

Online Appendix of  
“Mortgage Default during the U.S. Mortgage Crisis”

Thomas Schelkle\*

July 6, 2018

This online appendix of the paper, “Mortgage Default during the U.S. Mortgage Crisis”, presents further details and results on the data and empirical facts (appendix A), the reduced form models (appendix B) and the structural model (appendix C).

---


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

# APPENDIX A: DATA AND EMPIRICAL FACTS

The main analysis of the paper is focussed on loans with an initial loan-to-value (LTV) ratio above 95% for the substantive reasons explained in section 1.1 of the paper. This appendix shows that the general empirical patterns concerning the rise in default rates and relatively stable loan characteristics at origination across cohorts extend to other loans with lower LTVs as well. In other words the full sample exhibits broadly the same dynamics across cohorts than the sample of highly leveraged borrowers analysed in the paper, though there are of course various level differences.

Figure A1 presents the observed cumulative default rates for different loan cohorts for the full sample of loans with all possible initial LTVs. Of course the data still refers to prime, fixed-rate, 30-years mortgages as in the main text. Though the level of default rates is in general somewhat lower than for the sample of highly leveraged borrowers, the rise in default rates follows a very similar pattern.

Figure A1: Cumulative Default Rates in the Full Sample



Evidence on average loan characteristics at origination for different cohorts for the full sample are presented in Table A1. The observed pattern is similar to the one for highly leveraged borrowers in the main text. Average FICO credit scores are almost constant across cohorts. The average loan-to-value ratio changes a bit more across cohorts than for the above 95% LTV sample. But surprisingly in the full sample many of the later

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

Table A1: Average Loan Characteristics at Origination by Loan Cohort for the Full Sample

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

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

At least for the prime mortgage market the facts documented in this section are prima facie evidence against explanations that emphasize the role of a deterioration in lending standards or increased borrower leverage for the rise in default rates like Corbae and Quintin (2015). For the full sample I have also conducted an additional analysis of compositional changes with respect to initial LTVs. For this purpose I grouped all loans with respect to their initial LTV into 10 bins with a width of 10 percentage points (the two exceptions are the first bin with all LTVs between 1 – 14% and the last bin with all LTVs above 95%). Using the observed default rates of each of these bins for each cohort and the actual LTV composition of each cohort I then computed two counterfactual aggregate default rates for all cohorts, i.e. the equivalent to Figure A1. In the first counterfactual scenario I keep the default rates of each LTV bin fixed at their observed level of the

2002 cohort and only vary the LTV composition across cohorts according to the actual data. In the second counterfactual I only vary the default rates of each bin across cohorts as observed in the actual data, but counterfactually keep the LTV composition fixed as it is observed for the 2002 cohort. This decomposition analysis shows that for all practical purposes the aggregate default dynamics across cohorts witnessed in Figure A1 are entirely driven by increases in default profiles for each initial LTV bin and not by changes in the composition of initial LTVs. The same conclusion applies if one uses a cohort other than the 2002 one for fixing one of the two margins in the counterfactual calculations. These results reinforce the conclusion in the main paper that the fall in house prices and not compositional effects are key for understanding the rise in default rates at least for prime fixed-rate mortgages during the crisis.

## APPENDIX B: REDUCED FORM MODELS

### *B.1 Estimation Procedure*

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

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

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

## ***B.2 Robustness Checks***

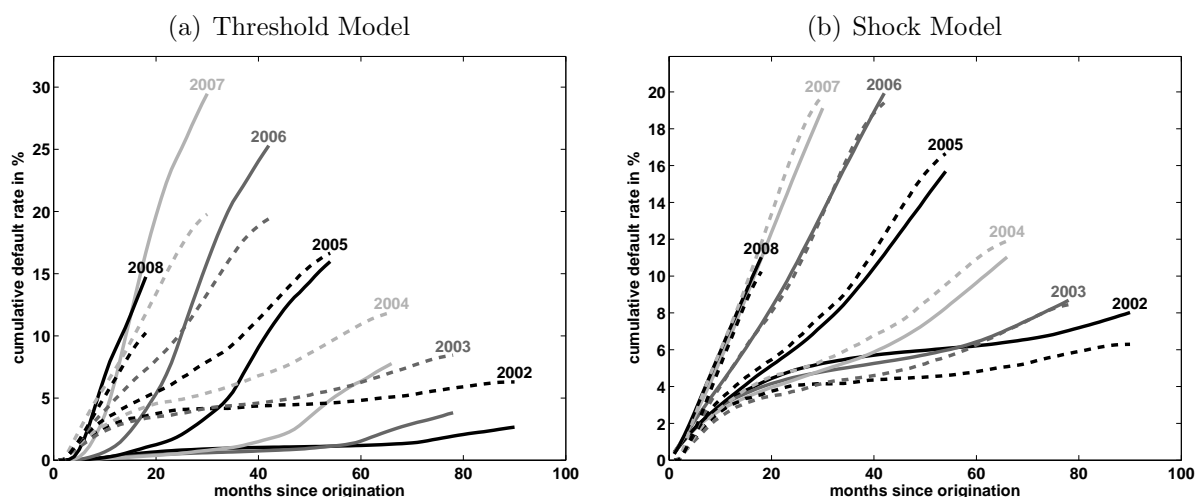
This section reports a battery of robustness checks that were performed to scrutinize the reduced form results. I find that the results are robust across all the modifications considered here. For brevity I do not report the graphs corresponding to Figure 3 for all these checks. But these are available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

