead, with positive equity selling is strictly preferred to defaulting because the borrower is the residual claimant of the house 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

[33]Results 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 which is 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 witnessed already in section 3.

Figure 4: Repayment Policy Function

(a) Employed in $t = 1$

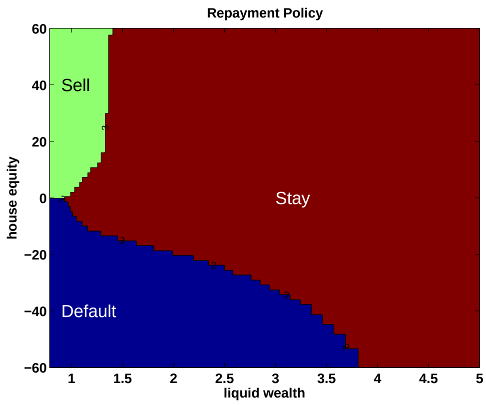

(b) Unemployed in $t = 1$

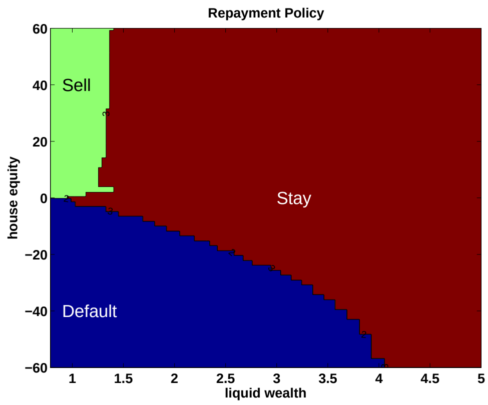

(c) Employed in $t = 20$

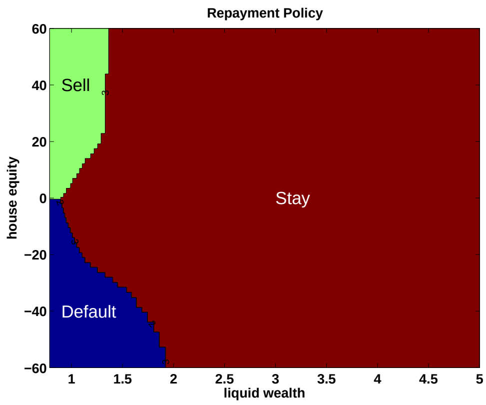

(d) Unemployed in $t = 20$

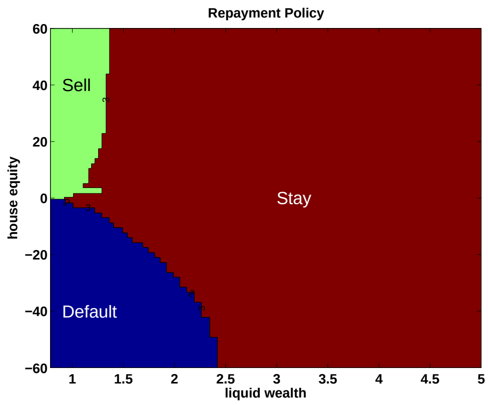

Notes: Repayment choice as a function of the state variables liquid wealth, house equity, employment status and time. Blue region: Default. Green region: Sell. Red region: Stay.

benefit of staying is that the borrower receives the utility benefit of owning a house and keeps the option to default, sell or stay later. Specifically, there are possible future states of the world with positive equity. But the probability of reaching these 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 the mortgage to increase current consumption. The other reason is that in these states it becomes very likely that borrowing constraints bind in the future and the agent is forced to terminate the mortgage then. But an anticipated future mortgage default creates an incentive to default already 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 that 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 a mortgage default. 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 5 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 5: Cumulative Default Rates of 2002 Cohort: Model vs. Data

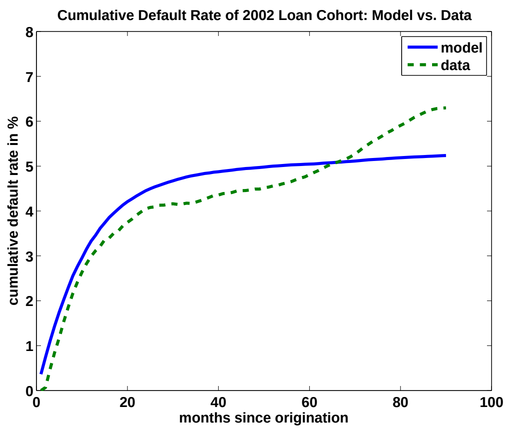

Though this cohort faces growing average house prices during the immediate time after origination as seen in figure 1(b), some individuals experience falling house prices and negative equity as a consequence of individual house price shocks. Households with negative equity default when prolonged stretches of unemployment have exhausted their buffer stock savings, cf. the default region of the state space in figure 4. 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.

