Correlation in Mortgage Defaults

Chao Ma*　　　　　Hongbiao Zhao†

Xiamen University　　*Shanghai University of Finance and Economics*

December 11, 2018

*Wang Yanan Institute for Studies in Economics & Department of Finance at School of Economics, Xiamen University, Xiamen, Fujian 361005, China. Email: ma.271@buckeyemail.osu.edu

†School of Statistics and Management, Shanghai University of Finance and Economics, No. 777 Guoding Road, Shanghai 200433, China. Email: h.zhao1@lse.ac.uk

Abstract

Previous studies have made numerous efforts in estimating the default risk for individual mortgages. However, the risk of mortgage-backed securities or banks' mortgage portfolios also depends on the correlation among individual mortgages. We conduct formal statistical tests and find that conditional on the typical observable factors affecting individual mortgage default risk, there is still a large degree of correlation in defaults that would generate time clustering in defaults. We further conduct a variety of robustness checks and find that this residual correlation cannot be explained by the missing observable macroeconomic variables (e.g. national GDP growth or stock market returns) or the unobserved location-specific time-invariant frailty. To quantify the degree of this residual correlation, we calibrate a residual Gaussian copula correlation parameter after conditioning on the individual default intensities estimated from the Cox hazard model.

Keywords: Mortgage default; Default correlation; Copula; Mortgage-backed securities

1 Introduction

Mortgage default risk is one of the most critical risks facing the financial system. The 2007 financial crisis was triggered by the meltdown of the mortgage market. Accurately analyzing mortgage default risk and correctly pricing the risk is important for households, mortgage lenders, and mortgage-backed security (MBS) investors. A rich literature has studied the econometric techniques in estimating individual mortgage default risk and the empirical factors that affect individual mortgage default risk (e.g. Deng et al. (2000), Clapp et al. (2006), and Guiso et al. (2013)). When lenders decide whether to approve a borrower's mortgage application and how much interest to charge, estimating individual-level mortgage default risk might provide sufficient information.

However, when evaluating the risk of MBSs or banks' mortgage loan portfolios, what matters is not only the default risk of each individual loan, but also the correlation among them. The probability density function of MBS returns will have a heavier left tail if there is default correlation among the mortgages in the pool. Even with accurate estimates on the default risk at the individual level, ignoring correlation in defaults will result in underestimating the MBS risk and overestimating the risk-diversification effect of pooling mortgages together and issuing MBSs, and thus cause overpricing on MBSs. Ignoring correlation in mortgage defaults will also cause banks to underestimate the tail risk of banks' mortgage loan portfolios and the level of economic capital required by the Basel Accords to withstand extremely adverse scenarios. In addition, ignoring correlation in mortgage defaults can cause other financial institutions to make poor risk-management decisions, such as Freddie Mac and Fannie Mae's decisions on mortgage-purchase activities and guarantee-fee charges and fund managers' risk-hedging strategies. Correlation in mortgage defaults can also enlarge the overall risk of the entire financial system, which should be taken into consider by the regulator when making policies.

There are multiple reasons that can cause correlation among defaults of different individual mortgages, especially the time clustering in defaults. First, risk-related observable covariates may be correlated across different loans. For example, local house prices and unemployment rates may be correlated across different locations. Second, unobserved time-variant common risk factors can cause defaults to cluster in time. It is inevitable that some relevant risk factors cannot be observed by researchers or investors but play important roles in affecting default intensities, such as soft information (Rajan et al., 2015), lenders' screening efforts (Keys et al., 2010, 2012), borrowers and lenders' strategic behavior in debt renegotiation (Piskorski et al., 2010; Mayer et al., 2014; Agarwal et al., 2017), and emotional and behavioral factors (Bhutta et al., 2017). Those unobserved factors can have co-movements across different individual mortgages.

The third reason is spillover effects or contagion through multiple channels. One channel is that the foreclosure of a property will reduce the value of other properties in the same neighborhood (Schuetz et al., 2008; Lin et al., 2009; Harding et al., 2009; Campbell et al., 2011; Li, 2017) and thus make them more likely to default. Another channel is that a high number of defaults in a location or the entire society will reduce the social stigma of strategic default and convey information to borrowers about the probability of being sued, which causes people to strategically choose to default at a lower threshold (Guiso et al., 2013).[1]

Theoretically, risk analyses based on traditional econometric models for estimating individual default risks (such as Cox proportional hazard models) can capture the default correlation coming from some sources (e.g. the comovements of observable covariates), but cannot capture the default correlation coming from other sources (e.g. unobserved time-variant common risk factors and spillover or contagion effects) if those sources exist.

In this study, we test whether the comovement of observable covariates of different mortgages is sufficient to account for the degree of time clustering in defaults that we observe in the mortgage performance data. We find that conditional on the typical observable factors that are frequently employed by the mortgage default literature to estimate individual default risk, there is still a large degree of correlation in defaults that would generate time clustering in defaults. To quantify the degree of this extra correlation, we calibrate

[1] We focus on correlation that accounts for time clustering in defaults. There are other factors that can generate default correlation but not default clustering in time. Kau et al. (2011) included local unobserved time-invariant frailties as random effects in the estimation of individual mortgage default risk. Deng et al. (2005) included individual unobserved time-invariant frailties that are spatially correlated with each other. As an MBS pool usually includes mortgages throughout the entire country, the risk caused by those unobserved time-invariant frailties can be efficiently diversified.

a residual Gaussian copula correlation parameter. We find that the value of this parameter is comparable to the degree of residual correlation in corporate defaults that has been documented in the corporate default literature.

The methodologies in this paper mainly follow those employed by Das et al. (2007) in testing correlation in corporate defaults.[2] A rich literature has examined correlation in corporate defaults. As a seminal paper in this area, Das et al. (2007) tested whether there is extra correlation accounting for time clustering in corporate defaults that are not captured by the traditional econometric models estimating individual default intensities with typical observable risk factors included. We employed a Cox proportional hazard model to estimate individual mortgage default intensities, which has greater flexibility than the econometric model used for individual corporate defaults in Das et al. (2007) and other studies on corporate defaults. The additional flexibility comes from the baseline hazard function in the Cox model, which is a function of time since loan origination to be estimated nonparametrically. The reason to add this flexibility is that unlike firms, which are supposed to exist forever unless they go bankrupt or get acquired, mortgages have a finite horizon, such as 30 years. Even if all the other risk factors have the same values, mortgages at different distance to maturity (or length of time since origination) should have different default intensities.

In the corporate default literature, given the positive results of the extra correlation tests in the seminal paper Das et al. (2007), many other papers structurally modeled the sources of this extra correlation and estimate their corporate default models. Duffie et al. (2009) and Nickerson and Griffin (2017) estimated corporate default models with a common time-variant unobserved risk factor that would generate time clustering in defaults. Chen and Wu (2014) extend models with a single common unobserved risk factor to models with

[2]Chernobai et al. (2011) employed similar methods to test correlation in bank operational risks.

multiple sectoral frailties. Azizpour et al. (2018) estimated a corporate default model with contagion effects (a default by one firm has a direct impact on the health of other firms) using Hawkes point processes. In addition, some papers built theoretical models to illustrate the corporate default correlation coming from certain potential sources, such as Phelan (2017) on the efficiency for financial intermediaries to manage similar investments together. Furthermore, many studies discussed or examined the impact of correlation in corporate defaults on the risk of portfolios or structured financial products of corporate bonds (see Giesecke (2004); Das et al. (2007); Yu (2007); Coval et al. (2009); Duffie et al. (2009); Driessen et al. (2009); Koopman et al. (2009); Buraschi et al. (2010); Lando and Nielsen (2010); Azizpour et al. (2011); Nickerson and Griffin (2017); Phelan (2017)). Therefore, the results on the extra correlation of mortgage defaults in our study would motivate a number of future studies to be conducted on mortgage defaults.[3]

Although correlation in mortgage defaults plays an important role in shaping the risk of MBSs, banks' mortgage portfolios, and the entire financial system, few studies have empirically investigated this correlation. There are a few theoretical papers modeling correlation in mortgage defaults. Liu et al. (2009) employed copula to model the co-movements among individual property values, which can generate correlation in mortgage defaults. Fan et al. (2012) employed counterparty risk to model the mortgage default dependence structure. Both papers calibrate the effect of correlation in mortgage defaults on the risk and price of MBSs. Our study provides empirical evidence of an extra correlation in mortgage defaults that cannot be captured by the comovements of observables and quantifies this extra correlation by calibrating a copula model.

There are also studies analyzing correlation in other asset markets and its implication

[3]Our study in this paper focuses on testing the correlation in mortgage defaults that cannot be explained by observables and on quantifying this extra correlation by the calibrations of a copula model. Explicitly and structurally modeling the sources of this extra correlation (such as common time-variant unobserved risk factors and contagion effects) and estimating the model are beyond the scope of this paper.

on the systematic risk, such as Christoffersen et al. (2012) on correlation among international equity markets, Patton (2006) on correlation among foreign exchange rates, DeFusco et al. (2013) on contagion among real estate markets in different MSAs, and Case et al. (2012) on correlation between Real Estate Investment Trust (REIT) and stock returns.

The remaining portion of this paper is organized as follows. In Section 2, we describe the data used in this paper. In Section 3, we estimate individual mortgage default intensities uisng a Cox proportional hazard model. In Section 4, we conduct formal statistic tests for the extra correlation and quantify the extra correlation by calibrating a copula model. Section 6 describes robustness checks. Section 7 provides further discussion. Then, we conclude in Section 8.