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

Another robustness check replaces the out-of-sample test with an in-sample test. Here, I estimate the two models on all cohorts and then examine the fit within that sample. This exercise is informative on the best possible fit both models can give to the data. These results are thus worth to report in more detail. The estimated parameters are then  $-15.9\%$  for  $\phi$  and  $1.36\%$  for  $\psi$ , and the results are shown in Figure B1. The threshold model still has considerable problems to match the data even under these most favorable circumstances. It generally undershoots earlier cohorts and the early months after origination for all cohorts and at the same time still overshoots the late months of the 2006 and 2007 cohorts. In contrast, the shock model gives a very good fit to the data. The conclusions across models are essentially unchanged.

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

Figure B1: In-Sample Fit of the Two Models



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

In section 2.4 of the paper, it was assumed that the individual house price shocks are normally distributed. The major argument supporting this choice is that by the central limit theorem the mean of individual shocks converges asymptotically to a normal distribution anyway. However, since the analysis also covers periods where  $t$  is still small, I perform an additional check here. Instead of using a normal distribution for the individual shocks I specify them as being uniformly distributed on the interval  $[-b_t, b_t]$ . The parameter  $b_t$  is then chosen such that the variance of the uniformly distributed shock in period  $t$  in the respective census division is identical to the one used in the standard framework. I find that the results are almost identical.

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

option model. There is also no strong reason why the default threshold parameter  $\phi$  and default shock probability  $\psi$  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  $\phi_t$  and  $\psi_t$  that are allowed to differ each period from  $t = 1, \dots, T$  when fitting the models to the 2002 cohort. Under these circumstances, both models almost perfectly match the 2002 cohort. The cumulative default rates simulated for the other cohorts then inherit the non-smoothness of the first differences of the cumulative default rate of the 2002 cohort. However, subject to that qualification the conclusions on the out-of-sample fit remain essentially unchanged. The threshold model still greatly overshoots the later cohorts. The shock model generates default rates comparable in magnitude to the benchmark.

### ***B.3 Using an alternative definition of default***

In this section, I use a different definition of default. Instead of using a 60 or more days definition of default as in the main text, I now consider a loan to be in default once it is at least 120 days past due.<sup>1</sup> This is a more demanding definition and by all accounts being 120 days past due is considered as a very serious delinquency. This change of definition affects the levels of the data on cumulative default rates which is used to estimate and test the models. The level of default rates is lower now. However, the broad dynamics across cohorts are similar to the ones analyzed in the main text.