Eventually, 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 6 compares 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 witnessed in figure 1(b).

The model is particularly successful in the early months after loan origination, but tends to predict a bit too infrequent default in later months. 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 and reacts too strongly compared to the data.[34] One possible interpretation is

[34]In the online appendix I show that the model gives a better fit to the data when it is calibrated to a lower inflation rate. It is also noteworthy that especially in the final years of the simulation period inflation was lower than the constant benchmark value

Figure 6: Cumulative Default Rates of 2002-2008 Cohorts: Model vs. Data

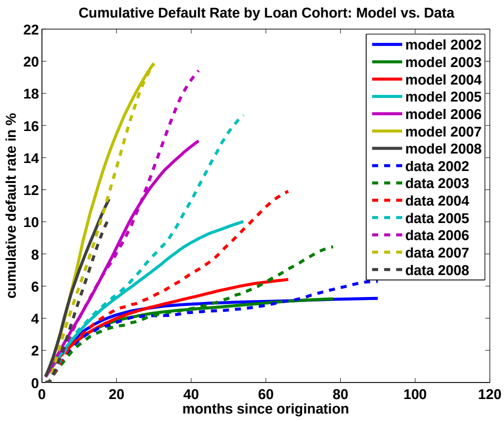

that in the real world borrowers do not fully understand and underestimate this effect relative to the rational agent in the model. Another plausible interpretation is that other life events like marital break-up or other income or expenditure shocks that were excluded from the model could be responsible 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 could definitely play an important role especially during the early months after origination. But assessing the role of other life events and a decomposition of actual default rates into the different causes within the double-trigger paradigm is an interesting area for future research.

7 Analysis of Two Crisis Mitigation Policies

This section applies the presented structural model for policy analysis. I study a situation where the government is concerned about a destabilization of the financial system due to the losses that mortgage lenders incur from mortgage default. Assume that the government decides to neutralize all these losses by a suitable bailout policy. The question is then: Should the government bail out lenders or homeowners?

In case lenders are bailed out the government needs to cover the negative equity of defaulters, i.e. by how much the outstanding mortgage balance

of 2.4%. On average between 2008 and 2010 inflation was 1.4%. It is possible that the model would perform better for these actual inflation rates.

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. 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 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 such that 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 absence of an intervention. For the bailout of lenders this simply amounts to the average negative equity of a defaulter which can readily be 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. 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 these calculation can only be as accurate as the model captures actual default behavior. There is also an argument to be made to focus more on the earlier cohorts that are observed for more time periods in order to accurately account for the delayed default effect of the subsidy.

Table 3 presents the results of this analysis for the different cohorts. Bailing out lenders implies average real present discounted costs between 4.5% and 9.7% of the initial house price per borrower who defaults. In contrast subsidizing homeowners on average only costs 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 9.8 times more expensive than subsidizing homeowners. These are large differences.

A couple of comments on these results are in order. First, these are partial equilibrium results. But it seems that general equilibrium effects

Table 3: Costs of Different Mitigation Policies

Cohort	2002	2003	2004	2005	2006	2007	2008
Bailout to Lenders	4.5	4.7	5.4	6.8	8.3	9.7	7.7
Subsidy to Borrowers	0.6	0.6	0.7	0.7	0.9	1.0	0.9
Ratio Bailout / Subsidy	7.1	7.6	8.2	9.1	9.5	9.8	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.

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 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 for example make unemployed borrowers more reluctant to accept new job offers and prolong their unemployment spells. 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 and not new loans.

Finally, one needs to keep in mind that I have only analyzed the choice between these two policies here and not the question whether the government should conduct such stabilization policies at all. Several of the raised points would merit further investigation. Nevertheless the above calculations using a model which is broadly consistent with empirical evidence are at the very least suggestive that there is considerable potential to improve on simply bailing out lenders in order to reduce the cost of mitigating a mortgage crisis for taxpayers.