2 Data

Freddie Mac Single Family Loan-Level Dataset The full dataset covers all the 30-year fixed-rate mortgages originated during 1999–2014 and purchased and guaranteed by Freddie Mac, which includes approximately 17 million loans. The data include loan-level origination information and monthly loan performance information. The origination information includes the FICO credit score at origination, the original loan-to-value ratio (LTV), the original debt-to-income ratio (DTI), the original unpaid balance, the metropolitan statistical area (MSA), the 3-digit zip code, and the state. The monthly loan performance information includes the current loan delinquency status and the loan age. Freddie Mac also created a smaller dataset by randomly selecting $50,000$ loans from each origination year. The analyses in this paper are based on this random sample.

Table 1: Descriptive statistics for the data

Variable	Mean	SD
FICO/100	7.3243	0.5391
Debt/income	0.3419	0.1161
Current LTV	0.6747	0.1962
MSA Unemployment rate	0.0671	0.0234
MSA Per capita income growth	0.0028	0.0322
National GDP growth	0.0408	0.0294
Monthly S&P500 return	0.0034	0.0434

Federal Housing Finance Agency (FHFA) House Price Index The FHFA publishes 3-digit zip code level house price indices. We match the 3-digit zip code level house price index to the Freddie Mac loan-level dataset by the 3-digit zip code to calculate the current house values. The purchase price or appraisal value of a house at the mortgage origination can be derived from the Freddie Mac loan-level dataset. Following the literature, we divide the house value at origination by the 3-digit zip code level house price index at origination and multiply it by the current house price index to construct the current house value. The current LTV is calculated as the current unpaid balance divided by the current house value.

Other Data We also extract MSA-level unemployment rates from the Bureau of Labor Statistics (BLS), MSA-level per capita income from the Bureau of Economic Analysis (BEA), and national GDP growth and S&P500 returns from the Federal Reserve.

Following the literature, default is defined as more than 90 days past due for a mortgage. At the end of the sample horizon, the cumulated default rate of the loans in the sample is 4.15%. Table 1 reports the descriptive statistics for key variables.

3 Estimating Individual Mortgage Default Intensity

In this section, we estimate a Cox proportional hazard model to obtain the default intensity of each individual mortgage. As a typical class of econometric methods dealing with duration data, Cox proportional hazard models have been wildly applied in analysing mortgage defaults (e.g. Deng et al. (2000)), corporate defaults (e.g. Doshi et al. (2013)), and bank operational risks (e.g. Chernobai et al. (2011)).[4]

The default intensity of mortgage i at calendar time t, $\lambda_{i,t}(\tau)$, is given by

$$\lambda_{i,t}(\tau) = \lambda_0(\tau) \exp\left(x_{i,t}\beta\right), \tag{1}$$

where τ is the length of time since origination, $\lambda_0(\tau)$ is the baseline hazard rate, and $x_{i,t}$ is the vector of time-variant and time-invariant covariates that proportionally shift the baseline hazard. $x_{i,t}$ include both mortgage-specific risk factors (FICO credit score, DTI, and LTV) and local economic conditions (unemployment rate and per capita income growth). The corporate default intensity model estimated by Das et al. (2007) did not have the baseline hazard function $\lambda_0(\tau)$ that varies in τ. The reason is that firms are supposed to exist forever unless they go bankrupt or get acquired. However, mortgages have a finite horizon, such as 30 years. Even if all the other risk factors have the same values, mortgages at different distance to maturity (or length of time since origination) should have different default intensities. Therefore, the baseline hazard rate (a function of time since origination) is included in the mortgage default model in the literature (e.g. Deng et al. (2000), Clapp et al. (2006), An et al. (2010), and Kau et al. (2011)).

The Cox proportional hazard model is estimated by a partial likelihood method, which has two steps. In the first step, β is estimated parametrically; in the second step, $\lambda_0(\tau)$ is estimated nonparametrically (see Cox (1975) for detailed estimation procedures). Table 2 displays the estimation results for β. All the parameters are consistent with theories, intuition, and previous empirical results reported in the mortgage default literature. A

[4] Shumway (2001) discussed the advantages of hazard models compared to single-period static models in predicting failures and proved the equivalence between hazard models and multi-period logit models.

Table 2: Estimation of a Cox proportional hazard model for individual mortgage default intensities

Variable	Estimate	SE
FICO/100	-1.1628***	0.0160
Debt/income	2.2205***	0.0811
Current LTV	2.9650***	0.0391
Unemployment rate	8.5599***	0.4323
Per capita income growth	-5.7537***	0.2569

*** denotes the 0.01% level of significance

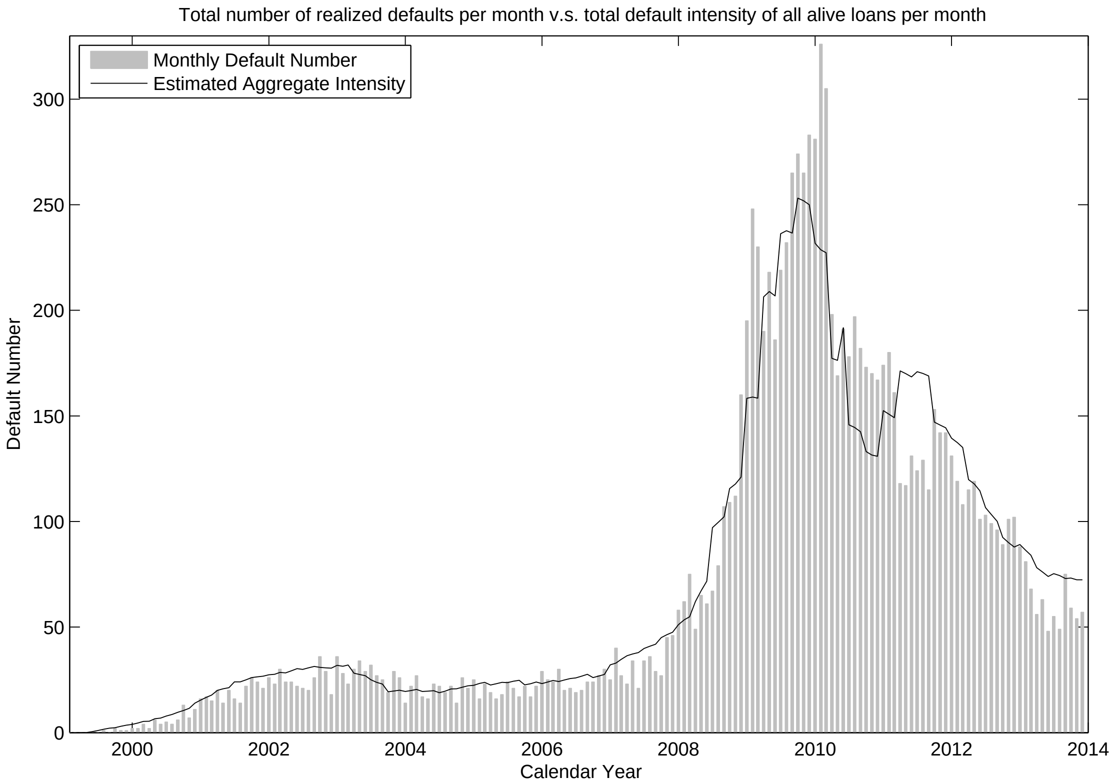

Figure 1: Aggregate intensities and defaults by month.

mortgage will have a higher default intensity if the credit score is low, the DTI or current LTV is high, the local unemployment rate is high, or the local per capita income growth is low. For each month, we aggregate the individual mortgage default intensities across all the mortgages that are active in that month. As shown in Figure 1, the monthly aggregate default intensity well matches the total number of defaults in each month.

4 Testing Extra Correlation

In this section, we test whether there is extra default correlation that is not captured by the empirical model in Section 3. Traditional methods of hypothesis testing require that under the null hypothesis, there should be multiple observations drawn from an identical distribution. However, in the data-generating process of mortgage defaults, different mortgages have different default intensities and the same mortgage at different months has different default intensities, which bring complications into the testing of default correlation.

We follow the time-rescaling method developed by Das et al. (2007) to address this issue. The basic idea of the method is as follows: first, the default intensities of all the active mortgages in a month are aggregated to obtain a monthly aggregated default intensity; second, time bins with different lengths are constructed such that the aggregated default intensity in each bin is identical across different bins. Thus, under the null hypothesis that there is no extra correlation of defaults that is not captured by the empirical model in Section 3, the realized total number of defaults in each time bin follows a standard Poisson distribution with a constant event rate. Section 4.1 below discusses the procedure of rescaling calendar time into time bins; Section 4.2 discusses the analyses based on the realized total number of defaults in each time bin.

4.1 Time Rescaling for Poisson Defaults

We suppose that the default time of the i^{th} loan, τ_i^*, is the first jump time of a point process with stochastic intensity process $\{\lambda_{i,t}\}_{t\in[0,T]}$ for any *loan index* $i = 1, 2, ..., n$, where n is the total number of loans in the sample. The total number of defaults in the sample up to time t is

$$N_t = \sum_{i=1}^{n} \mathbf{1}\{\tau_i^* \leq t\}, \qquad t \in [0, T],$$

where T is the terminal time. The total default intensity of all surviving loans at time t in the sample is

$$\lambda_t := \sum_{i=1}^{n} \lambda_{i,t} \mathbf{1}\{\tau_i^* > t\}, \qquad t \in [0, T].$$

For each loan i, $\lambda_{i,t}$ up to time τ_i^* have been estimated according to equation (1). The aggregated intensity λ_t and the total number of realized defaults for each month are plotted in Figure 1.

