Journal of Financial Economics 129 (2018) 154–183

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Exploring the sources of default clustering☆,☆☆

S Azizpour[a], K. Giesecke[b], G. Schwenkler[c,*]

[a] Apollo Global Management, 9 West 57th Street, New York, NY 10019, United States
[b] Department of Management Science & Engineering, Stanford University, 475 Via Ortega, Stanford, CA 94305, United States
[c] Department of Finance, Boston University Questrom School of Business, 595 Commonwealth Ave, Boston, MA 02215, United States

ARTICLE INFO

Article history:
Received 17 July 2015
Revised 20 March 2017
Accepted 18 May 2017
Available online 8 May 2018

JEL classification:
G12
G33
C5
C41

Keywords:
Default clustering
Contagion
Frailty
Correlated default risk

ABSTRACT

We study the sources of corporate default clustering in the United States. We reject the hypothesis that firms' default times are correlated only because their conditional default rates depend on observable and latent systematic factors. By contrast, we find strong evidence that contagion, through which the default by one firm has a direct impact on the health of other firms, is a significant clustering source. The amount of clustering that cannot be explained by contagion and firms' exposure to observable and latent systematic factors is insignificant. Our results have important implications for the pricing and management of correlated default risk.

© 2018 Elsevier B.V. All rights reserved.

1. Introduction

The US economy has repeatedly suffered significant clusters of corporate default events. Examples include the savings and loan crisis of the early 1990s, the burst of the dotcom bubble in 2001, and the financial crisis of 2007–2009. Fig. 1, which shows the annual number of defaults of US industrial and financial firms with Moody's rated debt between 1970 and 2012, illustrates these events. This paper analyzes the potential sources of the clusters observed in the default timing data. An understanding of these sources is crucial for the measurement of portfolio credit risk at financial institutions, the management of systemic financial risk, as well as the risk analysis and valuation of securities exposed to correlated default risk, such as collateralized debt obligations.

A major source of default clustering is the joint exposure of firms to common or correlated risk factors such as interest rates, stock returns, and GDP growth. The movements of these factors cause correlated changes

☆ This paper is a significant revision of an earlier paper by the first two authors titled "Self-Exciting Corporate Defaults: Contagion vs. Frailty."

☆☆ We are grateful for grants from Moody's Corporation and the Global Association of Risk Professionals, and for data from Moody's. Schwenkler was supported by a Mayfield Stanford Graduate Fellowship and a Lieberman Fellowship. We thank Richard Cantor for providing access to data, Antje Berndt, Jens Christensen, Darrell Duffie, Andreas Eckner, Jean-David Fermanian, Lisa Goldberg, Albert Kyle, Francis Longstaff, Michael Ohlrogge, George Papanicolaou, Bjorgvin Sigurdsson, Ken Singleton, Justin Sirignano, Ilya Strebulaev, Stefan Weber, Liuren Wu, Fan Yu, seminar participants at the 5th Bachelier World Congress, the 14th FDIC/JFSR Annual Bank Research Conference, the 2014 MFA Conference, the 2015 AFA Meetings, Boston University, Columbia University, the Federal Reserve Bank of San Francisco, Georgia Institute of Technology, the International Monetary Fund, Princeton University, Stanford University, the University of Illinois at Urbana-Champaign, the University of Southern California for discussions and comments, and Xiaowei Ding, Baeho Kim, Supakorn Mudchanatongsuk, and Hannan Zheng for excellent research assistance.

* Corresponding author.
E-mail address: gas@bu.edu (G. Schwenkler).
URL: http://www.gustavo-schwenkler.com (G. Schwenkler)

https://doi.org/10.1016/j.jfineco.2018.04.008
0304-405X/© 2018 Elsevier B.V. All rights reserved.

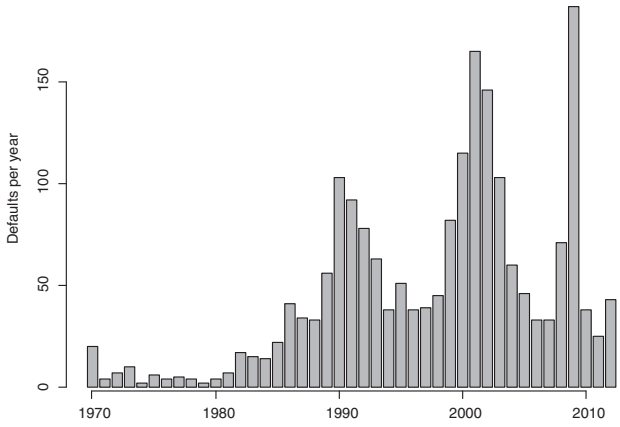


Fig. 1. Annual number of defaults of US firms with Moody's rated debt. Source: Moody's Default Risk Service.

in firms' conditional default rates. For example, strong economic growth often reduces the likelihood of default across the board. In an important paper, however, Das et al. (2007) provide strong evidence that this channel on its own cannot explain the degree of clustering observed in US industrial defaults between 1979 and 2004. The evidence suggests that there are additional sources of default clustering beyond firms' joint exposure to systematic factors. Using data on US industrial and financial default timing between 1970 and 2012, we test whether contagion, by which the default by one firm has a direct impact on the health of other firms, is a source of default clustering. We find strong evidence that contagion is a statistically and economically significant source of default clustering, after controlling for firms' joint exposure to observable and latent systematic factors.

The case for contagion as a potential source of clustering is easily made. Financial, legal, or business relationships between firms might act as a conduit for the spread of risk. For instance, a default by the protection seller in a credit swap could expose the buyer of protection and increase the default risk borne by the protection buyer's other counterparties (see Stulz, 2010). A related example is the collapse of Lehman Brothers in 2008, which pushed some of Lehman's creditors, trading partners, and clients into financial distress (see Aragon and Strahan, 2012; Brunnermeier, 2009; Chakrabarty and Zhang, 2012; Dumontaux and Pop, 2013; Fernando et al., 2012, and others). Contagion is not limited to the financial sector. For example, the default by major parts supplier Delphi in 2005 exposed General Motors, as indicated by a jump in GM's stock price and credit swap spreads. Bams et al. (2016), Boone and Ivanov (2012), Jorion and Zhang (2009), and Lang and Stulz (1992) provide evidence of default spillover effects on business partners. With contagion, the default by one firm can have a direct impact on the conditional default rates of other firms, as in the network models of Acemoglu et al. (2015), Eisenberg and Noe (2001), and Elliott et al. (2014). This impact causes correlation between firms' conditional default rates beyond the correlation due to firms' joint exposure to systematic factors.

To explore the role of contagion for default clustering, we propose a new reduced-form model of correlated default timing. The model addresses firms' exposure to observable systematic factors, a latent systematic factor, and failure events. The conditional rate of defaults is allowed to depend on time-varying factors that are observable throughout the sample period, a latent frailty factor with square-root dynamics that is not observable at all, and past failures. A default event ramps up the conditional arrival rate in a self-exciting way, representing a contagion effect.[1] The impact depends on the amount of defaulted debt and decays with time. Our model extends the standard doubly stochastic formulation widely used in theoretical and empirical analyses of correlated default risk.[2] In a doubly stochastic formulation, default arrivals are assumed to be conditionally Poisson given the paths of the observable risk factors, and the only potential source of default clustering is firms' joint exposure to these factors. Our model also extends the richer formulation of Duffie et al. (2009) with observable factors and a latent frailty factor, in which arrivals are assumed to be conditionally Poisson given the paths of all factors. The doubly stochastic assumption imposes strong restrictions on the conditional distribution of events since it precludes a direct influence of past failures on the conditional default rate. Because we do not make the doubly stochastic assumption, we avoid such restrictions. By allowing the conditional default rate to depend on past failures, our model of correlated default risk generates a much richer set of distributions.

We use our model to test several hypotheses regarding the role of contagion for default clustering. We first establish the presence of excess clustering in the default data that cannot be explained by firms' joint exposure to observable and latent systematic factors. This closes an important gap in the literature. Das et al. (2007) find that

[1] Self-exciting models that allow for the dependence of the conditional event rate on past events have been used by Aït-Sahalia et al. (2015) for analyzing the dynamics of asset returns with feedback jumps, by Bowsher (2007) for studying the dynamics of order book data, and by Jarrow and Yu (2001); Berndt et al. (2010), and others for pricing corporate bonds and credit derivatives. In contrast to these prior studies, we also model and estimate the dependence of the conditional event rate on a latent, dynamically varying frailty factor.

[2] See, for example, Chava and Jarrow (2004), Duffie and Garleanu (2001), Duffie et al. (2007), Jarrow et al. (2005), and Mortensen (2006).

the degree of clustering in the data cannot be explained by firms' joint exposure to systematic factors on its own. However, their tests ignore the potential influence of yet unidentified risk factors (i.e., frailty). There is in fact strong evidence that some of the systematic factors influencing corporate default rates have not been identified. Duffie et al. (2009), for example, provide evidence of the existence of a frailty factor that exhibits mean-reverting behavior. Chava et al. (2011), Koopman et al. (2008), Koopman et al. (2011), Monfort and Renne (2013), and others provide related evidence. This raises the question of whether there actually is excess clustering after properly controlling not just for the influence of observable factors but also for that of latent frailty factors. It is conceivable that the correlation caused by the future movements of the latent factors, as well as the uncertainty regarding the current value of these factors, is responsible for the excess clustering found by Das et al. (2007). To resolve this issue, we significantly extend the time-change tests developed by Das et al. (2007) for doubly stochastic models with observable risk factors. We show that under a change of time defined by the cumulative posterior mean of the default rate, default events follow a standard Poisson process. This result allows us to construct goodness-of-fit tests that also apply when the doubly stochastic assumption is violated because of the presence of latent frailty factors or the influence of past defaults. The tests provide strong evidence that, on its own, firms' joint exposure to observable and latent systematic risk factors cannot explain the degree of clustering in the data.

Having obtained strong evidence for the presence of excess clustering in the default timing data, we next test whether contagion is a significant source of default clustering after controlling for firms' joint exposure to observable and latent systematic factors. The likelihood estimates of our model highlight the significant dependence of conditional default rates on past failures. A default is found to have a persistent impact on the default rate, with a half-life of about three months. Upon a default event with $173 million of defaulted debt, the average in our data set, the default rate increases by roughly 2.8 events per year. A decomposition of conditional default rates indicates that the contagion channel plays a dominant role during clustering episodes, such as the crisis of 2007–2009. When the contagion channel is ignored, the role of frailty is overstated. These results strongly suggest that contagion is a statistically and economically significant source of clustering, after controlling for the influence of observable and latent systematic factors that prior studies have identified as predictors of US defaults.[3] They provide empirical support for a growing theoretical literature emphasizing a contagion channel, such as the network models of Acemoglu et al. (2015), Eisenberg and Noe (2001), Elliott et al. (2014), and many others. They also contribute to the current debate about the role of contagion in the financial versus the industrial sector:[4] we find that an industrial default event can trigger similar contagion effects as a similarly sized financial default event.

Finally, we test whether the default timing data exhibit excess clustering that cannot be explained by firms' joint exposure to observable and latent systematic risk factors as well as failure events. We study the distribution of model-implied time-changed default counts and inter-default times, again appealing to our time-change tests. We find that our fitted model, which captures the clustering caused by the influence of systematic risk factors and the contagion channel, almost perfectly matches the theoretical distribution of time-changed default counts and inter-default times implied by our time-change result. This suggests that the amount of default clustering in the data that cannot be explained by the sources we study in this paper is insignificant.

We explore the implications of our findings for the measurement of portfolio credit risk, risk premiums in credit markets, and default risk factors. Simulations are used to gauge the importance of the contagion channel for estimating portfolio credit risk out of sample. Tests of forecast accuracy indicate that a default timing model that ignores the contagion channel leads to excessively high and volatile forecasts of portfolio credit risk. This complements the results of Gouriéroux et al. (2014), who show that a CDS pricing model that ignores the contagion channel tends to overstate risk-neutral default probabilities. The accuracy tests also show that a model that does not control for the influence of latent frailty factors tends to significantly understate portfolio credit risk out of sample. Only after controlling for both frailty and contagion do we obtain accurate and sensible out-of-sample forecasts of correlated default risk in the US. These findings have important implications for the management of credit risk at financial institutions. They suggest it is necessary to address the contagion channel and control for frailty when measuring portfolio credit risk and estimating risk capital. The model of correlated default timing we propose in this paper meets this need and could replace conventional doubly stochastic models in risk management applications.

Our results also have important implications for the pricing of risk in credit markets. Several recent papers find evidence of contagion effects under the risk-neutral pricing measure. Jorion and Zhang (2007) empirically document spillovers of bankruptcies on intra-industry CDS spreads and equity prices. Berndt et al. (2010) develop a bond pricing model that allows for contagion and empirically show that the default of one firm can have a sizable impact on credit spreads of other firms. Bai et al. (2015) develop a general equilibrium model with contagion and show that pricing effects such as those found by Berndt et al. (2010) and Jorion and Zhang (2007) can be due to risk premiums for bearing exposure to contagion. Our results provide empirical support for the theoretical model of Bai et al. (2015), which assumes the existence of a contagion channel. Regressing weekly changes in bond spreads

[3] Focusing on US industrial defaults between 1982 and 2005, Lando and Nielsen (2010) find that the influence on default rates of past defaults is not significant. However, they do not control for firms' exposure to latent frailty factors.

[4] See Baur (2012), Chakrabarty and Zhang (2012), Helwege and Zhang (2016), and Lee and Poon (2014).

on weekly changes in the various components of the fitted default rate, we find that a part of the risk premiums on high-yield corporate bonds can be attributed to exposure to contagion risk.

Coval et al. (2009) provide evidence that CDO tranches referenced on the CDX North American Investment Grade Index were overpriced between 2004 and 2007, suggesting that investors did not receive adequate compensation for bearing exposure to clustered default risk. Fitting a richer model to the data, however, Collin-Dufresne et al. (2012) cannot find evidence of overpricing. Our approach allows us to measure the compensation investors require for bearing exposure to different sources of clustered default risk. Running regressions of weekly changes in CDX index and tranche spreads on weekly changes of the components of the fitted default rate that represent different sources of default clustering, we find that firms' exposure to observable systematic factors can explain a part of the changes of CDX index spreads. The contagion channel can explain a part of the spread changes of some senior CDX tranches. Consistent with Collin-Dufresne et al. (2012), these results suggest that investor may demand compensation for bearing exposure to clustered default risk. Moreover, the results indicate that different clustering sources can give rise to risk premiums that vary across securities and markets.