8 Conclusions

This paper has presented simulations of theoretical default 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. In contrast, the double-trigger hypothesis attributing default to the joint occurrence of negative equity and a life event is consistent with the data.

Based on this finding a structural dynamic stochastic model with liquidity constraints and unemployment shocks was presented to provide microfoundations 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 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.

The structural model was used to formally analyze two mitigation policies in a mortgage crisis. If the government desires to neutralize losses for lenders then subsidizing homeowners is much cheaper than bailing out lenders when liquidity problems are a key determinant of mortgage default. A related policy question to which the model could be applied in future work is how the design of unemployment insurance may help to prevent mortgage default.

The paper provides evidence that the observed aggregate house price dynamics play a very important role for the rise in mortgage default during the U.S. mortgage crisis. Together with the presented evidence on stable loan characteristics this cautions against attributing too much of the increase in default just to lax lending standards. Though the extreme movements of house prices were a rare historical event, the reaction of borrowers can be explained by a well-known theory of mortgage default. This finding may also help 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? Redefaults, 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 (2010). The Depth of Negative Equity and Mortgage Default Decisions. *Working Paper*.

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 (2011). A Model of Mortgage Default. *Working Paper*.

Carroll, C. D. (1997). Buffer-Stock Saving and the Life-Cycle/Permanent Income Hypothesis. *Quarterly Journal of Economics 112*(1), 1–55.

Chatterjee, S. and B. Eyigungor (2011). A Quantitative Analysis of the U.S. Housing and Mortgage Markets and the Foreclosure Crisis. *working paper*.

Corbae, D. and E. Quintin (2011). Mortgage Innovation and the Foreclosure Boom. *working paper*.

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.

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.

Foote, C. F., 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. F., K. S. Gerardi, L. Goette, and P. S. Willen (2009). Reducing Foreclosures: No Easy Answers. *NBER Macroeconomics Annual 24*, 89–138.

Foote, C. F., K. S. Gerardi, and P. S. Willen (2008). Negative Equity and Foreclosure: Theory and Evidence. *Journal of Urban Economics 64*(2), 234–245.

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.

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.

Kau, J. B., D. C. Keenan, and T. Kim (1994). Default Probabilities for Mortgages. *Journal of Urban Economics 35*, 278–296.

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.

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.

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 Reduced Form Models

A.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 $D_i = [d_{i1}, \ldots, d_{iT}]'$ for each individual. The mean of this vector $\overline{D} = \frac{1}{N}\sum_{i=1}^{N} D_i$ 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 $\widetilde{D}(\theta)$. The deviation of the model from the data is then given by $G(\theta) = \overline{D} - \widetilde{D}(\theta)$. 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

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

where $\overline{d_t}$ and $\widetilde{d}_t(\theta)$ are the t-th element in the vectors $\overline{D}$ and $\widetilde{D}(\theta)$, respectively. Here $\widetilde{d}_t(\theta)$ is evaluated using a frequency simulator such that $\widetilde{d}_t(\theta) = \frac{1}{S}\sum_{j=1}^{S} \widetilde{d_{jt}}(\theta)$ and $\widetilde{d_{jt}}(\theta)$ 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.

A.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. But 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. In fact 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 -13.8% for ϕ and 1.32% for ψ and the results are shown in figure 7. 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 an excellent fit to the data. The conclusions across models are essentially unchanged.

Figure 7: In-Sample Fit of the Two Models

(a) Threshold Model (b) Shock Model

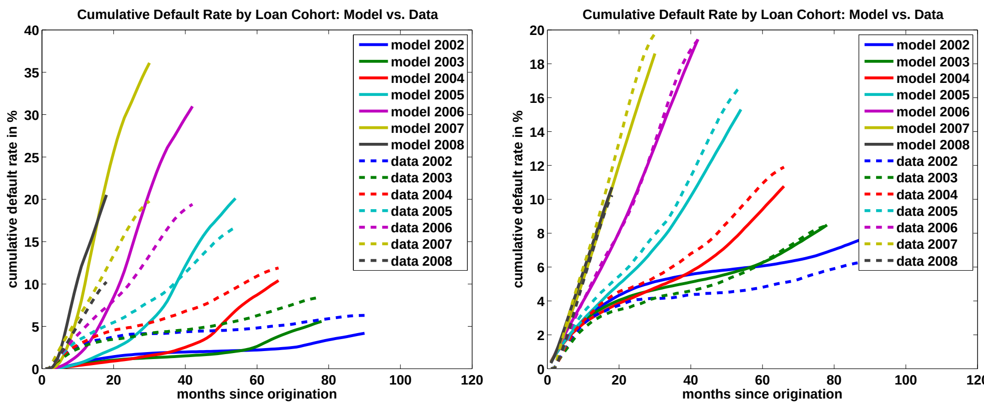

I also examine the role of the variation in mortgage rates and the distribution of loan-to-value ratios across cohorts in three alternative specifications. In the first specification, I keep the within cohort LTV distribution fixed across cohorts according to the average frequency. The second specification 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. But 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 $t = 1, \ldots, T$ 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. But 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.