Denote the *total accumulative default intensity* of all surviving loans in the pool by

$$\Lambda(t) := \int_0^t \lambda_s \mathrm{d}s, \qquad t \in [0, T].$$

To construct time bins such that each time bin contains the same amount of aggregated intensity, we cut the calendar time at the time points $\{t_0, t_1, ..., t_K\}$ with $t_0 = 0, t_K \leq T$ such that

$$\Lambda(t_k) - \Lambda(t_{k-1}) = c, \quad \text{for any bin index } k = 1, ..., K.$$

The constant c is referred as the bin size, i.e., the units of aggregated intensity in each bin. $K \in \mathbb{N}^+$ is the total number of bins cut from the sampling period (1999-2014). Denote the number of defaults within the k^{th} time bin, $[t_{k-1}, t_k)$, by

$$X_k := \sum_{i=1}^{n} \mathbf{1}\big\{t_{k-1} \leq \tau_i < t_k\big\}, \qquad k = 1, ..., K.$$

In theory (Meyer, 1971), under the null hypothesis that each default follows an independent process represented by the Cox proportional hazard model in equation (1), $\mathbf{X} \equiv \{X_k\}_{k=1,2,...,K}$ would be *i.i.d.* and should follow a standard Poisson distribution with a constant event rate c. We explore multiple approaches to test this null hypothesis and they all show the violation.

To ensure that our results are robust to the change in the bin size, we set $c = 200, 300, 400$, and 500, respectively, and conduct the analyses correspondingly. In each panel of Figure 2, vertical lines represent the time bin delimiters; each time bin has a different time length, but the aggregated intensity or the change of total accumulative intensity within each bin is always equal to c. The bin is narrower (wider) when the aggregated intensity is high (low) or the total accumulative intensity is steeper (flatter).

We do not set c beyond the 200-500 range for the following consideration. On the one hand, if c is too small, some bins would be narrower than one month, but mortgage payments are usually monthly; on the other hand, if c is too large, the number of bins cut from the sampling period (1999-2014) would be too small to guarantee the power of our statistical tests. In the corporate default literature (Das et al., 2007; Lando and Nielsen, 2010), much smaller numbers are used for c because the number of corporate bonds is much smaller than the number of mortgages, and the number of corporate defaults is much smaller than the number of mortgage defaults. However, the average time length of bins in our study is similar to that in Das et al. (2007) and Lando and Nielsen (2010).

4.2 Analyses based on Realized Numbers of Defaults in Rescaled Time Bins

4.2.1 Comparing Theoretical Distributions and Moments with Their Empirical Counterparts

For a given bin size c, we first compare the empirical distribution and moments of the realized number of defaults in each bin (X_k) with the simulated distribution and moments from a standard Poisson distribution with a constant event rate c. The empirical distributions

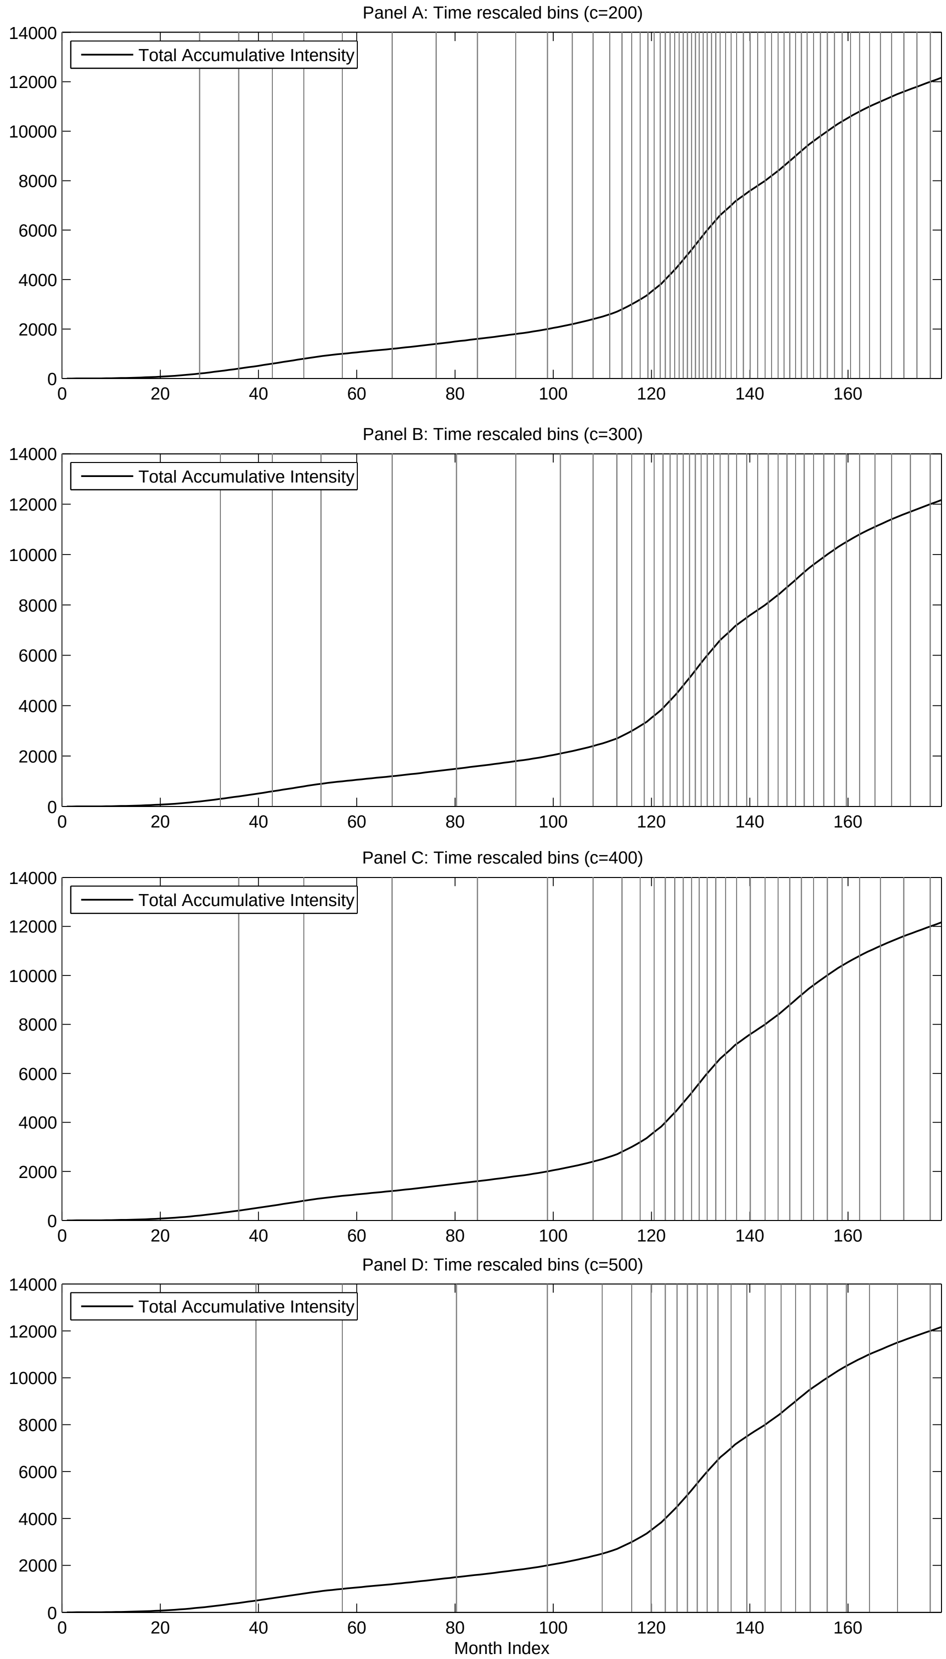

Figure 2: Time rescaled bins. The time bin delimiters at $\{t_k\}_{k=1,2,\ldots}$ over 1999-2014 (178 months) are marked by vertical lines. The curve represents the total accumulative default intensity. Panels A, B, C, and D display the time rescaled bins for the bin size $c = 200, 300, 400$ and 500, respectively. Given c, each bin is constructed such that the change of the total accumulative default intensity in each bin (or the aggregated default intensity within each bin) is always equal to c.

Table 3: Comparison of empirical (data) and theoretical (Poisson) moments
This table presents a comparison of empirical and theoretical moments for the distribution of defaults per bin. c is the bin size; K is the corresponding number of bin observations. The theoretical moments are either analytically derived or computationally simulated from the theoretical Poisson distributions under the hypothesis that there is no residual correlation after conditioning on the Cox hazard model in equation (1).

	c	K	Mean	Variance	Skewness	Kurtosis
Theoretical	200	60	200.00	200.00	0.07	3.00
Empirical	200	60	200.60	1,434.82	0.56	2.73
Theoretical	300	40	300.00	300.00	0.06	3.00
Empirical	300	40	300.93	2,997.76	0.52	2.39
Theoretical	400	30	400.00	400.00	0.05	3.00
Empirical	400	30	401.23	5,003.91	0.44	2.13
Theoretical	500	24	500.00	500.00	0.04	3.00
Empirical	500	24	501.54	7,118.78	0.34	2.46

of X_k for bin size $c = 200, 300, 400$, and 500 and the associated simulated Poisson distributions (with $10,000$ replications for each distribution) are plotted in Figure 3. Across all the choices for bin size c, the empirical distribution is more dispersed and has heavier tails than the simulated Poisson distribution. This pattern indicates that there exists extra correlation of defaults that is not captured by the default intensity model in equation (1).