Our analysis, which focuses on the economy-wide default rate rather than firm-level rates, also yields insights into the importance of firm-specific factors for default timing.[5] We show that our economy-wide model fits the event timing data as well as the firm-by-firm models of Duffie et al. (2007) and Lando and Nielsen (2010), which control for systematic and firm-specific factors. These authors find that distance-to-default, a volatility-adjusted measure of leverage, is the most statistically and economically significant predictor of firm-level default in the United States. However, when we aggregate the cross-section of distance-to-default of publicly traded firms between 1991 and 2012, and regress a cross-sectional average on an extensive list of standard macroeconomic default risk factors, we find that these factors can explain about 71% of the time variation of the cross-sectional distance-to-default average. We obtain similar results for other firm-specific default risk factors, such as leverage, volatility, and short-term debt ratio. Further, the residuals of these regressions can only explain a small fraction of the portion of the default rate not tied to the macroeconomic risk factors. Thus, firm-specific risk factors are less relevant for clustered default risk, after controlling for the influence of macroeconomic and frailty default risk factors as well as the contagion channel. This finding complements the results of Bharath and Shumway (2008), Campbell et al. (2008), Tian et al. (2015), and others, who find that distance-to-default only plays a minor role for forecasting firm-level defaults. Our results also suggest that there is extensive commonality in the cross-section of idiosyncratic risk factors influencing default timing. This observation complements recent findings by Atkeson et al. (2017), who show that there is strong comovement in the cross-section of a financial soundness measure known as distance-to-insolvency.

The remainder of this paper is organized as follows. Section 2 formulates several hypotheses regarding the sources of default clustering. Section 3 discusses the default timing data. Section 4 develops our reduced-form model of correlated default timing and discusses our estimation methodology. Section 5 tests our hypotheses. Section 6 analyzes the implications of our results for the estimation of portfolio credit risk, and Section 7 studies the implications for the pricing of clustered default risk in credit markets. Robustness tests are carried out in Section 8. Section 9 concludes. There is also an online appendix.

2. Hypotheses

The literature highlights the significance of macroeconomic factors for corporate default rates. For example, Duffie et al. (2007) find that default probabilities of US industrials rise when short-term interest rates decline. Lando and Nielsen (2010) show that US industrial default probabilities also rise when the Treasury yield curve becomes steeper and when the US industrial production growth rate falls. Giesecke et al. (2011) conclude that default rates increase during periods of GDP contraction, as well as during stock market crashes and periods of high stock market volatility. The importance of the macroeconomy for default timing is unsurprising: deteriorating macroeconomic conditions negatively impact corporate earnings and the valuation of corporate assets and liabilities. For example, Longstaff and Piazzesi (2004) document procyclicality in corporate earnings, Covas and Den Haan (2011) show that debt and equity issuance are also procyclical, forcing firms to reduce investments and operations in economic downturns, and Acharya et al. (2013) demonstrate that prices of speculative grade corporate bonds fall by excessive amounts during periods of macroeconomic stress.

The joint exposure of firms to macroeconomic risk factors is a major source of default clustering: factor movements induce correlated changes in firms' conditional default rates. However, Das et al. (2007) provide strong evidence that this source on its own does not explain the degree of clustering found in US default timing data. Unfortunately, the tests proposed by Das et al. (2007) do not control for the existence of latent frailty factors influencing corporate default rates. There is in fact strong evidence that some of the systematic factors influencing corporate default rates have not been identified. Duffie et al. (2009), for example, provide evidence of the existence of a frailty factor that exhibits mean-reverting behavior. Chava et al. (2011), Koopman et al. (2008), Koopman et al. (2011), Monfort and Renne (2013), and others provide related evidence. Several theoretical models have been proposed to show how frailty can arise endogenously. For example, Benzoni et al. (2015) develop an equilibrium model of default risk in which investors are uncertain about the underlying economic state of the world. In this model, a com-

[5] Bams et al. (2016), Bhansali et al. (2008), Cont and Minca (2013), Longstaff and Rajan (2008), Giesecke et al. (2011), Giesecke et al. (2014), and Hull and White (2008) also pursue an economy-wide approach to default timing. The advantages of this approach are its relative simplicity, ease of specification, and convenience of estimation.

mon frailty factor naturally arises as a driver of the cross-section of conditional default probabilities. Using a structural model of multi-firm default, Giesecke (2004) shows how uncertainty regarding firms' liabilities has an impact on conditional default probabilities across the board.

Failure to control for the influence of latent frailty factors when measuring the default clustering caused by firms' joint exposure to systematic risk factors can lead to misleading results. When latent factors are present, the tests of Das et al. (2007) may understate the importance for default clustering of firms' joint exposure to systematic factors because of omitted variable bias. To address this issue, we will devise tests that control for frailty when measuring the importance for default clustering of firms' joint exposure to systematic risk factors. We will test the following hypothesis:

Hypothesis 1: firms' joint exposure to observable and latent systematic risk factors on its own does not explain the degree of clustering found in US default timing data.

If we find evidence supporting Hypothesis 1, then this suggests that there exist additional sources of default clustering beyond firms' joint exposure to observable or latent systematic risk factors. An additional source of clustering might be contagion, by which the default by one firm may have a direct impact on the default rates of other firms. Financial, legal, or business relationships between firms might act as a conduit for the spread of risk. A widely studied example is the downfall of Lehman Brothers in 2008, which exposed many of the firm's creditors, clients, and trading counterparties. Chakrabarty and Zhang (2012) measure the impact of Lehman's default on several measures of stock market liquidity for Lehman's counterparties. They find that Lehman's counterparties experienced a more severe deterioration of liquidity following Lehman's collapse than comparable firms that had no counterparty relations with Lehman. In the same vein, Aragon and Strahan (2012) find that stocks held by hedge funds that had counterparty exposures to Lehman experienced abnormally large market liquidity deterioration after Lehman's downfall compared to stocks that had no exposure to Lehman. Fernando et al. (2012) show that industrial firms that had investment banking relations with Lehman experienced significantly negative abnormal stock returns immediately after Lehman's collapse. More generally, Hertzel et al. (2008) show that suppliers of a defaulted firm experience negative abnormal stock performance, suggesting that financial distress can spread across industries along the supply chain of a defaulted firm. Boone and Ivanov (2012) find that a default can have a negative impact on the financial performance of a firm's strategic partners because of the inability of the defaulting firm to fulfill its partnership agreements. Mistrulli (2011) finds that the default of a firm in a conglomerate may force the parent firm to bailout the defaulted firm, raising the parent firm's own default rate. Jorion and Zhang (2009) find evidence that a default may push the trade creditors of a defaulting firm into distress. The key feature of contagion is that a default has an impact on the default rates of other firms, as in the network models of Acemoglu et al. (2012), Eisenberg and Noe (2001), Elliott et al. (2014), and many others. This motivates the following hypothesis:

Hypothesis 2: after controlling for firms' joint exposure to observable and latent systematic risk factors, contagion is a statistically and economically significant source of corporate default clustering in the US.

A key feature of Hypothesis 2 is that we control for the influence of frailty when measuring the significance of the contagion channel. Failure to control for frailty may lead to biased estimates of the contagion channel given support for Hypothesis 1. Bias due to omitted variables may explain the findings of Lando and Nielsen (2010), who reject the hypothesis that contagion is a significant source of clustering. They only control for firms' exposure to observable risk factors.

The existence of evidence in support of Hypotheses 1 and 2 does not exclude the possibility of there being excess default clustering that cannot be accounted for by firms' joint exposure to observable and latent systematic factors as well as past failures (i.e., contagion). This motivates testing a third hypothesis:

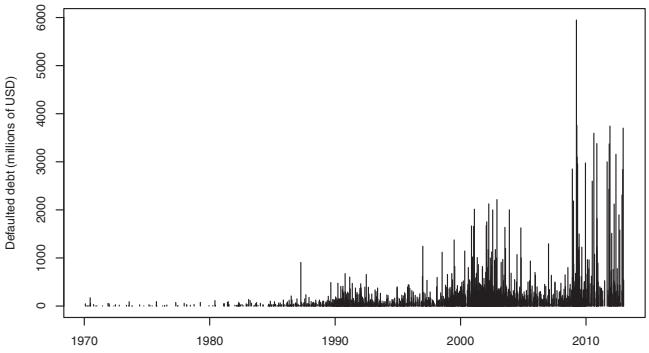
Hypothesis 3: there is significant excess default clustering not explained by firms' joint exposure to observable and latent systematic risk factors as well as contagion.

3. The data

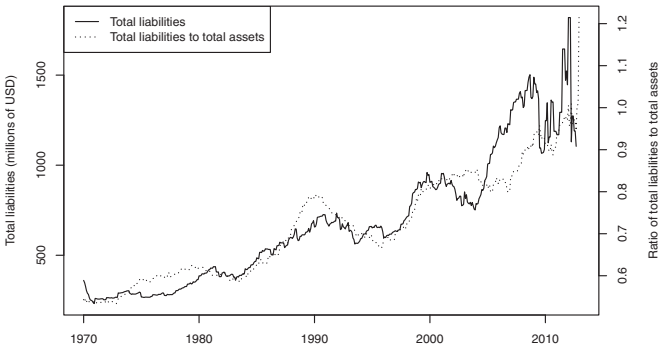
We will test the hypotheses developed in Section 2 using data on industrial and financial default timing in the US between 1/1/1970 and 12/31/2012. The data cover all defaults of firms domiciled in the US with debt rated by Moody's and were obtained from Moody's Default Risk Service.

We adopt Moody's definition of a default event (see Hamilton, 2005). A default is an event in any of the following categories: (1) a missed or delayed disbursement of interest or principal, including delayed payments made within a grace period; (2) bankruptcy (Section 77, Chapter 10, Chapter 11, Chapter 7, prepackaged Chapter 11), administration, legal receivership, or other legal blocks to the timely payment of interest or principal; or (3) a distressed exchange occurs where (i) the issuer offers debt holders a new security or package of securities that amount to a diminished financial obligation, or (ii) the exchange had the apparent purpose of helping the borrower avoid default. A repeated default by the same issuer is included in the set of events if it was not within a year of the initial event and the issuer's rating was raised above CAA after the initial default. This treatment of repeated defaults is consistent with that of Moody's.

We observe a total of 2001 defaults by 1294 firms. There are 1452 days with default events. Fig. 2 shows the face amount of defaulted debt for each day in our sample period. The date with the largest amount of defaulted debt is March 27, 2009, when Charter Communications Holdings LLC defaulted. April 8, 2009, the day on which Harrah's Operating Company collapsed, is next on the list. The third largest event is the default of American Airlines on November 29, 2011. We observe that the face value of defaulted debt tends to increase over the sample period. This is primarily due to the fact that the firms in our default data set have increased both their total liabilities and their leverage ratios over time; see Fig. 2(b).



(a) Time series of daily defaulted debt.



(b) Time series of total liabilities and leverage across all surviving firms in our default data set.

Fig. 2. Time series of defaulted debt and leverage.
In Fig. 2(a) we display the time series of the sum of defaulted debt on every day of our sample period. In Fig. 2(b) we show the time series of average total liabilities and average ratio of total liabilities to total assets across all surviving firms in our default data set. Data on total liabilities and total assets is obtained from the quarterly Compustat file when available and otherwise from the annually Compustat file.

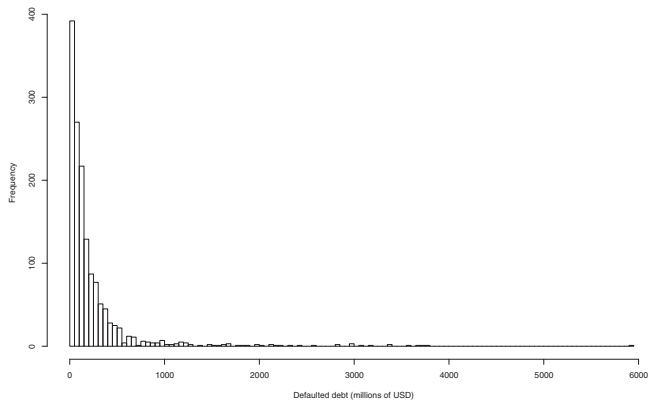
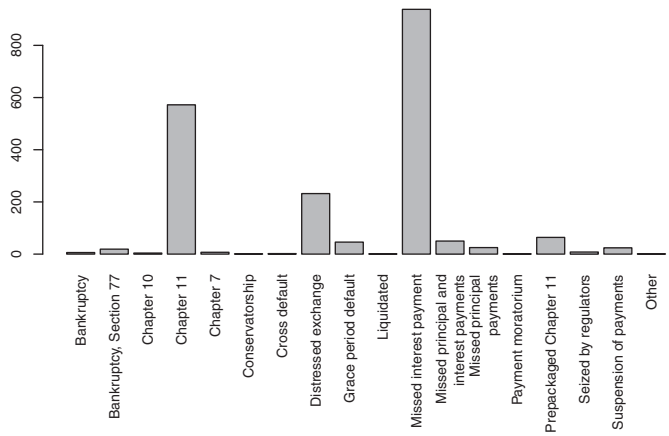


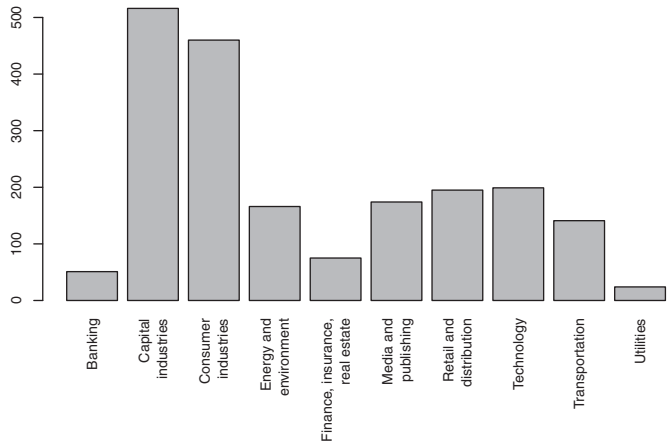
Fig. 3. Histogram of defaulted debt.

Fig. 3 reports the distribution of defaulted debt. About 95% of the defaults involve $500 million or less of defaulted debt. The average amount of defaulted debt is $173 million; the median is $100 million. Fig. 4 shows the number of defaults by event category. The vast majority of defaults is due to missed interest payments, Chapter 11 filings, or exchanges of distressed debt that resulted in losses to bond investors. Fig. 4 also indicates the number of events by sector. Most defaults occur in the capital industry, followed by the consumer industry, technology, retail and distribution, and media and publishing sectors.

Fig. 5 shows the joint distributions of the sector of a defaulting firm, the event category, the amount of defaulted debt, and the time until the next default. Defaults involving Chapter 11 filings, distressed exchanges, and missed principal payments tend to be quickly followed by additional defaults. These types of defaults also tend to be large in terms of the amount of defaulted debt. Similarly, defaults



(a) Number of defaults by Moody's event category.