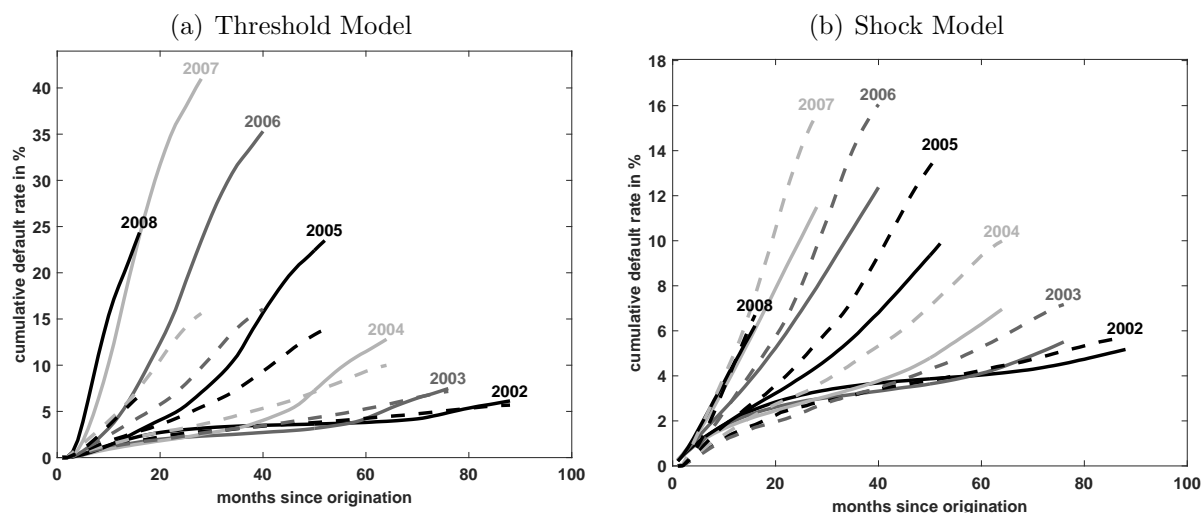
Again, I estimate both reduced form models on the 2002 cohort and use the remaining cohorts to test the estimated models. For the threshold model this yields an estimate of  $\phi$  of  $-11.9\%$  and for the shock model  $\psi$  is estimated as  $0.83\%$ . The results are reported in Figure B2. These 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

---

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

Figure B2: Using an alternative definition of default (120+ days)



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

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 2 of the paper.

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

The paper is focussed on loans with an LTV above 95% because these borrowers should be least likely to have a second mortgage on their home, cf. the discussion in section 1.1 of the paper. The question arises whether the results of the paper also generalize to loans with a lower LTV. This section provides some evidence on this by repeating the reduced-form analysis of section 2 for loans with an LTV of the first mortgage between 75% and 84%. Due to the discussed data problems this section is necessarily somewhat tentative. Nevertheless, some very interesting results emerge.

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



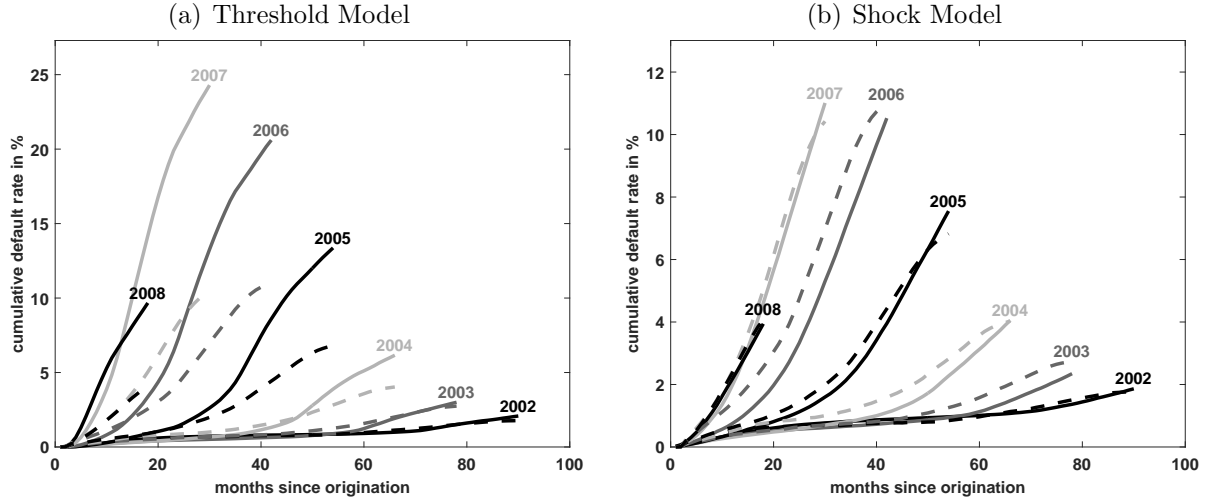
for the most extreme parameter values of  $\phi = 0$  and  $\psi = 1$ , both models undershoot the cumulative default rate of the 2002 cohort substantially for at least the first 60 months after origination. The reason is that the equity buffer generated by the down-payment is substantial for these borrowers.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

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

For this setup the reduced-form models are estimated again on the 2002 cohort. This yields estimates of  $\phi = -7.8\%$  and  $\psi = 2.25\%$ . The estimated models are again tested on their ability to predict out-of-sample. Figure B3 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. However, these results are at least suggestive that the main conclusions on the relative merit of the two theories may well extend to loans with a lower LTV.