The moments of a standard Poisson distribution of event rate c can be calculated analytically, and its mean, variance, skewness and kurtosis are c, c, $1/\sqrt{c}$ and $1/c + 3$, respectively. Then, we can carry out a comparison between these theoretical results and their empirical counterparts. As shown in Table 3, for each bin size c, the first-order moments (means) (marked by dashed lines "- - -") are very closely matched, which indicates that the default intensity model in equation (1) does not have systematic biases in estimating individual mortgage default intensities. However, the high order moments (variance, skewness, and kurtosis) are not well matched. Especially, the empirical variance is much larger than the theoretical variance, which indicates that there exists extra correlation not captured by the model equation (1), and this extra correlation generates higher uncertainty of aggregate-level default risks than the model in equation (1) does.

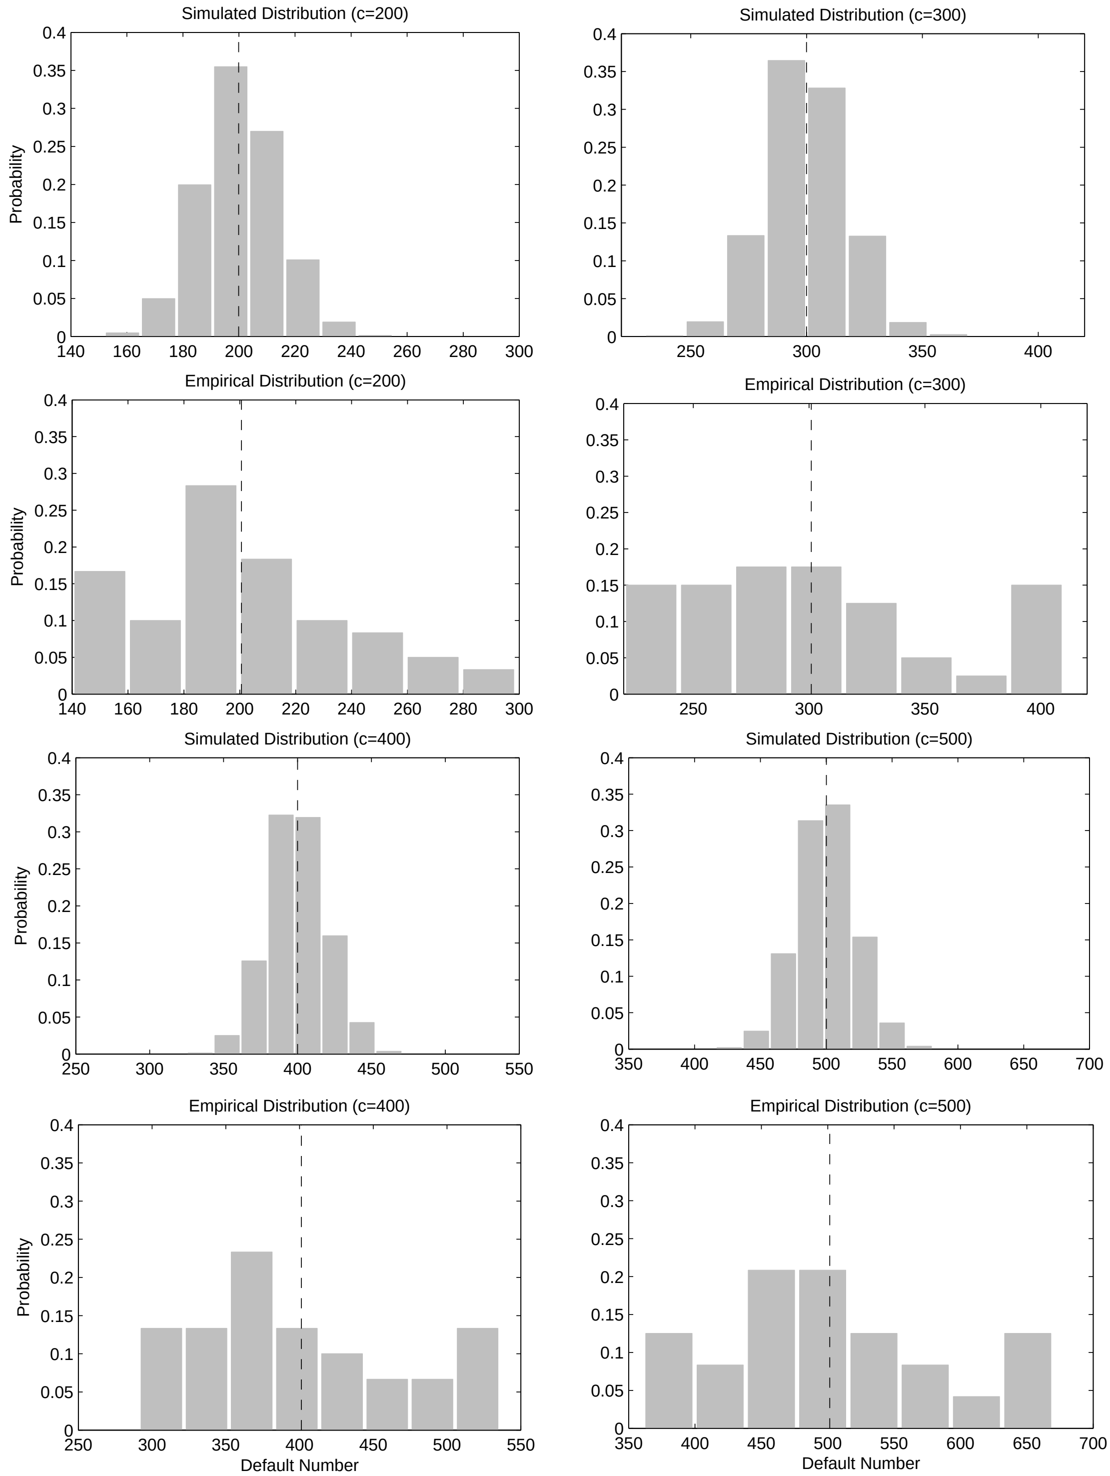

Figure 3: Comparison between the empirical distributions and the theoretical Poisson distributions of defaults per bin. For the bin size $c = 200$, 300, 400, and 500, respectively, the theoretical distribution is simulated with 10,000 replications from the Poison distribution under the hypothesis that there is no residual correlation after conditioning on the Cox hazard model in equation (1). Their means are marked by vertical dashed lines "- - -". The empirical distribution is more dispersed than its theoretical counterparts.

4.2.2 Fisher's Dispersion Test and Upper-Tail Test

Under the null hypothesis that each default follows an independent process represented by the Cox proportional hazard model in equation (1), $\mathbf{X} \equiv \{X_k\}_{k=1,2,\ldots,K}$ would be *i.i.d.* and should follow a standard Poisson distribution with a constant event rate c. To formally test this hypothesis, we first carry out the Fisher's dispersion test (Cochran, 1954) for a given bin size c with K bins. Under the null hypothesis that **X** are *i.i.d.* and follow a Poisson distribution of rate c, the W statistic defined in equation

$$W := \sum_{k=1}^{K} \frac{(X_k - c)^2}{c} \tag{2}$$

follows a chi-squared distribution of $K - 1$ degrees of freedom, i.e. $W \sim \chi^2_{K-1}$. The result of the Fisher's dispersion test is reported in the left part of Table 4. Across all the choices of bin size c, the null hypothesis is rejected at a 0.01% level of significance.

Alternatively, we also conduct the upper-tail test. For a given bin size c with K bins, denote M as the sample mean of the upper quartile of the empirical distribution of **X**. We simulate $10,000$ data sets, each consisting of K *i.i.d.* random variables drawn from a Poisson distribution with event rate c. Then, for each simulated data set, we compute the upper-quartile mean. The p-value is estimated as the fraction of the simulated data sets for which the sample upper-quartile mean is high than the empirical upper-quartile mean M.

The right part of Table 4 displays the result of the upper-tail test. The empirical upper-quartile mean is significantly higher than the simulated upper-quartile mean from a Poisson distribution at a level of 0.01%. This result indicates that the actual default risks have a fatter upper-quartile tail than predicted by the Cox model of (1).

Table 4: Fisher's dispersion test and upper-tail (UT) test

The columns on the left present Fisher's dispersion test. Under the null hypothesis that there is no residual correlation after conditioning on the Cox hazard model in equation (1), given the bin size c, the defaults per bin X_k are *i.i.d.* and follow a Poisson distribution of rate c, and thus the W statistic defined in equation (3) follows a chi-squared distribution of $K-1$ degrees of freedom. The columns on the right present the upper-quartile test. Under the null hypothesis that there is no residual correlation after conditioning on the Cox hazard model in equation (1), the simulated upper-quartile (UQ) mean of X_k from the theoretical Poisson distribution of rate c should well match its empirical counterpart.

	Fisher's Test		*UT Test*		
c	W	p-value	Empirical UQ mean	Simulated UQ mean	p-value
200	423.38	0.0000	252.80	209.42	0.0000
300	389.82	0.0000	378.10	311.51	0.0000
400	362.90	0.0000	505.57	413.27	0.0000
500	327.58	0.0000	613.33	514.72	0.0000

5 Measuring Extra Default Correlation via Copula

5.1 One-Correlation-Parameter Model

Our previous tests provide strong evidences that there exists extra correlation of mortgage defaults that is not captured by the traditional Cox hazard model in equation (1). The next step is to gauge the degree to which the correlation is not captured by the traditional Cox hazard model. Similar to Das et al. (2007), we calibrate the intensity-conditional copula model to obtain the pairwise residual copula correlation that must be added, after conditioning on the default intensities estimated by equation (1), to match the upper-quartile means of the empirical distribution of defaults per time bin.