(b) Number of defaults by industry sector according to Moody's 11 Code.

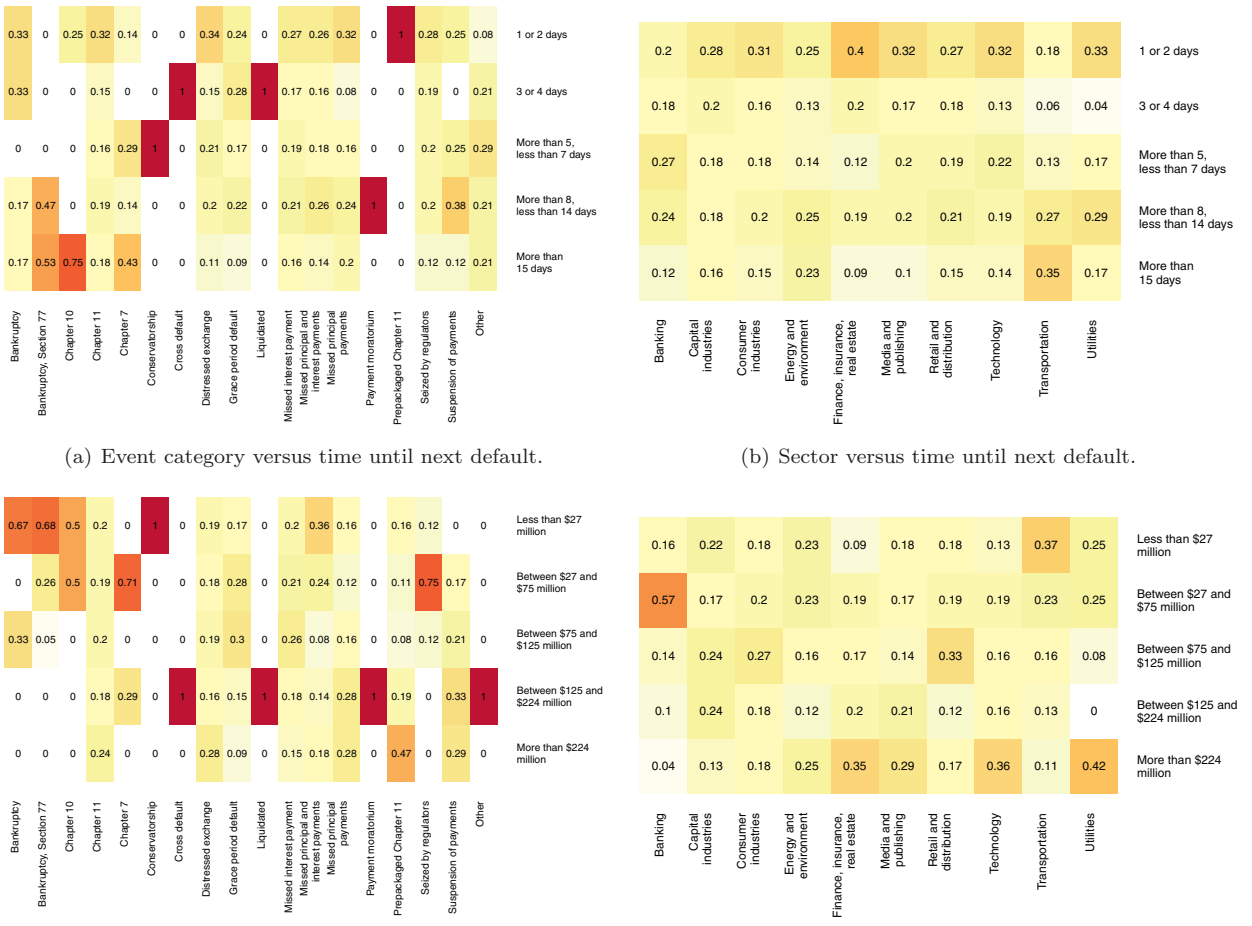
Fig. 4. Defaults by event category and industry sector.

in the finance, insurance, real estate, consumer industry, media and publishing, and technology sectors tend to be quickly followed by additional defaults. These sectors also tend to have large amounts of defaulted debt. Fig. 6 indicates that defaults with large amounts of debt tend to be quickly followed by additional defaults. Fig. 6 also suggests that the impact of large defaults on the arrival of additional defaults tends to be persistent: there is higher cumulated defaulted debt over the one-month period after a large default than in the one-month period after a small default. These observations will inform the design of our reduced-form model in Section 4.

We will study the influence on default timing of a number of explanatory variables (risk factors). We construct time series of variables from data obtained from the Federal Reserve Banks of New York and Saint Louis; Table 1 reports summary statistics. The variables we consider include the quarterly growth rate of the US industrial production calculated from final products and nonindustrial supplies (IP, observed monthly), the annualized quarterly growth rate of the US gross domestic product (GDP, observed quarterly), the trailing one-year return of the S&P 500 index (RET, observed daily), the trailing one-year volatility of the S&P 500 index (VOL, observed daily), the level of the yield curve measured by the yield of the three-month Treasury bill yield (LEVEL, observed daily), the slope of the yield curve measured by the spread between the ten-year Treasury note and the one-year Treasury bill (SLOPE, observed daily), the ten-year Treasury note yield (10Y, observed daily), Moody's AAA corporate yield (AAA, observed weekly), the credit spread between the Moody's BAA and AAA corporate yields (CRED, observed weekly), and a recession indicator according to the NBER business cycle dates (REC, observed monthly). Under different econometric assumptions and using different default timing data, Chava and Jarrow (2004), Das et al. (2007), Duffie et al. (2007), Campbell et al. (2008), Duffie et al. (2009), Lando and Nielsen (2010), Giesecke et al. (2011), and others have found some or all of these variables to be significant predictors of US defaults.

4. Reduced-form model

We develop a new reduced-form model of correlated default timing to test the hypotheses formulated in Section 2. Reduced-form models are widely used in theo-



(a) Event category versus time until next default.

(b) Sector versus time until next default.

(c) Event category versus quintiles of defaulted debt.

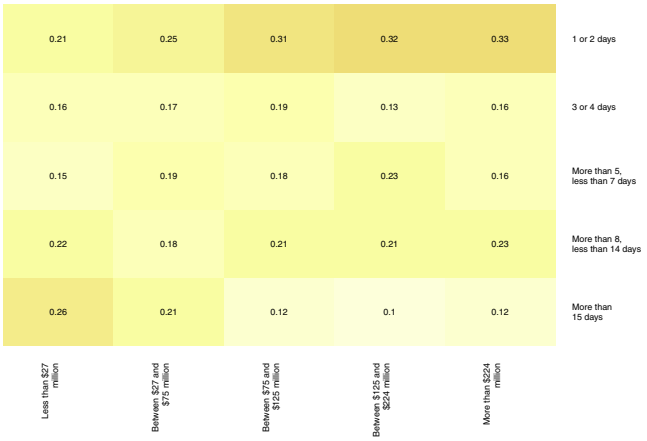
(d) Sector versus quintiles of defaulted debt.

Fig. 5. Joint distributions.
These charts show different empirical bivariate distributions of: (a) default type versus time until the next default, (b) sector of defaulted firm versus time until the next default, (c) default type versus amount of defaulted debt, and (d) sector of defaulted firm versus amount of defaulted debt. The breakdowns correspond to the quintiles of each variable.

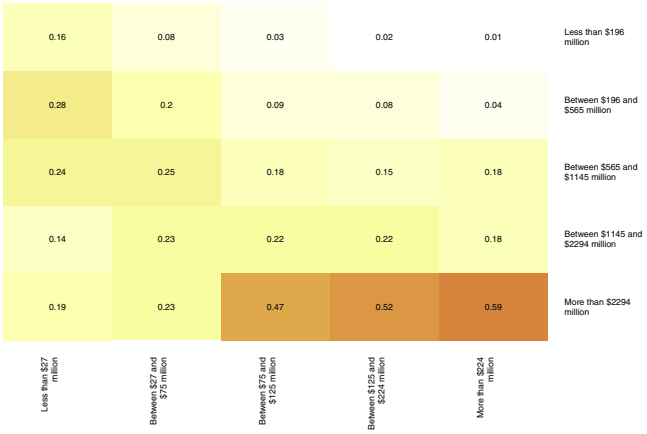
Table 1
Summary statistics of explanatory variables. This table provides summary statistics for the time series of macroeconomic explanatory variables.

Variable	Sample frequency	Mean	Standard deviation	Skewness	Kurtosis	Min.	First quartile	Median	Third quartile	Max.	Number of observations
IP	Monthly	2.36	4.11	−0.70	3.95	−11.78	0.62	2.76	4.90	11.22	516
GDP	Quarterly	6.71	4.27	0.75	5.99	−8.40	4.40	6.10	8.60	25.50	172
RET	Daily	8.23	17.32	−0.39	3.11	−48.82	−2.14	9.95	20.35	68.57	11218
VOL	Daily	17.46	7.50	0.44	2.50	2.11	11.06	16.97	22.43	37.07	11218
LEVEL	Daily	5.28	3.23	0.58	3.75	0.00	3.30	5.10	7.13	17.14	10740
SLOPE	Daily	−1.12	1.20	0.35	2.91	−3.53	−1.93	−1.20	−0.30	3.44	10740
10Y	Daily	6.98	2.79	0.61	3.27	1.43	4.89	6.79	8.33	15.84	10740
AAA	Weekly	8.05	2.41	0.78	3.53	3.26	6.45	7.73	9.04	15.85	2243
CRED	Weekly	1.11	0.46	1.78	7.26	0.52	0.80	0.98	1.30	3.47	2243
REC	Monthly	0.19	0.40	1.55	3.40	0.00	0.00	0.00	0.00	1.00	516

retical and empirical analyses of correlated default risk; see Berndt et al. (2010), Chava and Jarrow (2004), Das et al. (2007), Duffie et al. (2007), Duffie et al. (2009), Jarrow and Yu (2001), Jarrow et al. (2005), Lando and Nielsen (2010), Longstaff and Rajan (2008), and many others. Unlike the models developed in the literature, our model captures the influence of observable, time-varying risk factors, of a dynamically varying frailty, and of past defaults. The model provides portfolio loss forecasts for multiple future periods that take account of the time-series behavior of the observable and latent risk factors over the forecast period.



(a) Defaulted debt versus time until next default.



(b) Defaulted debt versus cumulated defaulted debt over the 30-day period post default.

Fig. 6. Impact of defaulted debt.
These charts show the empirical joint distributions of the amount of defaulted debt and the amount of days until the next default occurs (top chart), and the amount of defaulted debt and the cumulated defaulted debt over the 30-day period post default. The breakdowns of the buckets correspond to the quintiles of each variable.

4.1. Model formulation

The defaults in our sample arrive at an intensity that measures the conditional mean arrival rate of events per year (see Appendix A for details). To accommodate the potential sources of default clustering, we allow the intensity to depend on observable explanatory variables $X_1, \ldots, X_d$, a latent frailty variable Z, and past defaults. More precisely, if Z were observable, the intensity at time t would be of the form

$$\lambda_t = \exp\left(a_0 + \sum_{i=1}^{d} a_i X_{i,t}\right) + Y_t + Z_t, \tag{1}$$

where $a = (a_0, a_1, \ldots, a_d)$ is a vector of coefficients to be estimated, and Y, which will be specified below, models the influence of past defaults on the conditional default rate. The intensity at t given the actually available information is the posterior mean of λ_t; details are provided in Internet Appendix B.

The first term in Eq. (1) captures firms' exposure to the common observable risk factor $X = (X_1, \ldots, X_d)$, which represents the variables described in Section 3 above (see Table 1). As in Chava and Jarrow (2004), Das et al. (2007), Duffie et al. (2007), and other articles, this term takes the classical proportional hazards form of Cox (1972). When Y and Z are positive, this formulation ensures that the intensity is positive for any value of the factor vector X. Moreover, when $Y = 0$ and $Z = 0$, the influence of the risk factor X_i is multiplicative and the parameter a_i is a proportionality factor. If X_i increases by one unit, then the intensity increases by a factor of $e^{a_i} - 1$.

The term Z in Eq. (1) models the influence of a common latent frailty factor on the timing of defaults in our sample.[6] We propose a standard mean-reverting Cox-

[6] Due to identifiability restrictions, we can only allow for a one-dimensional frailty factor.

Ingersoll-Ross (CIR) model for the dynamics of the latent factor Z:

$$dZ_t = k(z - Z_t)dt + \sigma\sqrt{Z_t}dW_t, \tag{2}$$

for z, k, $\sigma \geq 0$, and $2kz \geq \sigma^2$.[7] Here, W is a standard Brownian motion independent of X, and Z is initialized at its stationary gamma distribution with shape parameter $\frac{2kz}{\sigma^2}$ and scale parameter $\frac{\sigma^2}{2k}$. The choice of a mean-reverting model is motivated by the results of Duffie et al. (2009), who find strong evidence for the existence of a mean-reverting frailty influencing US nonfinancial default timing between 1979 and 2004.

The term Y in Eq. (1) addresses the influence of past failures on default timing. We propose the classical self-exciting specification of Hawkes (1971), which is widely used in the empirical and theoretical contagion literature:[8]

$$Y_t = b\sum_{n:\,T_n \leq t} e^{-\kappa(t-T_n)}U_n, \tag{3}$$

where $b > 0$ and $\kappa > 0$ are parameters to be estimated. The increasing sequence of times T_n represent the dates in the sample with at least one default, and the variables U_n characterize the impact of defaults on the conditional default rate. The intensity is ramped up by bU_n given at least one event at T_n. The impact fades away with time at rate κ. Letting u_n be the sum of the face values of defaulted debt at T_n, we take

$$U_n = \max\{0, \log u_n\}. \tag{4}$$

We justify the choice in Eq. (4). Fig. 6 indicates that defaults with large amounts of defaulted debt tend to be quickly followed by additional defaults, and that the impact of defaults with large amounts of defaulted debt on the arrival of future defaults is persistent. Fig. 5 suggests that the variability in defaulted debt can capture the cross-sectional variability of the impact of the industry sector and event category on the arrival of future defaults. We interpret that defaults with large amounts of defaulted debt have a strong impact on the default rate, regardless of other default characteristics.