Figure B3: Results for borrowers with a first mortgage LTV of 75 – 84% taking second mortgages into account



## APPENDIX C: STRUCTURAL MODEL

### C.1 Value Functions

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

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

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

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

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

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

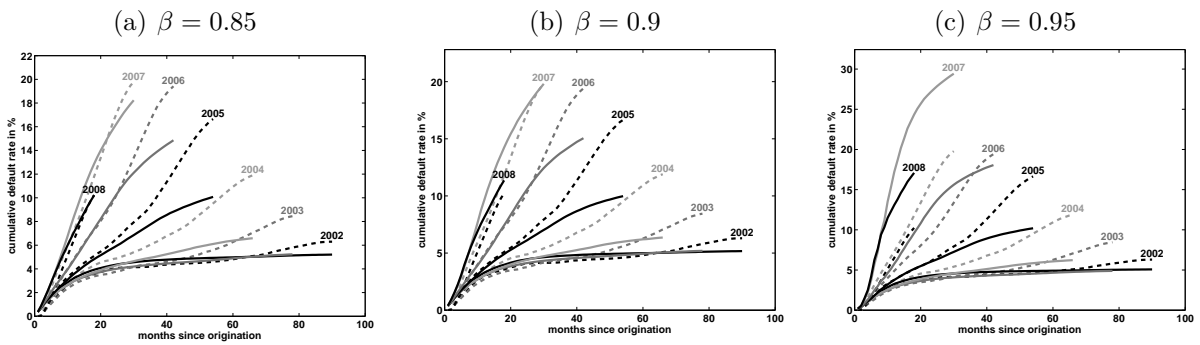
$$\begin{aligned}
V^r(X_t, L_t, t) &= \max_{C_t} \left\{ \frac{C_t^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}_t [V^r(X_{t+1}, L_{t+1}, t+1)] \right\} \\
\text{s.t. } X_{t+1} &= (1+r)(X_t - R - C_t) + Y_{t+1} \\
C_t &\leq X_t - R.
\end{aligned}$$

## C.2 Dependence on Preference Parameters

This section explores how the model depends on the predetermined preference parameters. Specifically, I compute results for alternative parameter values for  $\beta$  and  $\gamma$  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  $\beta$  equal to 0.85 and 0.95. For these value of  $\beta$ , the parameter  $\theta$  is then reestimated in order to fit the 2002 cohort. This yields values of  $\theta$  of 0.39 and 0.16 respectively. The results of these experiments are compared to the benchmark results in Figure C1.

Figure C1: Sensitivity to Preference Parameter  $\beta$

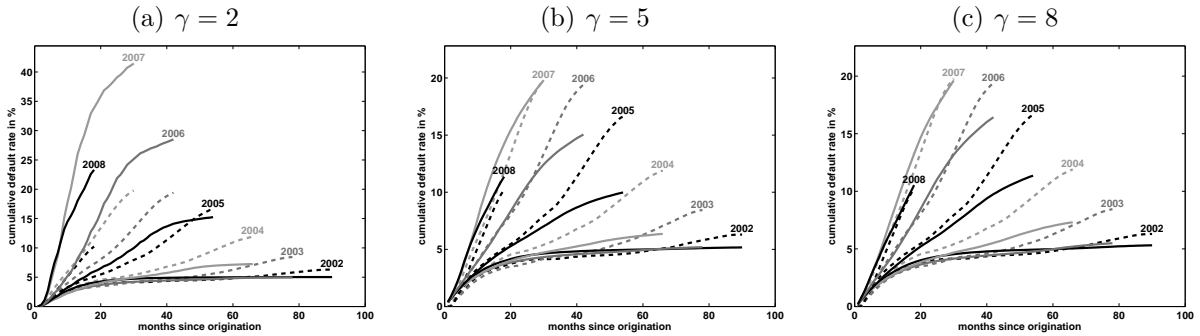


Notes: Solid lines: model. Dashed lines: data.

Next, I consider alternative values of  $\gamma$  equal to 2 and 8.  $\theta$  is then estimated as 0.06 and

0.64 respectively. Figure C2 compares these alternative calibrations to the benchmark.

Figure C2: Sensitivity to Preference Parameter  $\gamma$



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


These results show that the model works as well or better than in the benchmark calibration for higher values of  $\gamma$  and/or lower values of  $\beta$ . 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  $\gamma$  and higher values of  $\beta$  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 2 of the paper.