Similar to the corporate default literature, we use the upper-quartile mean as the benchmark to gauge the copula correlation parameter. The reason is that the upper-quartile mean is a measure of tail risk widely used in both the academia and the industry, and the tail risk of MBS returns or banks' loan portfolios generated by the correlation among individual mortgages have important implications for MBS pricing and financial institutions' risk-management decisions. Theoretically, different copula models would generate different numeric values for the residual copula correlation. In this paper, we employ the industry standard "Gaussian copula", which is widely used in the pricing of structured credit products such as collateralized debt obligations (CDOs). The calibrating algorithm is provided in Appendix A.

Table 5 reports the calibration results. Given a bin size c, for each gridded numeric value of the pairwise residual copula correlation parameter ρ, we simulate the total number of defaults for a bin and then compute the upper-quartile mean of the simulated sample. The numbers in bold represent the upper-quartile means of the simulated defaults per time bin that are closest to their empirical counterparts. Across all the choices for bin size c, the copula correlation parameter ρ at which the simulated upper-quartile mean best approximates its empirical counterpart is always 0.5%. The residual copula correlation calibrated by Das et al. (2007) for corporate defaults ranged from 1% to 4%, depending on the choice of the bin size; Nickerson and Griffin (2017) obtained a similar range for corporate defaults. The residual copula correlation calibrated in this paper for mortgage defaults is lower than those for corporate defaults. This is consistent with intuition because the number of mortgages is much larger than the number of corporate bonds, and mortgage borrowers have greater heterogeneities and are more loosely connected than firms.

However, a 0.5% pairwise residual copula correlation is large enough to dramatically increase the upper-tail risk of pool-level default risks. As shown in Table 5, ignoring the residual correlation (letting $\rho = 0$) would generate a upper-quartile mean that is lower than its empirical counterpart by $10-15\%$ ($(252.8-226.17)/252.8-1 = 10.53\%$ for $c = 200$; $(378.1 - 328.79)/378.1 - 1 = 13.04\%$ for $c = 300$; $(505.57 - 430.74)/505.57 - 1 = 14.80\%$ for $c = 400$; and $(613.33 - 531.96)/613.33 - 1 = 13.27\%$ for $c = 500$). Figure 4 displays the relationship between the upper-quartile mean and the residual Gaussian copula correlation for each bin size c. Given c, the upper-quartile mean is increasing in the residual Gaussian copula.

Table 5: Residual Gaussian copula correlation

Given a bin size c, for each gridded numeric value of the pairwise residual copula correlation parameter ρ, we simulate the total number of defaults for a bin and then compute the upper-quartile mean of the simulated sample. The numbers in bold represent the upper-quartile means of the simulated defaults per time bin that are closest to their empirical counterparts (the second column). The last column $\hat{\rho}$ represents the correlation parameter from linear interpolations.

		Pairwise residual copula correlation parameter									
c	Data	0.00%	0.25%	0.50%	0.75%	1.00%	1.25%	1.50%	1.75%	2.00%	$\hat{\rho}$
200	252.80	226.17	245.12	**258.21**	273.18	279.42	292.52	299.28	307.01	312.68	0.40%
300	378.10	328.79	357.57	**378.09**	394.97	409.75	422.52	434.61	448.87	450.93	0.50%
400	505.57	430.74	470.83	**496.86**	517.65	535.41	555.59	568.29	580.62	595.11	0.60%
500	613.33	531.96	582.16	**611.99**	637.11	656.36	681.62	701.31	719.58	732.94	0.51%

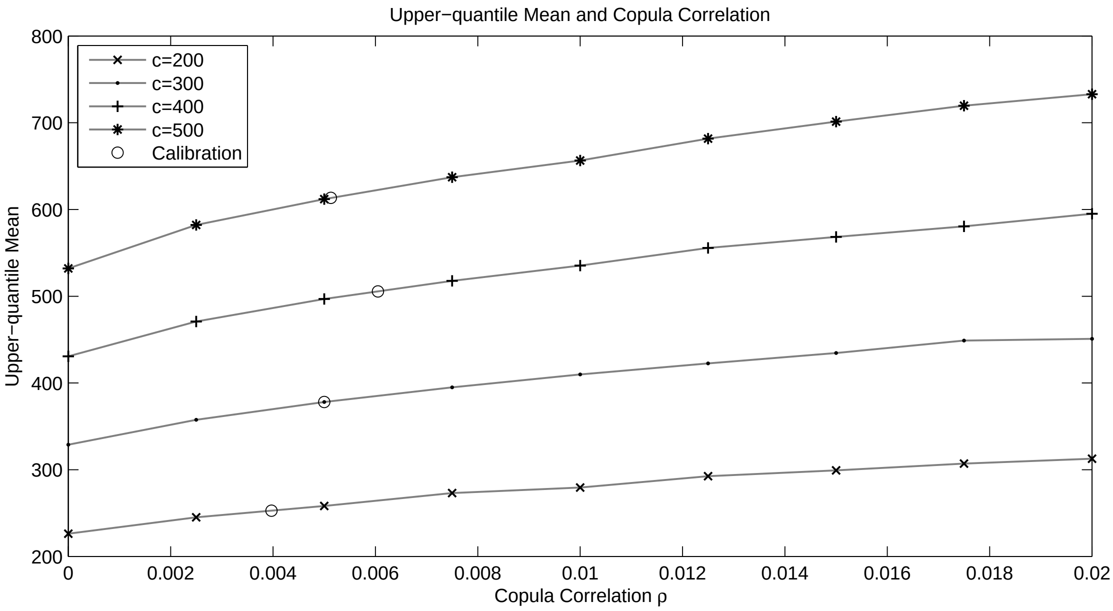

Figure 4: Relationship between residual Gaussian copula correlation and upper-quartile mean. Given a bin size c, the upper-quartile mean of the simulated defaults per bin (X_k) is increasing in the residual Gaussian copula correlation. The circle marks represent the empirical upper-quartile means.

5.2 Two-Correlation-Parameter Model

While any two mortgages can potentially be correlated in defaults, two mortgages in the same location may have a higher correlation. This fact is different from corporate defaults. The reason why any two mortgages can potentially be correlated in defaults is that some unobserved common risk factors are at the national level and that some spillover effects are transmitted through social network or social media rather than geographically (e.g., a high number of defaults in the entire society will reduce the social stigma of strategic default and convey information to borrowers about the probability of being sued, which causes people to strategically choose to default at a lower threshold (Guiso et al., 2013). The reason why two mortgages in the same location may have a higher correlation is that some other unobserved common risk factors are at the local level and that some other spillover effects are transmitted geographically.

In this subsection, we calibrate a two-correlation-parameter model to capture the special property of mortgage markets, which is different from the one-correlation-parameter model developed by Das et al. (2007) for corporate defaults. We assume that in a time bin, the correlation between any two mortgages across different locations is ρ_1 and the correlation between any two mortgages in the same location is $\rho_1 + \rho_2$. Therefore, ρ_1 captures the correlation between any two mortgages in the country; and ρ_2 captures the additional correlation within the same location. We calibrate ρ_1 and ρ_2 by matching the upper quartile mean of total default numbers and a measure of geographic dispersion of location-level default rates in the simulated data to their empirical counterparts. The higher ρ_2, the more likely that defaults are clustered in certain locations and thus the higher geographic dispersion of location-level default rates; the higher ρ_1 and ρ_2, the higher upper quartile mean of total default numbers.

The measure of geographic dispersion of location-level default rates in a time bin k is defined as follows:

$$\mathsf{dispersion}_k = \sum_{m=1}^{M} \Big(\text{default rate in location } m - \text{average default rate across locations}\Big)^2, \tag{3}$$

where

$$\text{default rate in location } m = \frac{\text{number of defaults in location } m}{\text{number of existing loans in location } m}.$$

For a given pair of copula correlation parameters ρ_1 and ρ_2 and a bin size c, we use the following procedure to simulate defaults:

1. Suppose the number of loans in bin k is n_k. For each loan index i and each bin index k, calculate C_i^k, which denotes the increment in the cumulative intensity for this loan within this bin, i.e.,

$$C_i^k := \int_{t_{k-1}}^{t_k} \lambda_{i,s} \mathbf{1}\{s < \tau_i^*\} \mathrm{d}s, \qquad i = 1, 2, ..., n_k, \qquad k = 1, 2, ..., K.$$

2. Equally likely draw one of the bins, say k. Draw $\{W_i\}_{i=1,2,...,n_k}$ from a joint standard normal distribution with the following correlation structure: for any two loans i and j, if they are in different locations, $\mathrm{corr}(W_i, W_j) = \rho_1$; if they are in the same location, $\mathrm{corr}(W_i, W_j) = \rho_1 + \rho_2$.[5]

3. For each loan i, let $U_i = \Phi(W_i)$, where $\Phi(.)$ is the standard normal cumulative

[5]We use a modified two-factor Gaussian copula model to generate $\{W_i\}_{i=1,2,...,n_k}$. Denote $\jmath(i) \in \{1, 2, ..., M\}$ as the *location index* of loan i, where M is the total number of locations. Randomly draw Z, $\{Z_m\}_{m=1,2,...,M}$, and $\{\varepsilon_i\}_{i=1,2,...,n_k}$ from the standard univariate normal distribution $\mathsf{N}(0, 1)$. Compute

$$W_i = \sqrt{\rho_1} Z + \sqrt{\rho_2} Z_{\jmath(i)} + \sqrt{1 - \rho_1 - \rho_2}\, \varepsilon_i, \qquad i = 1, ..., n_k.$$

distribution function. Loan i will default in bin k if $U_i > e^{-C_i^k}$ for $i = 1, 2, ..., n_k$, and will be active otherwise. Then, we obtain the total number of simulated defaults in a bin for the mortgage pool and calculate the geographic dispersion of location-level default rates.