Hertzel et al. (2008), Jorion and Zhang (2007, 2009) provide evidence that bankruptcies of firms with large liabilities induce stronger contagion effects than bankruptcies of firms with small liabilities. For the global network of large CDS counterparties, Yang and Zhou (2013) document that credit risk spillovers from distressed firms are more severe if the distressed firm has large leverage ratios. Using data of firm liabilities from Compustat, we find that the amount of defaulted debt is correlated with total liabilities at default. Among the defaulted firms in our dataset for which liabilities data is available, the correlation coefficient between the logarithm of defaulted debt and the logarithm of total liabilities at default is 0.39. Table 2 indicates that a regression of total liabilities at default on defaulted debt has an adjusted R-squared of 0.15 without controlling for the influence of the explanatory variables X and 0.27 after controlling for X. Fig. 7 indicates that the quantiles of the logarithm of defaulted debt closely match the quantiles of the logarithm of total liabilities at default, except for a few defaults with very small and very large amounts of defaulted debt. As a result, defaulted debt provides a proxy for total liabilities at the time of default. Based on these observations and because liability data is only available for about one-third of the firms in our data set, we specify the impact of a default event on the conditional default rate in terms of defaulted debt and Eq. (4). Overall, our contagion factor roughly captures aggregate credit exposure to defaulted firms, consistent with the contagion mechanism of Jorion and Zhang (2007, 2009).

Table 2
Coefficient estimates and standard errors of ordinary least squares regressions of total liabilities.
This table presents the results of ordinary least squares regressions of the logarithm of the total liabilities of a firm at default on the logarithm of the defaulted debt and our macroeconomic variables at the time of default. Data on total liabilities is obtained from the quarterly Compustat file when available and otherwise from the annually Compustat file. * indicates significance at the 95% level, ** significance at the 99% level, and *** significance at the 99.9% level.

Variable	Regression 1	Regression 2
Constant	*** 1.4334	*** 3.5208
	(0.2583)	(0.5284)
log(DEFDEBT)	*** 0.4556	*** 0.3581
	(0.0415)	(0.0408)
IP		0.0250
		(0.0197)
GDP		0.0078
		(0.0217)
RET		−0.0046
		(0.0041)
VOL		0.0094
		(0.0076)
LEVEL		0.0057
		(0.1904)
SLOPE		0.0985
		(0.2147)
10Y		−0.2689
		(0.1438)
CRED		−0.1444
		(0.1708)
AAA		0.0249
		(0.1971)
REC		0.0915
		(0.1762)
Adjusted R^2	0.155	0.269
Data points	2001	2001

4.2. Likelihood estimation

We estimate the parameter $\theta = (a, b, k, \kappa, \sigma, z, s)$ governing the reduced-form model of Eq. (1) by the method of maximum likelihood. Since the frailty Z influencing the intensity cannot be observed, the likelihood problem is not standard unless $z = \sigma = 0$. We implement a variant of the

[7] We have considered alternative parameterizations of the model (2), including one with unit volatility but Z in Eq. (1) replaced by σZ. The results for these alternatives are very similar to those for Eq. (2).

[8] See Aït-Sahalia et al. (2015), Bowsher (2007), Errais et al. (2010), and others.

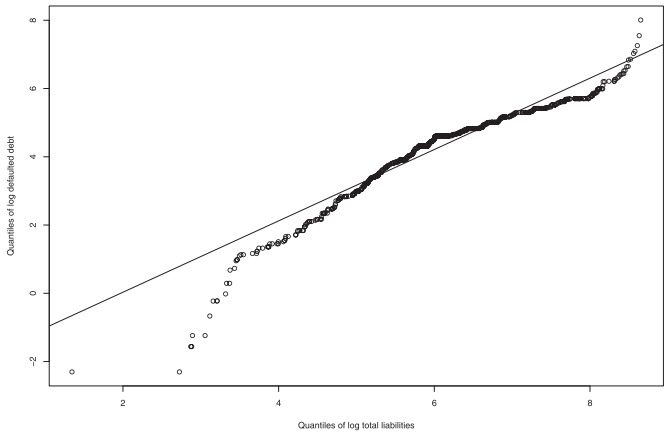


Fig. 7. QQ-plot of log defaulted debt versus log total liabilities.
This figure maps the realized quantiles of the logarithm of defaulted debt onto the quantiles of the logarithm of total liabilities at default of a firm. The slope of the line is 1.05. Number of observations: 650.

Table 3
Reduced-form model alternatives.
The table reports the intensity specification for each reduced-form model we estimate. Each model addresses a different set of clustering sources.

Model	Parameter restrictions	Intensity
Complete	None	$\lambda_t = \exp(a_0 + \sum_{i=1}^{d} a_i X_{i,t}) + Y_t + Z_t$
Base + Frailty	$b = 0$	$\lambda_t = \exp(a_0 + \sum_{i=1}^{d} a_i X_{i,t}) + Z_t$
Base + Contagion	$z = \sigma = 0$	$\lambda_t = \exp(a_0 + \sum_{i=1}^{d} a_i X_{i,t}) + Y_t$
Base	$b = z = \sigma = 0$	$\lambda_t = \exp(a_0 + \sum_{i=1}^{d} a_i X_{i,t})$

filtered likelihood estimation developed by Giesecke and Schwenkler (2018); see Internet Appendix A for details. The likelihood is the posterior mean of the complete-data likelihood given the observed data, which include the default dates, the defaulted debt, and time series of the explanatory variables $X_1, \ldots, X_d$ discussed in Section 3.[9] The posterior mean is evaluated using a quadrature method. This approach eliminates the need to perform Monte Carlo simulations of the frailty path over the 43-year sample period, which would be computationally burdensome. It also avoids the loss of estimation efficiency and potential biases associated with the simulated expectation-maximization (EM) algorithm often used for this type of problem.[10] The resulting likelihood estimators are consistent and asymptotically normal as the sample period grows (see Internet Appendix A.3).

4.3. Parameter estimates

We fit our model along with three nested (i.e., restricted) alternatives that are described in Table 3. A model for which $b = 0$ ignores contagion. A model for which $z = \sigma = 0$ ignores frailty. A base model for which $b = z = \sigma = 0$ captures only the clustering implied by firms' joint dependence on the observable risk factor X. Table 4 reports the parameter estimates and the corresponding asymptotic standard errors.

The self-excitation term Y is highly significant. Past defaults appear to be significant predictors of future defaults, after controlling for the influence of observable and frailty factors. A default has an economically significant and persistent impact on the intensity, with a half-life of about three months. Upon a default event with $173 million of defaulted debt, the average in our data set, the intensity increases by roughly 2.8 events per year. The fitted values of the parameters b and κ governing Y are very similar to the fitted values in the model omitting the frailty factor. This observation suggests that the term Y has been properly identified. In particular, the role of Y is not overstated at the expense of the frailty term. Identifiability is further discussed in Section 4.5.

The latent frailty Z exhibits significant mean-reverting behavior. A one standard deviation shock to Z increases the aggregate default rate by about 0.75 defaults per year, and this effect dissipates quickly with a half-life of about one day. Even though our data spans a longer time horizon and includes defaults of financial firms, the parameter estimates for our frailty factor in the restricted model are almost equal to the estimates for the frailty factor in the comparable model of Duffie et al. (2009).[11] However, the

[9] Following the empirical default timing literature, we interpolate the discrete observations of the explanatory variables to obtain continuous time series.

[10] Duffie et al. (2009); Koopman et al. (2008, 2011), and others use EM algorithms to estimate the influence of frailty on default timing. Giesecke and Schwenkler (2018) indicate that EM algorithms may be unable to precisely differentiate between jump and diffusive movements of the intensity when the intensity depends on a latent frailty factor and a self-exciting factor as in Eq. (1).

[11] Duffie et al. (2009) estimate that their frailty factor has a volatility of 0.125 and a mean-reversion speed of 0.018. We estimate that our frailty

Table 4
Parameter estimates for each of the four model alternatives described in Table 3. The asymptotic standard errors are given parenthetically. They are computed using the Hessian matrix of the log-likelihood at the parameter estimates reported. Some of the variables of Table 1 are not included here due to lack of significance. * indicates significance at the 95% level, ** significance at the 99% level, and *** significance at the 99.9% level.

	Complete	Base	Base + Contagion	Base + Frailty
Constant a_0	** −1.6926	*** 1.0141	** −1.6639	0.5649
	(0.6765)	(0.1349)	(0.6453)	(0.8971)
IP coefficient a_1	−0.0013	−0.0509	−0.0027	−0.0616
	(0.0334)	(0.0086)	(0.0318)	(0.0471)
GDP coefficient a_3	*** −0.1235	*** −0.1117	*** −0.1197	0.0031
	(0.0360)	(0.0098)	(0.0343)	(0.0391)
RET coefficient a_4	0.0017	* −0.0045	0.0015	−0.0193
	(0.0078)	(0.0018)	(0.0074)	(0.0173)
VOL coefficient a_5	0.0038	** 0.0112	0.0041	0.0123
	(0.0208)	(0.0038)	(0.0222)	(0.0246)
LEVEL coefficient a_6	0.0284	** −0.0466	0.0277	*** −0.6734
	(0.0612)	(0.0149)	(0.0857)	(0.1546)
SLOPE coefficient a_8	0.0007	* 0.0821	0.0019	−0.4748
	(0.1751)	(0.0352)	(0.2770)	(0.2836)
CRED coefficient a_9	−0.01804	*** −0.7305	−0.0192	−0.0971
	(0.3907)	(0.0741)	(0.4253)	(0.3035)
Frailty volatility σ	0.2221			*** 0.1288
	(0.1602)			(0.0146)
Frailty mean reversion rate k	* 6.0203			*** 0.0156
	(2.9545)			(0.0002)
Frailty mean reversion level z	* 0.0041			*** 0.5335
	(0.0020)			(0.0577)
Contagion sensitivity b	*** 0.0103		*** 0.0101	
	(0.0017)		(0.0017)	
Contagion decay rate κ	*** 0.0571		*** 0.0563	
	(0.0096)		(0.0098)	
Log-likelihood	606.29	446.12	606.21	599.24

fitted parameters in the unrestricted model differ by orders of magnitude from those in the model ignoring contagion. On one hand, the fitted frailty has small volatility and quickly reverts to a low but significant level in the unrestricted model. On the other hand, the fitted frailty is very persistent and has a high reversion level and high volatility in the restricted model. These findings suggest that a model ignoring the contagion channel may overstate the role of frailty.

Among the observable factors, GDP growth has a significant influence on the intensity. The sensitivity of the intensity to this factor is negative, consistent with the intuition that strong economic growth reduces the likelihood of defaults (see Giesecke et al., 2011). Moreover, the fitted value of the sensitivity parameter is similar to the values estimated for the restricted models (if these values are significant), suggesting that the impact of GDP growth on the intensity is not overstated at the expense of other risk factors or the frailty and self-exciting terms. A one standard deviation shock to GDP growth may increase the aggregate default intensity by up to 40%.

Unlike Duffie et al. (2007, 2009), and Lando and Nielsen (2010), we find that the S&P 500 return is a significant risk factor only in the base model. Our parameter estimates suggest that this is due to the fact that we control for frailty and contagion when measuring the impact of the explanatory variables. The aforementioned papers control for only one of these effects, or neither. Also the corporate bond yield spread is found to be insignificant. This finding complements previous results by Collin-Dufresne et al. (2001), Giesecke et al. (2011), and Gilchrist and Zakrajšek (2012) on the connection between credit spreads and the business cycle. It provides additional evidence that credit spreads contain little information about physical default probabilities.

4.4. Fitted intensity

Fig. 8 compares the fitted intensity to the default data. The fitted intensity at time t is given by the fitted posterior mean of λ_t given the data observed to time t. Without frailty ($z = \sigma = 0$), the posterior mean, h_t, is equal to λ_t. With frailty, the computation of h_t is a filtering problem; see Internet Appendix B.

The base model does a relatively poor job at explaining the default timing data. It lags clusters and significantly overstates the arrival rates during the first quarter of the sample period and during the financial crisis of 2007–2009. The richer models capture the substantial time-variation of realized arrival rates much better. The models addressing the contagion channel match the low arrival rates during the early part of the sample better than the models that ignore that channel.

factor in the Base + Frailty model has a volatility of 0.129 and a mean-reversion speed of 0.016.

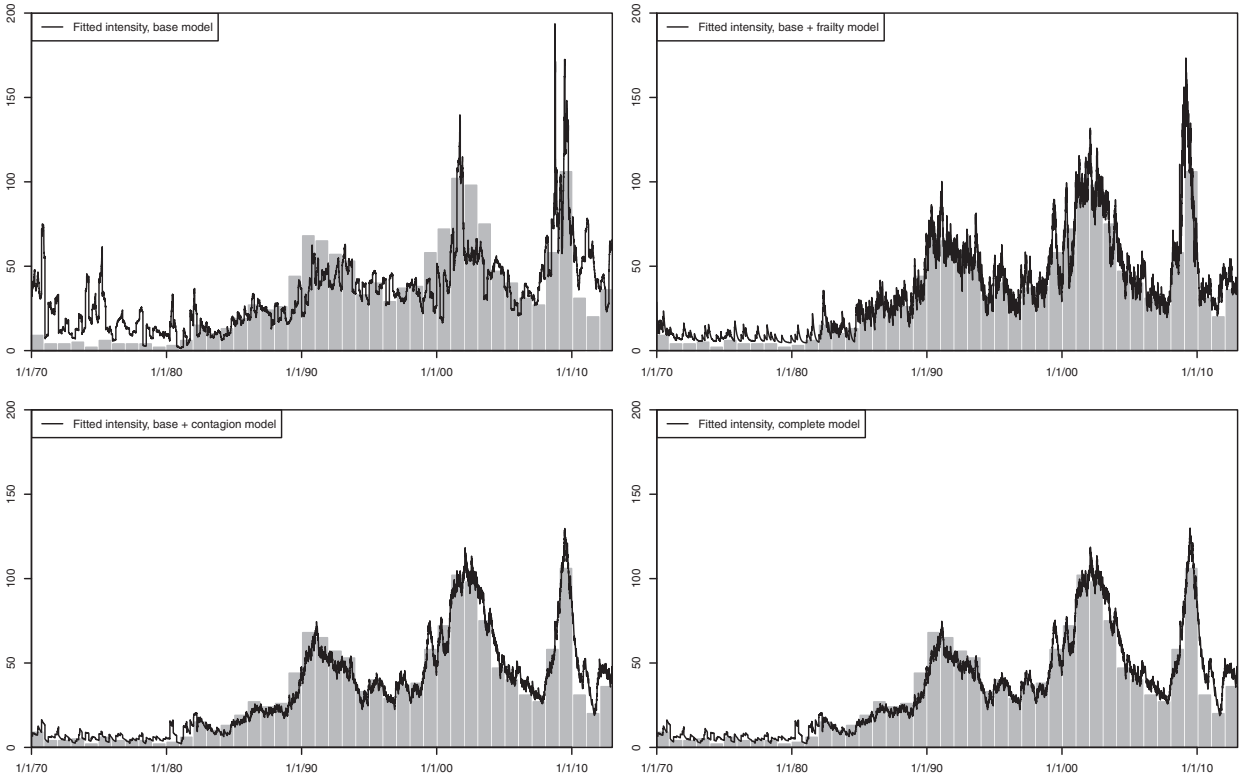


Fig. 8. Fitted intensities.
A figure shows the fitted intensity (measured in defaults per year) implied by the maximum likelihood estimates reported in Table 4. A gray bar represents the realized number of defaults per year.