A.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.[35] 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. Obviously 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 8. 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 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.

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

The paper is focussed on loans with a LTV above 95% because these borrowers should be least likely to have a second mortgage on their home, cf. the discussion in section 2. The question arises whether the results of the paper also generalize to loans with a lower LTV. This section provides

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

Figure 8: Using an alternative definition of default (120+ days)

(a) Threshold Model (b) Shock Model

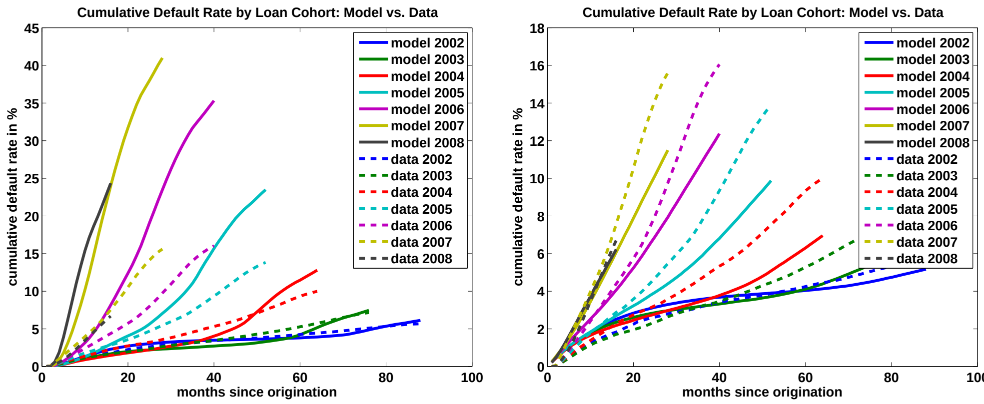

some evidence on this by repeating the reduced-form analysis of section 3 for loans with a LTV of the first mortgage between 75% and 84%. Due to the discussed data problems this section is necessarily somewhat tentative. Nevertheless, some very interesting results emerge.

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. But they neither report a break-down 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. But 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 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 9 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 explanation 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. But 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.

B Appendix to Structural Model

B.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

Figure 9: Results for borrowers with a first mortgage LTV of $75-84\%$ taking second mortgages into account

(a) Threshold Model (b) Shock Model

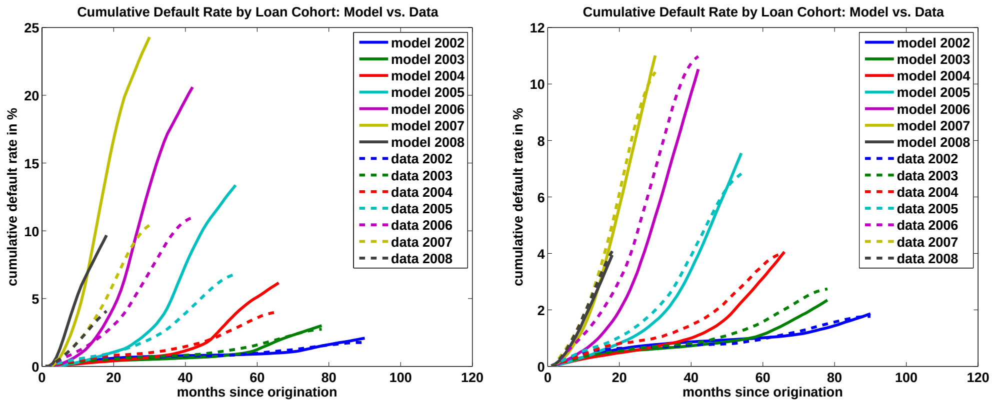

choice variables are consumption C_t and the mortgage termination choice. The value function of an owner $V^o(.)$ can then be written as

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

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

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

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

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

B.2 Computational Details

The borrower's optimization problem is characterized by four state variables (liquid wealth $X_t = A_t + Y_t$, 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. 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.