4. Repeat Steps 2 and 3 for $10,000$ times independently, and then calculate the upper-quartile mean of simulated defaults per bin and the average geographic dispersion of location-level default rates across these $10,000$ scenarios.

(The calibration results of ρ_1 and ρ_2 for different bin sizes are coming soon)

6 Robustness Checks

The results in the previous section indicate that there is a significant amount of default correlation that cannot be explained by the co-movement of observable risk factors. Therefore, the traditional risk model may perform well in predicting default risks for individual mortgages, but will underestimate the risk of banks' mortgage portfolios or MBSs based on mortgage pools. As one of the robustness checks, we examine whether the extra correlation comes from the nonlinear effects of observables on defaults. We try a variety of specifications for equation (1), including adding quadratic and cubic terms of observables or dividing the observables into bins and using bin fixed effects. Similar results on correlation tests and calibrations are obtained.[6] As another robustness check, we examine whether the extra correlation comes from other individual characteristics missing from the baseline regression in Table 2, including occupancy status (whether the property is owner occupied, second home, or investment property), whether originated through a broker, loan purpose (for house purchase or for refinance), whether first-time homebuyer, and

[6] The results are available upon request.

whether with a second mortgage. We add those variables into the estimation of equation (1) and similar results on correlation tests and calibrations are obtained.[7]

In this section, we examine whether the residual correlation can be explained by two other channels. One channel is the macroeconomic variables observable to researchers but missing in the model specifications. The other channel is the unobserved location-specific time-invariant frailty shared by mortgages originated from the same location.

6.1 Testing for Missing Covariates

It is possible that national level macroeconomic conditions can affect the default risks of mortgages all over the country and thus generate default correlation. One key variable measuring macroeconomic conditions is the GDP growth rate. In the previous model specification, we have already included local unemployment rate and per capita income growth, which are highly correlated with the national GDP growth rate. Here, we test whether the residual correlation of defaults can be caused by the missing national GDP growth, after controlling for local economic conditions. Following Das et al. (2007), given a bin size, we use the number of defaults in a bin in excess of the mean $(X_k - c)$ as the dependent variable and run regressions on the national GDP growth rates. As shown in Table 6, for all the bin sizes, GDP growth is not significant. As noted in Das et al. (2007), if an additional variable does not contribute to predicting the default risk after controlling for the existing variables, this additional variable should be uncorrelated with the number of defaults in a bin in excess of the mean.

Another macro variable we examine is S&P500 returns. First, stock market performance is also an important indicator of the macro economy. Second, losing money from the stock market may make households more likely default on their mortgages. As shown

[7]The results are available upon request.

Table 6: Testing for missing covariates

For each bin size c, we run regressions of the number of defaults in excess of the mean, $Y_k = X_k - c$, on the previous quarter's GDP growth rate and the previous month's return on S&P 500 index. The number of observations in each regression is the number of bins of size c. *** denotes the 0.01% level of significance. Coefficients without * are insignificant at a level of 10%.

c	Intercept	GDP growth (%)	S&P500 return (%)
200	208.8355***	-2.9731	
	(6.2988)	(1.4198)	
300	311.7991***	-3.9292	
	(11.2779)	(2.8202)	
400	411.2938***	-3.6407	
	(16.1073)	(3.7028)	
500	512.6674***	-3.7929	
	(23.5721)	(5.3663)	
200	201.1104***		-0.9895
	(5.088)		(1.4855)
300	301.2067***		-0.8309
	(8.9122)		(2.4823)
400	401.6607***		-2.5373
	(13.1075)		(3.9133)
500	502.4382***		-3.4222
	(17.4235)		(5.0756)
200	208.8102***	-2.9513	-0.0682
	(6.3425)	(1.4446)	(1.2718)
300	311.9551***	-4.019	0.2733
	(11.4894)	(2.8958)	(2.1957)
400	410.2706***	-3.1844	-1.4107
	(15.9464)	(3.73)	(3.7425)
500	510.4558***	-2.8508	-2.106
	(23.693)	(5.3606)	(4.7942)

in Table 6, S&P500 returns are insignificant in the regressions of $(X_k - c)$ for all the bin sizes.

We also add GDP growth and S&P 500 returns in the estimation of default intensity and repeat Fisher's dispersion test, the upper-tail test, and the calibration of copula correlation. The results are similar to those based on the default intensity estimation without including GDP growth and S&P 500 returns.[8]

6.2 Incorporating Location-specific Time-invariant Frailty

Mortgages originated in the same local area may share common unobserved factors that affect default intensities. This could be one source of mortgage default correlation. Kau

[8]The results are available upon request.

et al. (2011) estimated a Cox proportional hazard model with shared MSA-specific time-invariant frailty for mortgage defaults. In this subsection, we estimate a similar model and obtain the default intensities of each mortgage; then, we repeat Fisher's dispersion test, the upper-tail test, and the calibration of the copula correlation. The results are very similar to those based on the Cox model without shared frailty, which indicates that there is still a significant amount of default correlation coming from sources other than shared local time-invariant frailty.

Suppose the default intensity is determined by equation

$$\lambda_{i,t}(\tau) = \lambda_0(\tau) \exp\left(x_{i,t}\beta\right) v_m, \tag{4}$$

where v_m is the unobserved time-invariant risk factor shared by all the mortgages originated in location m, and v_m is different across different locations. Following the typical literature employing Cox models with shared frailty, v_m is assumed to follow a Gamma distribution $\Gamma\left(\frac{1}{\theta}, \theta\right)$ with $\theta > 0$. The parameter θ measures the heterogeneity of the unobserved time-invariant frailty across different locations. Because the variance of $\Gamma\left(\frac{1}{\theta}, \theta\right)$ is θ ($\frac{1}{\theta} \times \theta^2 = \theta$), this model will degenerate to a Cox model without shared frailty as θ approaches to zero.

The model is estimated with the expectation-maximization (EM) algorithm; see Dempster et al. (1977) and Therneau and Grambsch (2000, p.253–255) for detailed estimation procedures. The algorithm not only generates estimates of $\lambda_0(\tau)$, β, and θ, but also generates estimates of v_m for all the locations. Table 7 displays the estimates of β and θ, where a location is defined as a 3-digit zip code area.[9] All the β parameters are consistent with theories and intuition. The estimate of θ is 0.0613 and the standard error is 0.0104. Based on

[9]We also estimate a shared-frailty model in which a location is defined as an MSA. The corresponding results for Fisher's dispersion test, the upper-tail test, and the calibration of the copula correlation are similar.

a likelihood ratio test, θ is significantly positive at the level of 0.1%, which indicates that mortgages originated in the same location do share common unobserved time-invariant risk factors that affect default intensities. This result is consistent with the result in Kau et al. (2011).

The estimated default intensity of each mortgage at calendar time t is calculated by

$$\widehat{\lambda}_{i,t}(\tau) = \widehat{\lambda}_0(\tau) \exp\left(x_{i,t}\widehat{\beta}\right) \widehat{v}_m.$$

We aggregate $\widehat{\lambda}_{i,t}(\tau)$ of all the active mortgages in our sample at calendar time t to obtain the aggregated default intensity for calendar time t. Based on the aggregated default intensity for calender time t, we rescale the time and construct bins to repeat Fisher's dispersion test, the upper-tail test, and the calibration of the copula correlation.

As shown in Table 8, for each bin size c, the first-order moments (means) (marked by dashed lines "- - -") are very closely matched. However, the high order moments (variance, skewness, and kurtosis) are not well matched. Especially, the empirical variance is much larger than the theoretical variance. These results are similar to those in Table 3 for the Cox model without frailties. Table 9 displays the results of Fisher's dispersion test and the upper-tail test. The hypothesis that the default correlation is fully captured by the co-movement of observables $x_{i,t}$ and the existence of shared location-specific time-invariant unobserved frailty is rejected again at a very high significant level. There are only small changes in the statistics of those tests after adding shared frailty to the model, compared to Table 4. As shown in Table 10, compared to Table 5, the calibrated residual copula correlations also do not change much after adding shared frailty to the model.

Table 7: Estimation of a proportional hazard model with location-specific time-invariant frailty

Variable	Estimate	SE
FICO/100	-1.1628***	0.0160
Debt/income	2.2205***	0.0811
Current LTV	2.9650***	0.0391
Unemployment rate	8.5599***	0.4323
Per capita income growth	-5.7537***	0.2569
θ	0.0613**	0.0104

** denotes the 0.1% level of significance. *** denotes the 0.01% level of significance.

Table 8: Comparison of empirical (data) and theoretical (Poisson) moments (frailty case)
This table presents a comparison of the empirical and theoretical moments for the distribution of defaults per bin. c is the bin size; K is the corresponding number of bin observations. The theoretical moments are either analytically derived or computationally simulated from the theoretical Poisson distributions under the hypothesis that there is no residual correlation after conditioning on the proportional hazard model with location-specific time-invariant frailty according to equation (4).

	c	K	Mean	Variance	Skewness	Kurtosis
Theoretical	200	60	200.00	200.00	0.07	3.00
Empirical	200	60	200.53	1,292.29	0.30	2.92
Theoretical	300	40	300.00	300.00	0.06	3.00
Empirical	300	40	300.82	2,750.66	0.32	2.90
Theoretical	400	30	400.00	400.00	0.05	3.00
Empirical	400	30	401.10	4,496.44	0.39	3.15
Theoretical	500	24	500.00	500.00	0.04	3.00
Empirical	500	24	501.38	6,709.11	0.45	3.43

Table 9: Fisher's dispersion test and upper-tail (UT) test (frailty case)
The columns on the left present Fisher's dispersion test. Under the null hypothesis that there is no residual correlation after conditioning on the proportional hazard model with location-specific time-invariant frailty according to equation (4), given the bin size c, the defaults per bin X_k are *i.i.d.* and follow a Poisson distribution of rate c, and thus the W statistic defined in equation (3) follows a chi-squared distribution of $K-1$ degrees of freedom. The columns on the right present the upper-quartile test. Under the null hypothesis that there is no residual correlation after conditioning on the proportional hazard model with location-specific time-invariant frailty according to equation (4), the simulated upper-quartile (UQ) mean of X_k from the theoretical Poisson distribution of rate c should well match its empirical counterpart.