We examine the relative contributions to the fitted intensity h_t of each of the components of the unrestricted model by decomposing h_t into the sum of the contagion term Y_t, the filtered frailty factor $\mathbb{E}_t[Z_t]$ (the posterior mean of Z_t), and a term $e^{a_0+\sum_{i=1}^{d} a_i X_{i,t}}$ capturing the exposure to observable systematic factors. Fig. 9 shows that, since the mid 1980's, the latter component consistently accounts for 25% to 35% of the default rate. This finding is consistent with the estimates of Koopman et al. (2012) and suggests that we have properly controlled for the exposure to observable systematic risk factors. Fig. 9 also indicates that the contagion term represents a significant fraction of the default rate and that this fraction is especially large during clustering episodes. In contrast, the contribution of the frailty term to the default rate is small and only becomes larger during episodes of low default activity.

4.5. Identification

Simulations are used to study the ability of our estimation approach to identify the three different components of our model (observable factors, frailty, and contagion). We generate 50 sample data sets from each of the four models in Table 3 via Monte Carlo simulation, and then estimate the complete (i.e., unrestricted) model for each of these data sets. We seek to determine whether the complete model can identify a model component if it was present in the model used to sample the data.

Table 5 reports our findings. The exposure to the GDP growth rate is always found to be significant, and the latent frailty factor is found to be significant only when it is actually present in the data. This confirms our ability to identify the exposure to systematic risk factors, even if one of the factors is a frailty. The contagion factor is found significant only when it is actually present in the data. The z-scores of the estimates also indicate that the data-generating parameters are precisely estimated even when the estimated model is misspecified because it includes factors that are not present in the data. The fitted complete model does not overstate the role of the contagion term when this term is present in the data. The results of this analysis suggest that our approach can accurately identify and differentiate between the sources of clustering.

5. Hypotheses tests

5.1. Hypothesis 1

We develop time-scaling tests to test Hypothesis 1 regarding the ability of firms' joint exposure to observable and latent systematic risk factors to explain the degree of default clustering found in the default timing data.

An important result of Das et al. (2007) states that if there are no latent frailty factors and defaults are doubly

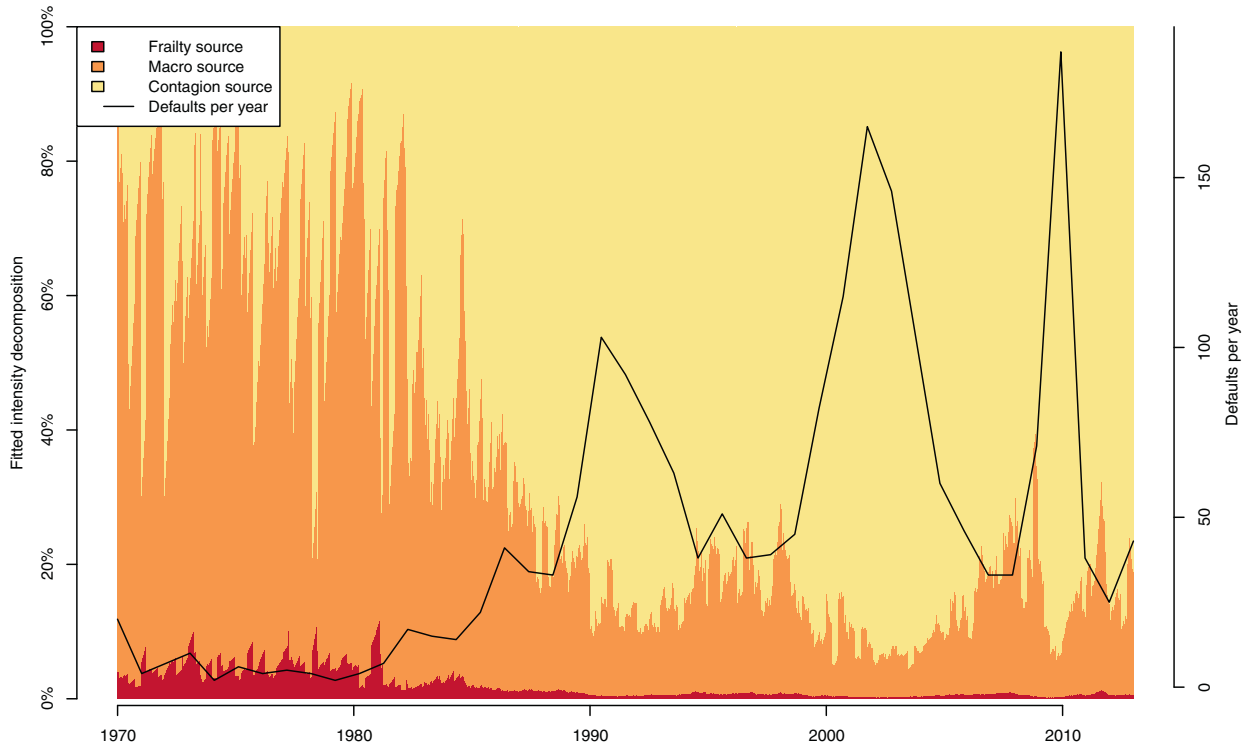


Fig. 9. Decomposition of fitted default intensity.
This figure decomposes the time series of the fitted default intensity, $\mathbb{E}_t[\lambda_t]$, into its macro, contagion, and frailty components, and compares to the number of defaults in each year. The fitted intensity at any point of time is normalized to unity so that each component is displayed as percentage of the fitted intensity. The fitted intensity and its components are computed as described in Appendix B using the parameter estimates of Table 4.

stochastic,[12] then the default count N_t can be transformed into a standard Poisson process by a change of time given by the cumulative intensity. Our model violates the doubly stochastic hypothesis because it addresses the contagion channel and it includes an unobservable frailty factor. Therefore, the tests of Das et al. (2007) do not apply. We generalize the time-change result of Das et al. (2007) in Proposition 1 in Appendix B. Our result also covers models in which the doubly stochastic assumption is violated because of the presence of a latent frailty factor or a self-exciting term. We show that the default count N_t can always be transformed into a standard Poisson process by a change of time given by the cumulative posterior mean h_t of the intensity. This result allows us to compare the goodness-of-fit of models addressing different sources of default clustering by testing the Poisson property of the time-scaled default times.

For a bin of size $w > 0$, we construct an increasing sequence of times t_i^w so that the cumulative posterior intensity mean h_t between two consecutive times is equal to w. Formally, $\int_{t_{i-1}^w}^{t_i^w} h_s \mathrm{d}s = w$ for each i. Let P_i^w be the number of defaults between t_{i-1}^w and t_i^w. Proposition 1 implies that the P_i^w are independent samples from a Poisson distribution with parameter w. We test this property for different bin sizes w using a series of tests described in Karlis and Xekalaki (2000) and Das et al. (2007). The Potthoff-Whittinghill-Bohning test evaluates the moments and the Kocherlakota-Kocherlakota test evaluates the generating function of the P_i^w. The other tests examine the independence, dispersion, tail properties, and distribution of the P_i^w.

Table 6 reports the p-values of the Poisson property tests. The base model, which only controls for firms' joint exposure to observable systematic factors, fails most of these tests at high significance levels. It generates binned event counts that are correlated and does not match the corresponding Poisson distributions. Table 7 suggests that the base model is rejected because it generates binned event counts that are overdispersed, with the variances of the P_i^w exceeding the theoretical counterparts. Overdispersed counts are a sign of model misspecification (see Cameron and Trivedi, 1998). Consistent with Das et al. (2007) but in contrast to Lando and Nielsen (2010), our results indicate that a model of clustered default risk that only controls for firms' exposure to observable systematic factors can be rejected because it is misspecified. Unlike the aforementioned studies, however, our rejection is

[12] A model with observable risk factors is doubly stochastic if default arrivals are conditionally Poisson given the paths of the factors. A doubly stochastic model addresses only the default correlation caused by firms' joint exposure to the risk factors.

Table 5
Average parameter estimates of a restricted version of the complete model of Table 3. The estimates in each column are derived from 50 data samples simulated from the following submodels:

Complete: $a_0 = -1.6$, $a_3 = -0.12$, $\sigma = 0.2221$, $k = 6.0203$, $z = 0.0041$, $b = 0.0103$, $\kappa = 0.0571$,
Base: $a_0 = -1.6$, $a_3 = -0.12$,
Base + Contagion: $a_0 = -1.6$, $a_3 = -0.12$, $b = 0.0101$, $\kappa = 0.0563$,
Base + Frailty: $a_0 = -1.6$, $a_3 = -0.12$, $\sigma = 0.1300$, $k = 0.0200$, $z = 0.5300$.

These choices are inspired by the estimates of Table 4. For simplicity, we restrict the complete model to only include the GDP growth rate as covariate given that this is the only significant factor in Table 4. The time horizon is the same as in Table 4. Empirical standard deviations of the estimated parameters are given parenthetically. We compute significance levels from a centralized normal distribution with standard deviations given by the empirical standard deviations of the parameter estimates. * indicates significance at the 95% level, ** significance at the 99% level, and *** significance at the 99.9% level. The z-score of the parameter estimates relative to the true data-generating parameters are given in brackets when applicable.

	Data sampled from the model:			
	Complete	Base	Base + Contagion	Base + Frailty
Constant a_0	*** −1.5214	−2.4421	*** −1.5404	*** −1.6012
	(0.1964)	(2.0002)	(0.1860)	(0.3265)
	[0.4000]	[−0.4210]	[0.3206]	[−0.0037]
GDP coefficient a_3	*** −0.1194	* −0.1714	** −0.0993	** −0.1129
	(0.0044)	(0.0838)	(0.0272)	(0.0543)
	[0.1422]	[−0.6130]	[0.7632]	[0.1305]
Frailty volatility σ	*** 0.2150	0.0583	0.0004	*** 0.1294
	(0.0135)	(0.1472)	(0.0015)	(0.0009)
	[−0.5255]			[−0.6113]
Frailty mean reversion rate k	* 6.9728	0.5052	0.8419	*** 0.0246
	(3.1274)	(0.2650)	(0.8660)	(0.0064)
	[0.3046]			[0.7126]
Frailty mean reversion level z	** 0.0040	0.1347	0.0034	*** 0.3999
	(0.0015)	(0.4778)	(0.0228)	(0.0658)
	[−0.0685]			[−1.9789]
Contagion sensitivity b	*** 0.0084	0.0018	*** 0.0089	0.0004
	(0.0012)	(0.0029)	(0.0012)	(0.0007)
	[−1.3774]		[−1.0332]	
Contagion decay rate κ	*** 0.0508	1.0070	*** 0.0554	0.1727
	(0.0070)	(2.4606)	(0.0068)	(0.2451)
	[−0.7861]		[−0.1375]	

based on a significantly longer sample period that encompasses the financial crisis and also includes financial firms.

The model that controls for observable and latent systematic risk factors fails the Potthoff-Whittinghill-Bohning tests of the moments of the Poisson distribution. It also fails a chi-squared goodness-of-fit test. Table 7 suggest that the model fails these tests because it tends to generate underdispersed binned event counts, with empirical variances of the P_i^w that are much lower than the corresponding empirical means. Underdispersion is a sign of positive correlation between model-implied inter-default times (see Winkelmann, 1995). We conclude that a model that only controls for firms' joint exposure to observable and latent systematic factors does not properly capture the correlation structure of the default times in the data set.

The results of the time-scaling tests support Hypothesis 1: there is excess clustering in the default timing data that cannot be explained by firms' exposure to observable and latent systematic factors. Next we examine whether the contagion channel can explain the excess clustering.

5.2. Hypothesis 2

We test Hypothesis 2 regarding the significance of the contagion channel for default clustering, after controlling for firms' exposure to observable and latent systematic factors. The parameter estimates for the unrestricted model, which are discussed in Section 4.3 above, indicate that the contagion term is statistically and economically significant. A default has a persistent impact on the intensity, with a half-life of about three months. Upon a default event with $173 million of defaulted debt, the average in our data set, the intensity increases by roughly 2.8 events per year.

To further assess the statistical significance of the contagion channel, we use likelihood ratio tests that measure by how much the goodness-of-fit is improved when adding factors to a restricted model. The test results are reported in Table 8. Relative to the base model, adding the frailty term leads to a significant improvement in model fit. This indicates the importance of controlling for latent systematic factors when measuring the degree of clustering due to exposure to systematic factors. Relative to the model that includes observable and latent factors, the inclusion of the self-excitation term representing an additional source of clustering yields further improvements. Overall, these findings provide strong support of Hypothesis 2. Contagion is a significant source of clustering, after controlling for firms' joint exposure to observable and latent systematic factors, both from a statistical and from an economic perspective.

Table 6
Poisson distribution tests.
p-values of Poisson distribution tests for the counts of the binned event counts P_i^w. The Potthoff-Whittinghill-Bohning test compares the empirical moments of the P_i^w to the theoretical moments. Its asymptotic distribution is standard normal (Karlis and Xekalaki (2000)). The Kocherlakota-Kocherlakota test evaluates the empirical moment generating function of the P_i^w and compares to its theoretical counterpart, see Karlis and Xekalaki (2000). The upper tail test examines the fatness of the right tail of the empirical distribution of the P_i^w (see Das et al., 2007). We compute empirical p-values for the Kocherlakota-Kocherlakota and the upper tail tests via bootstrapping with 100,000 samples. The independence test is a BDS test with a distance parameter of one standard deviation. The Fisher dispersion test examines the variance of the realized P_i^ws and is described in Das et al. (2007). Finally, we also perform a standard chi-squared goodness of fit test for the Poisson distribution. The number of samples of P_i^w for each model are as follows (ordered by increasing bin size). Base: 727, 364, 243, 182, 146. Base + frailty: 775, 388, 259, 194, 155. Base + contagion: 729, 365, 243, 183, 146. Complete: 730, 365, 244, 183, 146. * indicates rejection with 95% confidence, ** rejection with 99% confidence, and *** rejection with 99.9% confidence.