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

In this section, I show that versions of the model which abstract from a direct utility benefit of living in the bought house are excessively sensitive to changes in aggregate house prices. As an illustration I conduct simulations of the model where the utility flow parameter  $\theta$  is set to 0. I then adjust the discount factor  $\beta$  to again match the default rates of the 2002 cohort for  $\theta = 0$ . Figure C3 presents these results for the benchmark value of the CRRA coefficient  $\gamma = 5$ , but also for values of  $\gamma$  of 2 and 8. Each of those experiments is associated with a different value of  $\beta$ . These results confirm the key role of including a direct utility benefit of living in the bought house in order to prevent employed and not currently liquidity-constrained agents from defaulting when house prices decline

strongly. Without this feature the model cannot accurately capture the sensitivity of default rates to aggregate house prices.

Figure C3: Performance of the model without a direct utility benefit of living in the bought house ( $\theta = 0$ ) for different values of  $\gamma$



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


### C.4 Role of Inflation

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

Figure C4 presents the results of this exercise. The fit of the model improves in the later period after origination relative to the benchmark results. This shows that the strength of the mortgage tilt effect is responsible for the problems of the model in the main text to capture default decisions in later periods after origination.

Using this alternative calibration I have also repeated the policy analysis of section 6

Figure C4: Performance of the Model for a Lower Inflation Rate




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

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

### *C.5 Using an alternative definition of default*

This section reports the results for the structural model when the 120 days definition instead of the 60 days definition is used to measure default empirically as in section B.3. All other procedures are as in the main text. The estimate of  $\theta$  is then 0.32. Figure C5 reports the results for all cohorts. The fit of the model to this data is qualitatively similar to the one of the main text that uses a 60 days default definition. The results of the policy analysis are also essentially unchanged and bailing out lenders is found to be between 8.7 and 11.2 times more expensive than subsidizing homeowners. This analysis shows that the main results of the structural model are robust to using an alternative reasonable definition of default.

Figure C5: Model Results for the 120+ Days Default Definition



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

### C.6 A Structural Single-Trigger Model

This section investigates a structural single-trigger model in the spirit of the frictionless option theoretic literature. The aim is to confirm that the conclusions on the relative merit of the two theories drawn in the reduced form section 2 also carry over to a comparison of structural models of these theories.


For this purpose, it is convenient that the structural double trigger model of sections 3 and 4 in fact nests a single-trigger model if one suitably modifies some model parameters. The single-trigger paradigm relies on the assumption that borrowers have access to perfect and complete financial markets such that they can perfectly insure against income fluctuations. Accordingly, mortgage default decisions are unaffected by income fluctuations and liquidity problems. The structural model nests such a case when the following changes to model parameters are made. The net replacement rate in case of unemployment  $\rho$  is set to 1 such that income is identical in the employed and unemployed state. Given the actual replacement rate of unemployment insurance and the unemployment rate implied by the job separation and finding rates from the benchmark calibration, this insurance costs about 2% of net income when employed. Accordingly, employed net income  $(1 - \tau)Y_0$  is lowered by 2% such that it is net of insurance fees, which is achieved by resetting the debt to income ratio to a slightly higher value of  $DTI = 0.4085$ . These

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

Figure C6 presents the resulting fit of the structural single-trigger model to all cohorts. The model strongly overpredicts the rise in default rates. Its performance is broadly comparable to the one of the reduced-form single-trigger model. When comparing the structural single and double-trigger models to each other, one again finds that the double-trigger model explains the data much better than the single trigger model. The impression from visually inspecting Figure C6 and Figure 7 of the main text is confirmed by simple measures of goodness of fit of the model predictions to the data. For instance the root mean squared error (mean absolute error) of the out-of-sample forecast is 3.5 (2.8) times higher for the structural single-trigger than for the structural double-trigger model. Thus, this exercise using structural models confirms the reduced form results on the relative merit of the two theories.



Figure C6: Performance of a Structural Single-Trigger Model



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

## REFERENCES

- Corbae, Dean, and Erwan Quintin (2015). Leverage and the Foreclosure Crisis. *Journal of Political Economy* 123(1), 1–65.
- Elul, Ronel, Nicholas S. Souleles, Souphala Chomsisengphet, Dennis Glennon, and Robert Hunt (2010). What ‘Triggers’ Mortgage Default? *American Economic Review Papers and Proceedings* 100(2), 490–494.