	Fisher's Test		*UT Test*		
c	W	p-value	Empirical UQ mean	Simulated UQ mean	p-value
200	381.31	0.0000	248.07	209.39	0.0000
300	357.68	0.0000	369.60	311.54	0.0000
400	326.08	0.0000	494.43	413.31	0.0000
500	308.71	0.0000	606.67	514.82	0.0000

Table 10: Residual Gaussian copula correlation (frailty case)
Given a bin size c, for each gridded numeric value of the pairwise residual copula correlation parameter ρ, we simulate the total number of defaults for a bin and then compute the upper-quartile mean of the simulated sample. The numbers in bold represent the upper-quartile means of the simulated defaults per time bin that are closest to their empirical counterparts (the second column). The last column $\hat{\rho}$ represents the correlation parameter from linear interpolations

		Pairwise residual copula correlation parameter									
c	Data	0.00%	0.25%	0.50%	0.75%	1.00%	1.25%	1.50%	1.75%	2.00%	$\hat{\rho}$
200	248.07	229.40	**249.02**	263.25	273.91	285.48	292.03	303.59	311.86	319.00	0.24%
300	369.60	329.43	359.07	**377.45**	395.20	410.69	421.86	437.96	444.12	460.52	0.39%
400	494.43	432.07	473.95	**498.20**	519.83	537.12	550.74	575.23	585.69	596.66	0.46%
500	606.67	533.81	584.30	**614.56**	640.71	660.93	683.61	702.83	719.89	739.48	0.43%

7 Further Discussion

This paper provides empirical evidence of extra correlation in mortgage defaults that cannot be captured by the traditional Cox hazard models. This extra correlation does not come from the co-movements of observable risk factors and cannot be explained by either observable macroeconomic variables missing from the model specification or shared unobservable location-specific time-invariant frailty.

One concern is that with the technology of big data, more detailed information about borrowers can be collected, and fewer unobservables will be left in the residuals (for example, some studies used borrowersąŕ mobile-phone usage and social network information to predict loan defaults). Consequently, is it still worthwhile to calibrate the residual correlation of defaults?

First, although researchers can obtain many additional variables with the technology of big data, for each variable, usually there are a large proportion of observations with missing values. Those additional variables are helpful in predicting individual default risks for the sample with non-missing values. However, for the risk analysis at the portfolio or MBS level, all the individual mortgages in the pool matter. Calibrating the residual correlation is still helpful to improve the portfolio or MBS level risk analysis.

Second, those additional variables regarding borrowersąŕ information are usually costly to obtain and maintain. The costs include not only monetary costs but also time costs to go through the legal process because of privacy-protection laws. Practitioners need to make balance between obtaining these costly additional variables to improve the prediction of individual default risk and consuming more computation power in calibrating the residual correlation to improve the portfolio or MBS level risk analysis.

Third, to predict future mortgage default risks and the resulting portfolio or MBS level risks, the future movements of observables need to be predicted. Macroeconomic-level observables are much easier to predict than detailed-level observables. First, macroeconomic variables usually have longer historical data. Second, the performance of forecasting models for macroeconomic variables are more stable, robust, and reliable. Therefore, using a limited set of observables and taking into account extra residual correlations would be enough in predicting portfolio or MBS level risk. In the corporate default correlation literature, Das et al. (2007) and Duffie et al. (2009) used two firm-level observables and two macroeconomic observables; Nickerson and Griffin (2017) used one firm-level observables and six macroeconomic observables.

The results in this paper would motivate multiple future research topics regarding mortgage defaults that are beyond the scope of the current paper.

One future research topic is to estimate structural models to explore the sources and structure of this extra correlation, which is currently the frontier of studies on the extra correlation in corporate defaults. For example, Duffie et al. (2009) and Nickerson and Griffin (2017) estimated corporate default models with a common time-variant unobserved risk factor that could generate time clustering in defaults. Azizpour et al. (2018) estimated a corporate default model with contagion effects (the default by one firm has a direct impact on the health of other firms). Analyzing the extra correlation in mortgage defaults along those directions has the potential to provide deeper understanding of the sources and structure of correlation, better estimates for the degree of correlation, and more accurate predictions for pool-level default risks.

Another future research topic is to study how MBSs should be priced when taking into account the extra correlation or how much MBSs were mispriced previously for not considering the extra correlation. In the corporate default literature, Broer (2018) built a theoretical model to analyze the mis-pricing of CDOs due to correlation in corporate defaults. In the current industrial practice, the calibrated residual copula correlation has been applied in pricing corporate-debt financial products, such as CDOs, but not for mortgage related financial products (MBSs). The forecasting of pool-level mortgage default risks is usually conducted using the following procedure (see e.g. Duarte and McManus (2016)). In the first step, Cox hazard models are fitted for individual mortgage defaults. Second, time series models are fitted for the movement of observable risk factors (e.g. local house price index and unemployment rates). Third, given the current and past values of those observable risk factors, future values are simulated based on the time series models for a large number of paths. Fourth, the simulated future value of those observable risk factors are plugged into the Cox hazard model to predict the future default risk for each individual mortgage, and the pool-level future default risk is obtained by simply aggregating individual mortgage risks. Therefore, the risk forecasting procedure in current industrial practice only captures the mortgage default correlation generated from the co-movements of observable risk factors. The MBS pricing based on the risk analyses ignoring the extra residual correlation should be biased, which could play an important part in causing the 2007 financial crisis.

8 Conclusion

When evaluating the risk of MBSs or banks' mortgage loan portfolios, in addition to the default risk of each individual loan in the pool, the correlation among them should matter. Ignoring correlation among individual mortgage defaults will lead to underestimating the MBS risk and overestimating the risk-diversification effect of pooling mortgages together and issuing MBSs, and thus cause overpricing on MBSs. Ignoring correlation in mortgage defaults will also cause banks to underestimate the tail risk of banks' mortgage loan portfolios and the level of economic capital required by the Basel Accords to withstand extremely adverse scenarios. In addition, ignoring correlation in mortgage defaults can cause other financial institutions to make poor risk-management decisions, such as Freddie Mac and Fannie Mae's decisions on mortgage-purchasing activities and guarantee-fee charges and fund managers' risk-hedging strategies. Correlation in mortgage defaults can also enlarge the overall risk of the entire financial system, which should be taken into consider by the regulator when making policies.

There are multiple reasons that can cause correlation among defaults of individual mortgages, especially the time clustering in defaults. These reasons include comovements of observable covariates of different mortgages, unobserved time-variant common risk factors, and the spillover or contagion effects of defaults. Theoretically, risk analyses based on traditional econometric models for estimating individual default risks, such as Cox proportional hazard models, can capture the default correlation coming from the comovements of observable covariates, but they cannot capture the default correlation coming from other sources, if those sources exist.

In this study, we extend the methodologies used in the corporate default correlation literature to the residential mortgage market. We conduct formal statistical tests and find that conditional on the typical observable factors that are frequently employed by the mortgage default literature in estimating individual default risk, there is still a large degree of correlation in defaults that would generate time clustering in defaults. We further conduct a variety of robustness checks and find that this residual correlation cannot be explained by the missing observable macroeconomic variables (e.g. national GDP growth or stock market returns) and the shared unobservable location-specific time-invariant frailty. To quantify the degree of this residual correlation, we calibrate a residual Gaussian copula correlation parameter. The parameter value is comparable to the degree of residual correlation in corporate defaults that has been documented in the corporate default literature.

These results motivate multiple future research topics, such as estimating structural mortgage default models with a common time-variant unobserved risk factor or a contagion process, the methodologies of which have been newly developed in the corporate default literature.

Appendices

A Simulation Algorithm

According to Lando (1998), if each default follows a Cox process, then the default time τ_i^* of the i^{th} loan for any i can be defined by

$$\tau_i^* = \inf \left\{ t : \int_0^t \lambda_{i,s} \mathrm{d}s \geq E_i \right\},$$

where E_i is an exponential random variable of rate 1, i.e., $E_i \sim \mathsf{Exp}(1)$, and $\{E_i\}_{i=1,2,...,n}$ are *i.i.d.* It is well known that E_i can be simply constructed by

$$E_i = -\ln U_i,$$

where U_i is a standard continuous uniform random variable, i.e., $U_i \sim \mathsf{U}[0,1]$, and $\{U_i\}_{i=1,2,...,n}$ are *i.i.d.* Therefore, we can rewrite the default time τ_i^* above by

$$\tau_i^* = \inf\left\{t : \exp\left(-\int_0^t \lambda_{i,s}\mathrm{d}s\right) \leq U_i\right\}.$$

Conditional on the realization of intensity path $\{\lambda_{i,s}\}_{0\leq s\leq t}$ for any $i = 1,2,...,n$, all default times $\{\tau_i^*\}_{i=1,2,...,n}$ should be independent of each other under the Cox framework. However, an extra dependency for $\{\tau_i^*\}_{i=1,2,...,n}$ can be further constructed easily via a dependent series $\{U_i\}_{i=1,2,...,n}$ using a n-dimensional copula function, for example, the industrial standard Gaussian copula.