Potthoff-Whittinghill-Bohning test				
w	Base	Base + frailty	Base + contagion	Complete
2	*** 0.000	*** 0.000	*** 0.000	*** 0.000
4	*** 0.000	*** 0.000	* 0.012	0.126
6	*** 0.000	*** 0.000	0.253	0.499
8	*** 0.000	** 0.001	0.737	0.717
10	*** 0.000	** 0.002	0.660	0.820

Kocherlakota-Kocherlakota test				
w	Base	Base + frailty	Base + contagion	Complete
2	0.575	0.508	0.688	0.693
4	0.158	0.466	0.719	0.807
6	* 0.015	0.482	0.833	0.916
8	** 0.001	0.515	0.965	0.949
10	*** 0.000	0.489	0.952	0.978

Independence test				
w	Base	Base + frailty	Base + contagion	Complete
2	*** 0.000	0.911	0.961	0.517
4	*** 0.000	0.655	0.148	0.163
6	*** 0.000	* 0.030	0.118	0.064
8	*** 0.000	0.074	0.296	** 0.004
10	*** 0.000	** 0.002	* 0.018	** 0.001

Fisher dispersion test				
w	Base	Base + frailty	Base + contagion	Complete
2	*** 0.000	1.000	1.000	1.000
4	*** 0.000	1.000	0.997	0.949
6	*** 0.000	1.000	0.884	0.771
8	*** 0.000	1.000	0.647	0.657
10	*** 0.000	1.000	0.674	0.592

Upper tail test				
w	Base	Base + frailty	Base + contagion	Complete
2	** 0.004	1.000	0.994	0.985
4	*** 0.000	1.000	0.921	0.796
6	*** 0.000	1.000	0.742	0.630
8	*** 0.000	1.000	0.741	0.672
10	*** 0.000	1.000	0.769	0.632

Chi-squared test				
w	Base	Base + frailty	Base + contagion	Complete
2	*** 0.000	*** 0.000	*** 0.000	*** 0.000
4	*** 0.000	*** 0.000	** 0.008	0.115
6	*** 0.000	*** 0.000	0.248	0.517
8	*** 0.000	*** 0.000	0.763	0.746
10	*** 0.000	*** 0.000	0.683	0.849

5.3. Hypothesis 3

We study whether excess clustering remains in the default data that cannot be explained by firms' exposure to systematic factors and contagion. We begin by testing the unrestricted model using the time-scaling tests introduced in Section 5.1. The results are reported in Table 6. The unrestricted model fails few tests. It exhibits only a minor degree of under or overdispersion, with empirical means and variances of the P_i^w in Table 7 that are closest to each other when compared to the nested submodels.

The time-scaling tests evaluate specific distributional features of time-scaled default events. The unrestricted model may allow for unexplained excess clustering even if it does not fail the tests of Section 5.1 if, for example, it fails at matching distributional features not evaluated by the time-scaling tests. To show that this is not the case, we look at quantile-quantile (QQ) plots of time-scaled default events and inter-default times. Under the null hypothesis of a correctly specified model, Proposition 1 in Appendix B implies that the time-scaled binned default events P_i^w are independent samples from a Poisson distribution with rate w. Proposition 1 also implies that the time-scaled inter-default times $\tau_n = \int_{T_{n-1}}^{T_n} h_s ds$ are independent samples from a standard exponential distribution. If the quantiles of observed time-scaled binned defaults and inter-default times match the corresponding theoretical quantiles within the bounds of statistical significance, then this is an indication that the unrestricted model is adequately specified for the default timing data. Such a result would suggest that all default clustering in the data can be explained by firms' joint exposure to systematic factors and contagion.

Fig. 10 shows the QQ plots. We see that the quantiles of binned default events match the theoretical quantiles for all bin sizes. We also see that the quantiles of inter-default times match the corresponding theoretical quantiles of the standard exponential distribution for a majority of the observations. Put together, the QQ plots suggest that the unrestricted model appropriately captures the default clustering in the data. These findings provide evidence against Hypothesis 3.

6. Implications for portfolio credit risk

We study the implications of our results for the estimation of portfolio credit risk. To this end, we test the ability of our model and its nested alternatives to forecast clustered default risk out of sample. We use Monte Carlo simulation to generate the conditional distribution of the number of defaults in the year ahead, given a model fitted to the data available at the begin-

Table 7
Moments of binned event counts.
This moments of the binned event counts P_i^w for bin sizes $w \in \{2, 4, 6, 8, 10\}$ for each alternative model. We also report the moments of the theoretical Poisson distributions with rate w.

$w = 2$					
	Observations	Mean	Variance	Skewness	Kurtosis
Theoretical		2.000	2.000	0.707	3.500
Base model	727	1.997	2.617	1.134	6.161
Base + contagion model	729	1.992	1.585	0.565	3.268
Base + frailty model	775	1.874	1.302	0.613	3.606
Complete model	730	1.989	1.597	0.611	3.206
$w = 4$					
	Observations	Mean	Variance	Skewness	Kurtosis
Theoretical		4.000	4.000	0.500	3.250
Base model	364	3.989	6.997	0.846	4.581
Base + contagion model	365	3.978	3.247	0.391	2.989
Base + frailty model	388	3.742	2.466	0.393	2.999
Complete model	365	3.978	3.527	0.592	3.301
$w = 6$					
	Observations	Mean	Variance	Skewness	Kurtosis
Theoretical		6.000	6.000	0.408	3.167
Base model	243	5.975	13.983	0.800	4.274
Base + contagion model	243	5.975	5.371	0.343	2.854
Base + frailty model	259	5.606	3.720	0.309	3.203
Complete model	244	5.951	5.586	0.232	2.690
$w = 8$					
	Observations	Mean	Variance	Skewness	Kurtosis
Theoretical		8.000	8.000	0.354	3.125
Base model	182	7.978	21.900	0.570	3.618
Base + contagion model	183	7.934	7.655	0.206	3.055
Base + frailty model	194	7.485	4.987	0.069	2.572
Complete model	183	7.934	7.633	0.236	2.488
$w = 10$					
	Observations	Mean	Variance	Skewness	Kurtosis
Theoretical		10.000	10.000	0.316	3.100
Base model	146	9.945	33.087	0.502	3.353
Base + contagion model	146	9.945	9.431	0.095	3.093
Base + frailty model	155	9.368	6.104	0.168	2.876
Complete model	146	9.945	9.680	0.130	2.549

ning of the year. We then run a series of forecast accuracy tests. Our approach, detailed in Internet Appendix C, eliminates the need to simulate the paths of the frailty variable during the forecast period. It is computationally efficient.

The validity of Hypothesis 2 tells us that the unrestricted model should outperform all nested model alternatives in predicting clusters of defaults out of sample. Moreover, the rejection of Hypothesis 3 suggests that the unrestricted model, which allows for only an insignificant degree of unexplained excess clustering, should generate forecasts of clustered default risk that are matched by the data. It should only be rejected by few forecast accuracy tests. In contrast, the validity of Hypothesis 1 implies that the nested models that control only for the influence of systematic factors (observable or latent) should generate biased forecasts of clustered default risk out of sample. These nested model alternatives should be rejected by our forecast accuracy tests.

Fig. 11 contrasts the forecast distribution with the realized number of defaults, for each year between 1991 and 2012, and for each alternative model. The base model, which only considers observable systematic factors, lags default clusters. It fails to predict the clusters associated with the burst of the dotcom bubble and the financial crisis. The model including frailty but ignoring contagion performs better at predicting these clustering episodes. However, it generates forecast distributions with excessively heavy tails. Moreover, the means of the forecast distributions show very little variation through the test period. Thus, if the mean were used as a point forecast, then the forecast would change little every year, despite the significant variation of observed default rates. The model that includes the contagion channel but ignores frailty generates much more concentrated forecast distributions. The forecast mean matches the realized number of defaults quite well. The unrestricted model performs better still, produc-

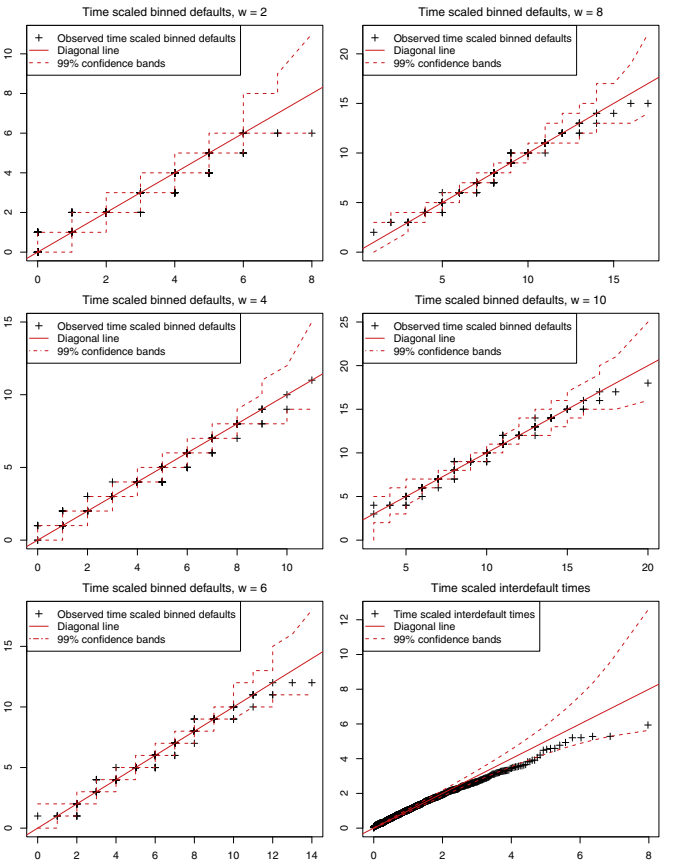


Fig. 10. QQ plots.
This figure compares the quantiles of observed time-scaled binned defaults events P_i^w for $w \in \{2, 4, 6, 8, 10\}$, as well as the quantiles of observed time-scaled inter-default times $\tau_n = \int_{T_{n-1}}^{T_n} h_s \mathrm{d}s$, to the corresponding theoretical quantiles under the null hypothesis of correct specification of the unrestricted model. Under this null hypothesis, Proposition 1 in Appendix B implies that P_i^w are i.i.d. samples of the Poisson distribution with rate w, and τ_n are i.i.d. samples of the standard exponential distribution. Confidence bands are computed via bootstrap with 10^5 bootstrap samples.

ing forecast distributions with lower variance during the first third of the test period.

To test forecast accuracy, we consider the 99% value at risk (VaR) of the forecast distribution, a standard measure of portfolio credit risk. For a given model, we evaluate the sequence of hit indicators associated with violations of the VaR in different periods. Under the null, these indicators are independent draws from a Bernoulli distribution with success probability $(1 - -0.99) = 0.01$. Table 9 reports the violation rates and the p-values of several independence and Bernoulli distribution tests. The base model is rejected by these tests due its sluggishness. The model including the contagion channel but ignoring frailty is rejected because its VaR forecasts are too low.

The tests of the hit indicators may unduly favor the frailty model because it generates forecast distributions with excessively heavy tails and VaR forecasts that are excessively large. If one were to estimate risk capital according to the VaR, one would end up holding too much capital relative to the actual risk. While one would be covered in most situations, scarce risk capital would be wasted. To address this issue, we consider the mean relative bias and the root mean square relative bias of the VaR forecasts (see Hendricks, 1996). The former is the percentage deviation of a forecast from the average forecast of all models, averaged over all forecast periods. The latter accounts for the standard deviation of the former. The outcomes of these measures, reported in Table 9, indicate that the frailty model indeed tends to overstate the VaR and produce the largest forecast volatility. The unrestricted model generates VaR forecasts with the least amount of bias and volatility.

All things considered, our out-of-sample analysis corroborates the results of our hypotheses tests. We find that a model that captures default clustering caused by firms' joint exposure to observable and latent systematic factors, as well as contagion, provides balanced and accurate out-of-sample forecasts of correlated default risk. Extending the results of Das et al. (2007), Duffie et al. (2009), and Lando and Nielsen (2010), our findings indicate that failure to control for both latent factors, as well as contagion, introduces bias into forecasts of portfolio credit risk. A model that ignores the contagion channel tends to significantly overstate correlated default risk, while a model that ig-

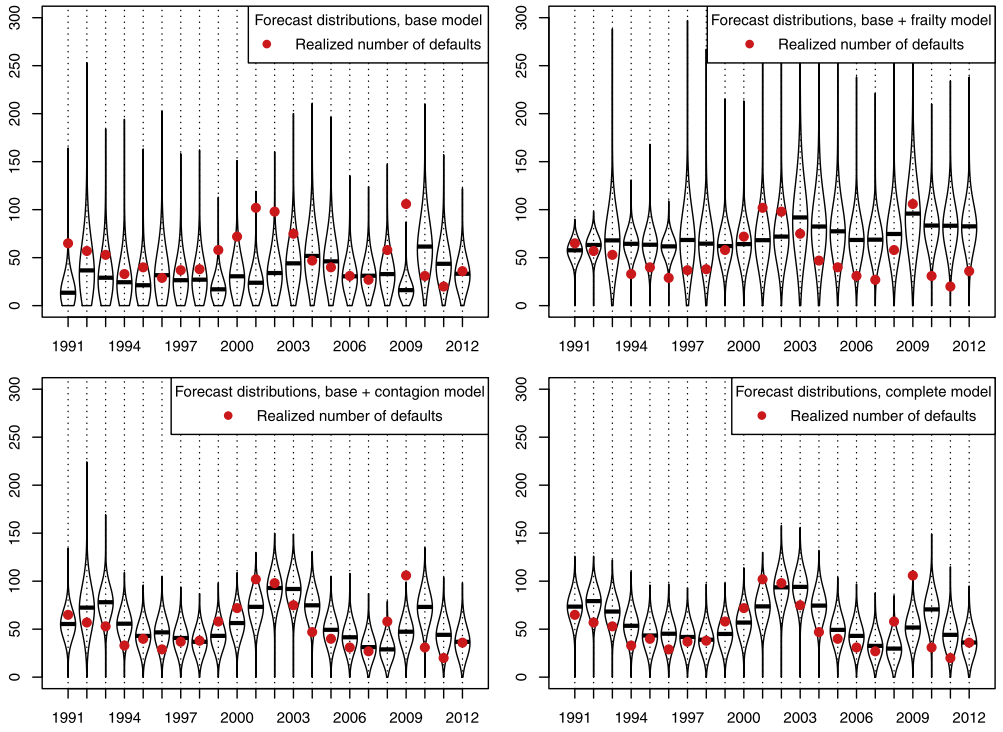


Fig. 11. Out-of-sample forecast distribution of the number of defaults in the year ahead.
The plot shows the Gaussian kernel-smoothed forecast distributions, with a horizontal line indicating the mean. The dot indicates the realized number of defaults. The forecast distributions are computed via Monte Carlo simulation as described in Internet Appendix C.