B.3 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 0.4 and 0.16 respectively. The results of these experiments are compared to the benchmark results in figure 10.

Figure 10: Sensitivity to Preference Parameter β

(a) $\beta = 0.85$ (b) $\beta = 0.9$ (c) $\beta = 0.95$

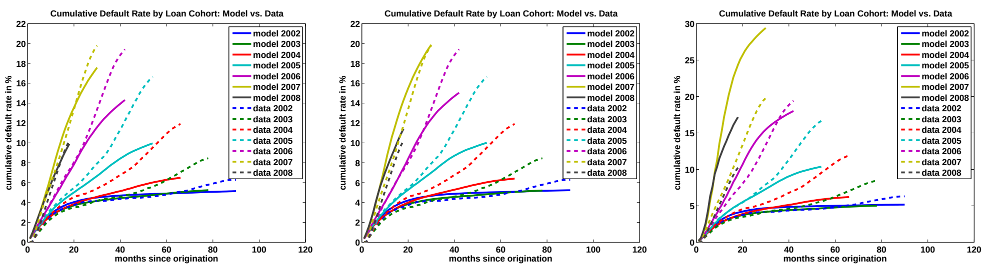

Next I consider alternative values of γ equal to 2 and 8. θ is then estimated as 0.06 and 0.64 respectively. Figure 11 compares these alternative calibrations to the benchmark.

Figure 11: Sensitivity to Preference Parameter γ

(a) $\gamma = 2$ (b) $\gamma = 5$ (c) $\gamma = 8$

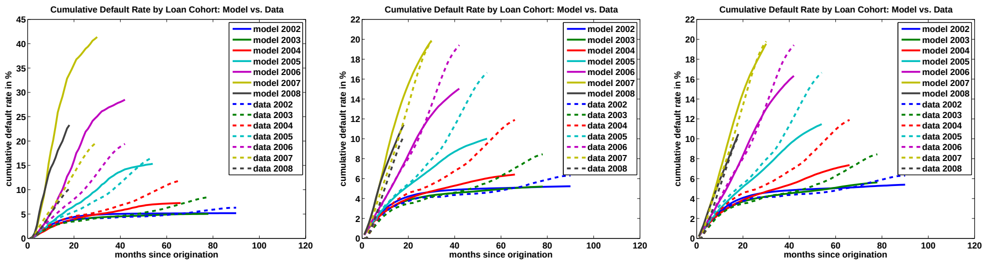

These results show that the model works as well or better than in the benchmark calibration for higher values of γ and/or lower values of β. These parameter changes make the agent less willing to substitute intertemporally and/or more impatient to consume today. This worsens the liquidity problem caused 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.

B.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. I simply change the inflation rate π ad-hoc to 1% instead of 2.4% in the benchmark calibration. All other parameters are unchanged, but θ is reestimated at a value of 0.44 to fit the 2002 cohort. Figure 12 presents these results. The fit of the model improves in the later period after origination relative to the benchmark results.

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 abso-

Figure 12: Performance of the Model for a Lower Inflation Rate

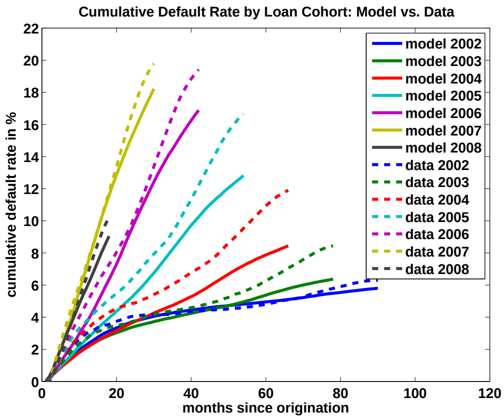

lute costs of both policies tend to increase a bit relative to the benchmark and the relative cost of the bailout to lenders also increases. Bailing out lenders is then between 8.6 and 11.2 times more expensive than subsidizing homeowners using this alternative calibration. Thus the conclusions across policies are robust or even strengthened relative to the benchmark calibration.

B.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 A.3. All other procedures are as in the main text. The estimate of θ is then 0.32. Figure 13 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 8.8 and 11.1 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 13: Model Results for the 120+ Days Default Definition