Therefore, following the algorithm developed by Schönbucher and Schubert (2001) and applied by Das et al. (2007), we use the procedure below to simulate the defaults for a given copula correlation ρ and a bin size c:

1. For each loan index i and each bin index k, calculate C_i^k, which denotes the increment in the cumulative intensity for this loan within this bin, i.e.,

$$C_i^k := \int_{t_{k-1}}^{t_k} \lambda_{i,s}\mathbf{1}\{s < \tau_i^*\}\mathrm{d}s, \qquad i = 1,2,...,n, \qquad k = 1,2,...,K.$$

2. Equally likely draw one of the bins, say k. Draw $\{W_i\}_{i=1,2,...,n}$ from a joint standard normal distribution with $\mathrm{corr}(W_i, W_j) = \rho$ whenever $i \neq j$.

3. For each loan i, let $U_i = \Phi(W_i)$, where $\Phi(.)$ is the standard normal cumulative distribution function. Loan i will default in bin k if $U_i > e^{-C_i^k}$ for $i = 1,2,...,n$, and will be active otherwise. Then, we obtain the total number of simulated defaults in a bin for the mortgage

pool.

4. Repeat Steps 2 and 3 for $10,000$ times independently, and then calculate the upper-quartile mean of simulated defaults per bin across these $10,000$ scenarios, which is reported in Table 5.

References

Agarwal, S., Amromin, G., Ben-David, I., Chomsisengphet, S., Piskorski, T., and Seru, A. (2017). Policy intervention in debt renegotiation: Evidence from the home affordable modification program. *Journal of Political Economy*, 125(3):654–712.

An, X., Clapp, J. M., and Deng, Y. (2010). Omitted mobility characteristics and property market dynamics: application to mortgage termination. *Journal of Real Estate Finance and Economics*, 41(3):245–271.

Azizpour, S., Giesecke, K., and Kim, B. (2011). Premia for correlated default risk. *Journal of Economic Dynamics and Control*, 35(8):1340–1357.

Azizpour, S., Giesecke, K., and Schwenkler, G. (2018). Exploring the sources of default clustering. *Journal of Financial Economics*, 129(1):154–183.

Bhutta, N., Dokko, J., and Shan, H. (2017). Consumer ruthlessness and mortgage default during the 2007 to 2009 housing bust. *Journal of Finance*, 72(6):2433–2466.

Broer, T. (2018). Securitization bubbles: Structured finance with disagreement about default risk. *Journal of Financial Economics*, 127(3):505–518.

Buraschi, A., Porchia, P., and Trojani, F. (2010). Correlation risk and optimal portfolio choice. *Journal of Finance*, 65(1):393–420.

Campbell, J. Y., Giglio, S., and Pathak, P. (2011). Forced sales and house prices. *American Economic Review*, 101(5):2108–31.

Case, B., Yang, Y., and Yildirim, Y. (2012). Dynamic correlations among asset classes: REIT and stock returns. *Journal of Real Estate Finance and Economics*, 44(3):298–318.

Chen, P. and Wu, C. (2014). Default prediction with dynamic sectoral and macroeconomic frailties. *Journal of Banking & Finance*, 40:211–226.

Chernobai, A., Jorion, P., and Yu, F. (2011). The determinants of operational risk in U.S. financial institutions. *Journal of Financial and Quantitative Analysis*, 46(6):1683–1725.

Christoffersen, P., Errunza, V., Jacobs, K., and Langlois, H. (2012). Is the potential for international diversification disappearing? a dynamic copula approach. *Review of Financial Studies*, 25(12):3711–3751.

Clapp, J. M., Deng, Y., and An, X. (2006). Unobserved heterogeneity in models of competing mortgage termination risks. *Real Estate Economics*, 34(2):243–273.

Cochran, W. G. (1954). Some methods for strengthening the common χ^2 tests. *Biometrics*, 10(4):417–451.

Coval, J., Jurek, J., and Stafford, E. (2009). The economics of structured finance. *Journal of Economic Perspectives*, 23(1):3–25.

Cox, D. R. (1975). Partial likelihood. *Biometrika*, 62(2):269–276.

Das, S. R., Duffie, D., Kapadia, N., and Saita, L. (2007). Common failings: How corporate defaults are correlated. *Journal of Finance*, 62(1):93–117.

DeFusco, A., Ding, W., Ferreira, F. V., and Gyourko, J. (2013). The role of contagion in the last American housing cycle. Wharton School Working Paper.

Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):1–38.

Deng, Y., Pavlov, A. D., and Yang, L. (2005). Spatial heterogeneity in mortgage terminations by refinance, sale and default. *Real Estate Economics*, 33(4):739–764.

Deng, Y., Quigley, J. M., and Order, R. (2000). Mortgage terminations, heterogeneity and the exercise of mortgage options. *Econometrica*, 68(2):275–307.

Doshi, H., Ericsson, J., Jacobs, K., and Turnbull, S. M. (2013). Pricing credit default swaps with observable covariates. *Review of Financial Studies*, 26(8):2049–2094.

Driessen, J., Maenhout, P. J., and Vilkov, G. (2009). The price of correlation risk: Evidence from equity options. *Journal of Finance*, 64(3):1377–1406.

Duarte, J. and McManus, D. A. (2016). Mortgage-related securities (MRSs). In *Handbook of Fixed-Income Securities*, pages 53–74. John Wiley & Sons.

Duffie, D., Eckner, A., Horel, G., and Saita, L. (2009). Frailty correlated default. *Journal of Finance*, 64(5):2089–2123.

Fan, G.-Z., Sing, T. F., and Ong, S. E. (2012). Default clustering risks in commercial mortgage-backed securities. *Journal of Real Estate Finance and Economics*, 45(1):110–127.

Giesecke, K. (2004). Correlated default with incomplete information. *Journal of Banking & Finance*, 28(7):1521–1545.

Guiso, L., Sapienza, P., and Zingales, L. (2013). The determinants of attitudes toward strategic default on mortgages. *Journal of Finance*, 68(4):1473–1515.

Harding, J. P., Rosenblatt, E., and Yao, V. W. (2009). The contagion effect of foreclosed properties. *Journal of Urban Economics*, 66(3):164–178.

Kau, J. B., Keenan, D. C., and Li, X. (2011). An analysis of mortgage termination risks: a shared frailty approach with MSA-level random effects. *Journal of Real Estate Finance and Economics*, 42(1):51–67.

Keys, B. J., Mukherjee, T., Seru, A., and Vig, V. (2010). Did securitization lead to lax screening? evidence from subprime loans. *Quarterly Journal of Economics*, 125(1):307–362.

Keys, B. J., Seru, A., and Vig, V. (2012). Lender screening and the role of securitization: evidence from prime and subprime mortgage markets. *Review of Financial Studies*, 25(7):2071–2108.

Koopman, S. J., Lucas, A., and Schwaab, B. (2009). Modeling frailty-correlated defaults using many macroeconomic covariates. *Journal of Econometrics*, 162(2):312–325.

Lando, D. (1998). On Cox processes and credit risky securities. *Review of Derivatives Research*, 2(2-3):99–120.

Lando, D. and Nielsen, M. S. (2010). Correlation in corporate defaults: Contagion or conditional independence? *Journal of Financial Intermediation*, 19(3):355–372.

Li, L. (2017). Why are foreclosures contagious? *Real Estate Economics*, 45(4):979–997.

Lin, Z., Rosenblatt, E., and Yao, V. W. (2009). Spillover effects of foreclosures on neighborhood property values. *Journal of Real Estate Finance and Economics*, 38(4):387–407.

Liu, Z. Y., Fan, G.-Z., and Lim, K. G. (2009). Extreme events and the copula pricing of commercial mortgage-backed securities. *Journal of Real Estate Finance and Economics*, 38(3):327–349.

Mayer, C., Morrison, E., Piskorski, T., and Gupta, A. (2014). Mortgage modification and strategic behavior: evidence from a legal settlement with countrywide. *American Economic Review*, 104(9):2830–57.

Meyer, P.-A. (1971). Démonstration simplifiée d'un théorème de Knight. *Séminaire de Probabilités V Université de Strasbourg*, 191:191–195.

Nickerson, J. and Griffin, J. M. (2017). Debt correlations in the wake of the financial crisis: What are appropriate default correlations for structured products? *Journal of Financial Economics*, 125(3):454–474.

Patton, A. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review*, 47(2):527–556.

Phelan, G. (2017). Correlated default and financial intermediation. *Journal of Finance*, 72(3):1253–1284.

Piskorski, T., Seru, A., and Vig, V. (2010). Securitization and distressed loan renegotiation: Evidence from the subprime mortgage crisis. *Journal of Financial Economics*, 97(3):369–397.

Rajan, U., Seru, A., and Vig, V. (2015). The failure of models that predict failure: Distance, incentives, and defaults. *Journal of Financial Economics*, 115(2):237–260.

Schönbucher, P. J. and Schubert, D. (2001). Copula-dependent default risk in intensity models. Working Paper, Bonn University.

Schuetz, J., Been, V., and Ellen, I. G. (2008). Neighborhood effects of concentrated mortgage foreclosures. *Journal of Housing Economics*, 17(4):306–319.

Shumway, T. (2001). Forecasting bankruptcy more accurately: A simple hazard model. *Journal of Business*, 74(1):101–124.

Therneau, T. M. and Grambsch, P. M. (2000). *Modeling Survival Data: Extending the Cox Model*. Springer Science & Business Media.

Yu, F. (2007). Correlated defaults in intensity-based models. *Mathematical Finance*, 17(2):155–173.