Table 8
Likelihood ratio tests.
This table reports likelihood ratio test statistics, *p*-values, and degrees of freedom for the corresponding asymptotic distributions. A likelihood ratio test evaluates the fit of an alternative model relative to a nested benchmark model. The test statistic is given by twice the difference between the maximum log-likelihood of the alternative model and the benchmark model. The log-likelihood values are reported in Table 4. The test statistic has, asymptotically, a chi-squared distribution with degrees of freedom equal to the number of additional parameters included in the alternative. *** indicates significance at the 99.9% level.

Benchmark model	Base	Base	Base	Base + frailty
Alternative model	Base + contagion	Base + frailty	Complete	Complete
Test statistic	320.18	299.60	320.34	14.10
Degrees of freedom	2	3	5	2
p-value	*** 0.000	*** 0.000	*** 0.000	*** 0.001

nores firms' exposure to latent systematic factors tends to significantly understate it.[13]

7. Pricing implications

We examine the implications of our results for the pricing of clustered default risk in corporate credit markets. Jorion and Zhang (2007) empirically document spillovers of bankruptcies on intra-industry CDS spreads. Berndt et al. (2010) calibrate to bond prices a model of clustered default risk that allows for contagion. They find that the default of one firm can have a sizable impact on credit spreads of other firms. These papers provide important evidence of contagion effects under the pricing measure. A theoretical explanation for the findings of Jorion and Zhang (2007) and Berndt et al. (2010) is provided by Bai et al. (2015). These authors study production-based general equilibria for large economies and show that in the presence of a contagion channel, exposure to contagion demands a large risk premium. Our results empirically support the model of Bai et al. (2015). We find evidence of a contagion channel under the physical probability measure.

[13] We attribute the underperformance of the model that ignores frailty to the underdispersion for small bin sizes detected by our time-scaling tests; see Tables 6 and 7. Underdispersion is a sign of positive correlation between model-implied inter-default times (see Winkelmann, 1995). Therefore, underdispersion for small bin sizes suggests that long model-implied inter-default times tend to be followed by long model-implied inter-default times, leading to an underestimation of clustered default risk.

Table 9
Out-of-sample forecast accuracy tests.
The table reports *p*-values of various tests of forecast accuracy and other measures of accuracy. The violation rate indicates the number of times that the realized number of defaults in a given year exceeds the forecast value-at-risk. The null hypothesis of the unconditional coverage test of Kupiec (1995) is that the violation rate does not exceed 1%. The null hypothesis of the independence test of Christoffersen (1998) is that a binary first-order Markov chain for the hit indicators has transition matrix given by the identity matrix. The asymptotic distribution of these tests is chi-squared with one degree of freedom. The Markov test combines the coverage and independence tests. It takes the null of the coverage test and the alternative of the independence test (see Christoffersen, 1998). The asymptotic distribution of this test is chi-squared with two degrees of freedom. The null hypothesis of the dynamic quantile test due to Engle and Manganelli (2004) is that there is no correlation between the hit indicators and the number of defaults per year. The asymptotic distribution is chi-squared with three degrees of freedom. Finally, the mean relative bias is the average relative difference of the value-at-risk predicted by each model, compared to the average prediction of all models. The root mean squared relative bias is the standard deviation of the latter. * indicates rejection with 95% confidence, *** indicates rejection with 99% confidence, and *** rejection with 99.9% confidence.

Model	Base	Base + frailty	Base + contagion	Complete
Violation rate	5/22	0/22	2/22	1/22
Unconditional coverage test	*** 0.000	0.506	* 0.020	0.221
Independence test	0.951	1.000	0.107	0.752
Markov test	*** 0.000	0.802	* 0.018	0.450
Dynamic Quantile Test	* 0.020	1.000	0.318	0.972
Mean relative bias	−11.98%	26.44%	−8.13%	−6.32%
Root mean squared relative bias	17.31%	33.30%	14.37%	13.77%

We empirically validate the findings of Bai et al. (2015). We run random coefficient panel regressions of changes in corporate bond spreads on changes in the different components of our model: a factor $e^{a_0+\sum_{i=1}^{d} a_i X_{i,t}}$ that captures exposure to observable systematic factors, a filtered frailty factor $\mathbb{E}_t[Z_t]$ that captures exposure to latent systematic factors, and a factor Y_t that captures exposure to the contagion channel. If exposure to contagion risk indeed demands a large risk premium, we should see significantly positive regression coefficients on our contagion factor. This should hold even when we control for the component of credit spreads that is attributed to default probabilities. Recent evidence by Hilscher and Wilson (2017) suggests that high-yield bonds carry higher risk premiums during periods of high default probabilities. Based on this evidence and the fact that the contagion factor spikes up during periods of default clustering (see Fig. 9), we expect that changes in the contagion factor Y_t will be particularly important for explaining credit spreads changes of high-yield bonds. We should see significant variation across rating classes in the sensitivity of credit spread changes to changes in our contagion factor.

Table 10 reports the results of the regressions, in which we also include a series of control variables as well as a nonparametric measure of the default rate that serves as a control for default probabilities.[14] We find that, on average, changes in US corporate bond spreads react significantly and positively to changes in our contagion factor. Across rating classes, we see that the coefficients on the contagion factor are positive whenever they are significant. Consistent with Hilscher and Wilson (2017), we find that changes in our contagion factor explain a significant fraction of the variation of credit spreads of some high-yield bonds. An *F*-test of whether the sensitivity to changes in our contagion factor is constant across rating classes has a *p*-value of 0.073, suggesting that there is some variation across rating classes. The results of the regressions provide empirical support for the findings of Bai et al. (2015).

Next, we analyze the pricing of clustered default risk in collateralized debt obligation (CDO) markets. Coval et al. (2009) study CDO tranches referenced on the CDX North American Investment Grade Index between 2004 and 2007. They find that those CDO tranches were significantly overpriced, suggesting investors received little compensation for bearing clustered default risk. Collin-Dufresne et al. (2012), however, fitting an extension of the model of Coval et al. (2009) with catastrophic systematic jumps to CDX term structure data, cannot find evidence of overpricing. Using a different modeling approach, Azizpour et al. (2011) consider tranches referenced on both the Investment Grade and High Yield CDX Indices and document substantial risk premiums.

Going beyond these papers, our approach enables us to analyze the compensation CDX investors receive for bearing exposure to different components of clustered default risk. We regress changes in CDX index and tranche spreads on changes in the different components of the unrestricted model as well as relevant control variables. Consistent with the evidence in favor of Hypotheses 2 and 3, we should see significantly positive regression coefficients on some of these factors after controlling for aggregate default probabilities if clustered default risk is priced and carries a risk premium. Given that Collin-Dufresne et al. (2012) require catastrophic market jumps to match senior CDX tranche

[14] We construct a weekly nonparametric measure of the default rate by measuring the ratio of days within a week for which at least one default observation is recorded in our data.

Table 10
Random coefficient panel regressions of weekly changes in corporate credit spreads on weekly changes of our three default clustering factors. We run these regression within each rating class, as well as on average across rating classes. "Macro" stands for the factor $e^{a_0+\sum_{i=1}^{d} a_i X_{i,t}}$, "Contagion" stands for the factor Y_t, and "Frailty" stands for the filtered frailty $\mathbb{E}_t[Z_t]$. We use the parameter estimates of Table 4 to construct the time series of the contagion, frailty, and macro factors. We include a series of control variates taken from Collin-Dufresne et al. (2001). "Level" denotes the ten-year Treasury contant maturity rate, while "Slope" denotes the difference between the ten-year and the two-year Treasury constant maturity rates. "Leverage" stands for the average leverage ratio in a rating class, where we rate US firms according to their S&P Domestic Long Term Issuer Credit Rating (Compustat item SPLTICRM). Leverage ratios are computed as in Collin-Dufresne et al. (2001). We also control for changes in a nonparametric measure of the aggregate default rate computed as the ratio of days of a week with at least one default observation in our data sample. Weekly changes are measured from Wednesday to Wednesday. Standard errors are computed following Swamy (1970). · indicates significance at the 90% level, * indicates significance at the 95% level, ** significance at the 99% level, and *** significance at the 99.9% level.

	Average	AAA	AA	A	BBB	BB	B
Constant	0.167	*** 1.002	0.001	0.001	0.000	0.002	−0.004
	(0.167)	(0.002)	(0.003)	(0.003)	(0.003)	(0.006)	(0.008)
Default rate	* 0.033	0.013	0.022	0.007	* 0.039	0.022	· 0.095
	(0.013)	(0.013)	(0.016)	(0.017)	(0.020)	(0.037)	(0.050)
Macro	** 0.018	*** 0.011	** 0.010	** 0.009	** 0.011	* 0.015	*** 0.049
	(0.006)	(0.003)	(0.004)	(0.003)	(0.004)	(0.007)	(0.010)
Contagion	** 0.003	0.000	0.002	· 0.003	* 0.004	* 0.007	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)	(0.003)	(0.004)
Frailty	0.033	−0.155	−0.318	−0.285	−0.107	0.248	0.838
	(0.180)	(0.196)	(0.240)	(0.256)	(0.291)	(0.545)	(0.751)
Leverage	· 6.086	−0.109	** −1.851	2.110	** 3.901	*** 14.117	*** 19.235
	(3.460)	(0.177)	(0.682)	(1.305)	(1.241)	(1.624)	(1.687)
S&P 500	−0.259	** −0.447	*** −0.800	−0.209	−0.283	0.570	−0.254
	(0.184)	(0.138)	(0.191)	(0.269)	(0.303)	(0.523)	(0.663)
VIX	*** 0.003	0.001	* 0.003	** 0.005	0.002	0.004	0.002
	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.003)	(0.004)
Level	* −0.339	*** −0.072	*** −0.091	** −0.084	*** −0.137	*** −0.632	*** −1.029
	(0.163)	(0.020)	(0.024)	(0.026)	(0.030)	(0.056)	(0.077)
Slope	* 0.206	* 0.070	** 0.092	* 0.085	* 0.101	*** 0.305	*** 0.587
	(0.084)	(0.028)	(0.034)	(0.036)	(0.041)	(0.077)	(0.107)
R^2	0.042	0.109	0.172	0.157	0.184	0.429	0.527
Number of obs.	783	783	783	783	783	783	783
Start date	1997-01-02	1997-01-02	1997-01-02	1997-01-02	1997-01-02	1997-01-02	1997-01-02
End date	2012-12-20	2012-12-20	2012-12-20	2012-12-20	2012-12-20	2012-12-20	2012-12-20

spreads, and given that Fig. 9 shows that the contagion factor spikes up during periods of default clustering, we expect to see that changes in the contagion factor Y_t are particularly important for explaining changes in senior CDX tranche spreads.

Table 11 shows the results of ordinary least squares regression for the five-year and ten-year investment-grade CDX index spreads, and of random coefficient panel regressions for the five-year investment-grade CDX tranche spreads. We find that a part of the CDX index spread changes can be explained by changes in the exposure to observable systematic factors.[15] We also see that, on average, CDX tranche spreads react significantly and positively to changes in the contagion factor. Changes in the spreads of some senior CDX tranches tend to be sensitive to changes in the contagion factor. This observation is validated by a Z-test of whether the sensitivity to changes of in our contagion factor is higher for senior CDX tranches (capturing more than 3% of the losses) than for junior CDX tranches (capturing up to 3% of the losses). Such a Z-test has a p-value of 0.026. We cannot attribute any spread changes to latent systematic factors as measured by the filtered frailty.

Consistent with Azizpour et al. (2011) and Collin-Dufresne et al. (2012), our results indicate that CDX investors indeed require risk premiums for bearing clustered default risk. Extending Collin-Dufresne et al. (2012), we show that part of the risk premiums for senior CDX tranches can be attributed to contagion risk. Overall, the results of this section suggest that the pricing of clustered default risk is nuanced; different clustering sources give rise to risk premiums that differ across markets, rating classes, and seniority.

8. Robustness checks

8.1. Firm-specific variables

The literature has identified several firm-specific risk factors that are significant predictors of corporate default probabilities; see Das et al. (2007) and Lando and Nielsen (2010), and the references therein. These factors include firm size, firm stock returns and volatilities, as well as distance to default (a volatility-adjusted measure of leverage). One could argue that our results might be biased because we do not control for the influence of firm-specific risk factors. If variations in the cross-section of firm-specific risk

[15] We attribute the insignificance of changes in the macro factor in the regression of changes in the five-year CDX index to the short data time span.

Table 11
Regressions of weekly changes in CDX index and tranche spreads for different rating classes on weekly changes of our three default clustering factors. We run ordinary least squares regressions for the five-year and ten-year CDX indexes, and random coefficient panel regressions for CDX tranches. For the ten-year horizon we employ the Markit North America Investment Grade index (Bloomberg item IBOXUMAJ). For the five-year horizon we employ the five-year Dow Jones North American Investment Grade index (Bloomberg item DJCDXNI). "Macro" stands for the factor $e^{a_0+\sum_{i=1}^{d} a_i X_{i,t}}$, "Contagion" stands for the factor Y_t, and "Frailty" stands for the filtered frailty $\mathbb{E}_t[Z_t]$. We include several control variates from Coval et al. (2009). We use the yield on the three-month Treasury bill as a proxy for the risk-free rate. "Swap" denotes the five-year swap rate. "Bid ask" is the bid-ask spread of the corresponding security. We also control for a nonparametric measure of the aggregate default rate computed as the ratio of days of a week with at least one default observation in our data sample. Weekly changes are measured from Wednesday to Wednesday. · indicates significance at the 90% level, * indicates significance at the 95% level, ** significance at the 99% level, and *** significance at the 99.9% level.

	10-year CDX	5-year CDX	Average	Tranches on 5-year CDX 0%–3%	3%–7%	7%–10%	10%–15%
Constant	0.001	· 0.004	* 0.008	0.003	0.018	0.007	0.004
	(0.004)	(0.002)	(0.004)	(0.002)	(0.013)	(0.005)	(0.002)
Default rate	0.030	0.018	0.019	0.016	0.071	−0.014	0.003
	(0.024)	(0.017)	(0.018)	(0.015)	(0.092)	(0.035)	(0.017)
Macro	* 0.016	−0.007	* −0.014	· −0.009	−0.030	−0.015	−0.004
	(0.007)	(0.005)	(0.006)	(0.005)	(0.029)	(0.011)	(0.005)
Contagion	0.000	0.001	· 0.004	0.001	· 0.012	· 0.004	0.002
	(0.002)	(0.001)	(0.002)	(0.001)	(0.007)	(0.002)	(0.001)
Frailty	0.084	−0.124	−0.118	−0.038	−0.025	−0.428	−0.043
	(0.350)	(0.222)	(0.098)	(0.194)	(1.231)	(0.462)	(0.223)
Risk-free rate	*** −0.120	* −0.089	· −0.151	−0.048	· −0.406	−0.127	−0.049
	(0.024)	(0.043)	(0.085)	(0.038)	(0.238)	(0.089)	(0.043)
S&P 500	*** −1.014	−0.448	* −1.760	· −0.455	** −4.156	** −1.879	** −0.877
	(0.277)	(0.281)	(0.827)	(0.245)	(1.554)	(0.579)	(0.280)
VIX	*** 0.009	0.004	0.001	0.003	0.006	−0.004	−0.001
	(0.002)	(0.002)	(0.002)	(0.002)	(0.014)	(0.005)	(0.002)
Swap	−0.008	−0.007	** −0.048	−0.014	−0.083	−0.064	−0.038
	(0.027)	(0.022)	(0.015)	(0.019)	(0.123)	(0.046)	(0.022)
Bid ask		*** 0.010	** 1.087	** 0.459	*** 2.070	** 1.373	** 0.632
		(0.003)	(0.370)	(0.154)	(0.623)	(0.447)	(0.216)
R^2	0.409	0.338	0.201	0.295	0.315	0.304	0.314
Number of obs.	289	132	132	132	132	132	132
Start date	2003-10-30	2004-08-05	2004-08-05	2004-08-05	2004-08-05	2004-08-05	2004-08-05
End date	2011-03-24	2007-07-26	2007-07-26	2007-07-26	2007-07-26	2007-07-26	2007-07-26

factors contain information about systematic default risks that are not captured by the systematic risk factors we include in our model (see Table 1), then firm-specific risk factors may be important for default clustering. To show that this is not the case, we test whether the cross-section of omitted firm-specific default risk factors contains information not captured by our systematic variables.

For the companies in our default data set, we construct monthly time series of cross-sectional averages of significant firm-specific default risk factors.[16] Table 12 provides summary statistics. The firm-specific variables we consider are the quick ratio (QUICK),[17] ratio of short term debt to total debt (SD), total liabilities (TL), distance to default (DTD), log book asset value (LBA), net income to total assets ratio (NITA), total liabilities to total assets ratio (TLTA), log relative market capitalization with respect to the total market capitalization of all NYSE/Nasdaq/Amex listed firms (RSIZE), one-year trailing stock return (RET), and the realized standard deviation of monthly returns over rolling horizons of 12 months (VOL). Next, we regress the cross-sectional averages on our macroeconomic variables. Table 13 gives a summary of the regression results. We see that our macroeconomic variables can capture at least 52% of the monthly variation of cross-sectional averages of firm-specific default risk factors. Surprisingly, our macroeconomic variables capture 71% of the monthly fluctuations of average distance to default in our data set. Duffie et al. (2007, 2009), Lando and Nielsen (2010), and many others find that distance to default is the most economically and statistically significant determinant of US corporate default probabilities. Yet, our regression analysis states that the cross-sectional average of distance to default is almost perfectly spanned by our macroeconomic variables.

Next we regress the frailty and self-exciting terms of the intensity $(Y_t + Z_t)$ on the residuals of the regressions in Table 13. We find that the residuals can only explain about 22% of the variation of $Y_t + Z_t$ estimated for the complete model over the sample period. We conclude that while firm-specific default risk factors may contain significant information about idiosyncratic default risk, they only contain little information about systematic default risk not captured by our macroeconomic variables. Thus, the absence of firm-specific variables in our model is unlikely to bias our results.

[16] Proposition 1 of Duffie (1998) implies that the only relevant firm-specific data for modeling the conditional default rate in our sample are the cross-sectional averages of firm-specific default risk factors after censoring defaulted firms.

[17] The quick ratio is the ratio of a firm's cash, short-term investments, and total receivables over its current liabilities.

Table 12
Summary statistics of firm-specific default risk factors.
This table provides summary statistics for the time series of firm-specific default risk factors. Except for distance to default, we construct time series from accounting and stock market data from the joint Compustat/CRSP database quarterly file when available, and otherwise from the annually file. If we were to construct our own distance to default time series using Compustat/CRSP data, we would have to neglect about two-thirds of our default events due to data unavailability. To circumvent this issue, we obtain distance to default data from the Credit Research Initiative (CRI) at the National University of Singapore, which is only available starting in the year 1991. The moments and quantiles are computed from the cross-section of firms in our default data set with available data on CRSP/Compustat/CRI. We exclude the 5% cross-sectional extreme values to ensure continuity of our time series.

Variable	Sample frequency	Mean	Standard deviation	Skewness	Kurtosis	Min.	First quartile	Median	Third quartile	Max.	Number of obs. across firms
QUICK	Monthly	0.92	0.58	0.84	3.55	0.00	0.50	0.84	1.23	2.89	95739
SD	Monthly	0.18	0.23	1.66	5.06	0.00	0.02	0.09	0.25	0.96	101739
TL	Monthly	628.71	951.41	2.63	10.50	0.00	75.20	260.65	708.35	5851.77	106048
DTD	Monthly	3.16	2.27	0.72	3.76	−2.20	1.50	2.99	4.43	13.06	8751
LBA	Monthly	4.68	1.58	−0.04	2.47	1.00	3.56	4.72	5.81	8.51	83906
NITA	Monthly	0.00	0.03	−2.03	9.89	−0.20	−0.01	0.00	0.02	0.08	106590
TLTA	Monthly	0.72	0.26	0.28	3.20	0.00	0.54	0.70	0.89	1.52	106036
RSIZE	Monthly	−10.05	2.69	−1.27	4.74	−19.74	−11.13	−9.52	−8.27	−5.64	80636
RET	Monthly	−1.10	5.62	−1.19	5.10	−24.23	−3.39	0.00	2.36	10.44	75531
VOL	Monthly	16.26	10.76	1.68	5.95	0.00	9.18	13.15	19.71	60.76	75529

Table 13
Coefficient estimates and standard errors of ordinary least squares regressions of firm specific default risk factors.
This table presents the results of ordinary least squares regressions of the cross-sectional averages of key firm-specific default risk factors on our macroeconomic variables. We censor firms from the cross-sectional averages after a default. * indicates significance at the 95% level, ** significance at the 99% level, and *** significance at the 99.9% level.

Variable	QUICK	SD	TL	DTD	LBA	NITA	TLTA	RSIZE	RET	VOL
Constant	* 0.021	0.018	0.008	* 0.004	0.002	0.016	0.014	−0.004	0.003	0.001
	(0.010)	(0.011)	(0.011)	(0.002)	(0.009)	(0.011)	(0.011)	(0.006)	(0.006)	(0.006)
IP	−0.001	−0.001	−0.001	0.000	−0.000	−0.001	−0.001	0.000	0.000	0.000
	(0.000)	(0.001)	(0.001)	(0.000)	(0.000)	(0.001)	(0.001)	(0.000)	(0.000)	(0.000)
GDP	* −0.001	** −0.001	* −0.001	** 0.000	−0.001	*** −0.002	** −0.001	0.000	0.000	0.000
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)	(0.000)
RET	0.000	0.000	0.000	*** 0.000	0.000	0.0000	0.0000	0.0000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
VOL	** 0.001	** 0.001	*** 0.001	*** 0.000	*** 0.001	*** 0.001	*** 0.001	*** 0.001	*** 0.000	** 0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
LEVEL	−0.001	−0.003	−0.002	−0.001	−0.004	−0.003	−0.002	−0.003	−0.003	−0.003
	(0.004)	(0.005)	(0.005)	(0.001)	(0.004)	(0.005)	(0.005)	(0.003)	(0.003)	(0.003)
SLOPE	0.007	* 0.011	0.008	*** 0.003	* 0.010	0.009	0.009	* 0.007	* 0.007	* 0.007
	(0.005)	(0.005)	(0.006)	(0.001)	(0.004)	(0.006)	(0.006)	(0.003)	(0.003)	(0.003)
10Y	*** −0.025	*** −0.029	*** −0.029	0.000	*** −0.016	*** −0.030	*** −0.030	** −0.010	** −0.008	** −0.009
	(0.005)	(0.006)	(0.006)	(0.001)	(0.004)	(0.006)	(0.006)	(0.003)	(0.003)	(0.003)
CRED	*** −0.060	*** −0.064	*** −0.068	*** −0.002	*** −0.048	*** −0.071	*** −0.067	*** −0.033	*** −0.031	*** −0.031
	(0.004)	(0.005)	(0.005)	(0.001)	(0.004)	(0.005)	(0.005)	(0.003)	(0.003)	(0.003)
AAA	*** 0.044	*** 0.051	*** 0.051	*** 0.002	*** 0.037	*** 0.054	*** 0.052	*** 0.027	*** 0.023	*** 0.025
	(0.005)	(0.005)	(0.005)	(0.001)	(0.004)	(0.005)	(0.005)	(0.003)	(0.003)	(0.003)
REC	*** −0.022	*** −0.026	*** −0.025	* 0.002	*** −0.016	*** −0.026	*** −0.026	** −0.008	* −0.007	* −0.006
	(0.005)	(0.005)	(0.005)	(0.001)	(0.004)	(0.005)	(0.005)	(0.003)	(0.003)	(0.003)
Adjusted R^2	0.516	0.522	0.550	0.705	0.602	0.536	0.539	0.656	0.652	0.660
Data points	516	516	516	264	516	516	516	516	516	516

8.2. Economy-wide versus firm-level default rates

Rather than modeling firm-by-firm default rates, our analysis focuses on the economy-wide default rate that describes events in the pool of Moody's rated issuers. To show that this approach is appropriate for measuring correlated default risk and does not bias our results, we compare the goodness-of-fit of our simple base model to the goodness-of-fit of the leading firm-by-firm reduced-form models of Das et al. (2007) and Lando and Nielsen (2010). Das et al. (2007) study 495 US non-financial default events during the time period January 1979 to October 2004. A firm's default intensity depends on the distance to default, the firm's trailing one-year stock return, the three-month Treasury bill rate, and the one-year trailing return of the S&P 500 index. Lando and Nielsen (2010) analyze 370 US nonfinancial default events between January 1982 and December 2005. A firm's default intensity depends on the one-year return on the S&P 500 index, the three-month US Treasury rate, one-year percentage change in US industrial production, the spread between the ten-year and one-year Treasury rate, a firm's one-year equity return, one-year distance-to-default, quick ratio, ratio of short-term debt to total debt, and log book value.

Although our base model does not include firm-specific factors, it fits the data at least as well as the firm-by-firm models of Das et al. (2007) and Lando and Nielsen (2010). To show this, we reestimate our base model from a sub-

Table 14
Comparison to Lando and Nielsen (2010). *p*-values of Poisson distribution tests for the base model and Models I and II of Lando and Nielsen (2010). We compute tests statistics and *p*-values as indicated in Lando and Nielsen (2010) after fitting the base model to the data subsample that mirrors the data used by Lando and Nielsen (2010). We copy *p*-values for Models I and II from Table 4 of Lando and Nielsen (2010). The number of samples for the base model are as follows (ordered by increasing bin size): 407, 204, 136, 102, 82. * indicates rejection with 95% confidence, ** rejection with 99% confidence, and *** rejection with 99.9% confidence.

Fisher dispersion test				Serial correlation test 1			
w	Base	Model I	Model II	*w*	Base	Model I	Model II
2	0.981	0.110	0.324	2	0.324	* 0.047	0.289
4	0.775	* 0.021	0.377	4	0.059	* 0.011	0.392
6	0.278	* 0.013	0.406	6	* 0.022	** 0.004	0.722
8	0.111	** 0.002	0.113	8	** 0.010	** 0.003	0.170
10	* 0.038	** 0.004	* 0.037	10	** 0.008	** 0.002	0.177
Upper tail test (mean)				**Upper tail test (median)**			
w	Base	Model I	Model II	*w*	Base	Model I	Model II
2	0.891	0.476	0.428	2	1.000	1.000	0.728
4	0.523	0.211	0.447	4	0.275	0.972	0.972
6	0.268	0.062	0.501	6	* 0.031	0.136	0.734
8	0.138	* 0.012	0.233	8	0.369	0.260	0.257
10	0.224	* 0.013	0.120	10	0.555	* 0.025	0.566
Potthoff-Whittinghill-Bohning test				**Kocherlakota-Kocherlakota test**			
w	Base	Model I	Model II	*w*	Base	Model I	Model II
2	* 0.022	0.156	0.582	2	* 0.023	0.152	0.558
4	0.443	0.023	0.625	4	0.353	* 0.028	0.768
6	0.513	* 0.012	0.696	6	0.656	* 0.012	0.787
8	0.189	** 0.001	0.154	8	0.305	** 0.002	0.192
10	0.051	** 0.004	0.051	10	0.071	** 0.006	0.112

sample of our data that mirrors the data sets used by Das et al. (2007) and Lando and Nielsen (2010). Using the same time-change tests as these authors (see Section 5 above), Table 14 indicates that our base model performs slightly better than Model I and similarly to Model II of Lando and Nielsen (2010). Further, Table 15 indicates that our base model slightly outperforms the model of Das et al. (2007). Thus, our base model of aggregate defaults performs as well as alternative firm-by-firm models of defaults that have been proposed in the literature. The choice of our modeling approach does not bias our results.

8.3. Financial firms

Defaults of financial firms are special as they tend to be relatively rare, significant, and surrounded by government intervention. Among the 2001 distinct default events in our data set, there are only 126 default events by financial firms. The mean defaulted debt is $240 million, which is larger than the mean defaulted debt for nonfinancial firms ($169 million). Further, all defaults of the type "Seized by regulators" in our data set are by financial firms. Despite the peculiarity of financial firm defaults, empirical evidence on the impact of financial firm defaults on aggregate default rates is sparse. Giesecke et al. (2014) find evidence that, over the long term, the impact of financial default crises on the real economy is far more severe than the impact of nonfinancial default crises. This evidence is consistent with the financial accelerator mechanism of Bernanke et al. (1999), which hypothesizes that distress of financial firms may amplify and propagate shocks to the real economy. Still, Helwege (2010) argues that financial firm distress may not significantly spread to the real economy given that financial firms are able to endure large shocks because they are well regulated and diversified. Further, Helwege and Zhang (2016) provide empirical evidence indicating that historically, financial firm defaults have only triggered modest contagion effects.

We test if our findings are sensitive to the exclusion of financial firms. Table 16 provides the maximum likelihood parameter estimates of the complete model and the three nested alternatives after excluding financial firm defaults from our data. We find that the parameter estimates are essentially the same as those for all (financial and nonfinancial) firms; compare with the estimates in Table 4. We conclude that the inclusion of financial firms is not a key driver of our results.

8.4. Days with multiple defaults

Our data set includes days with multiple defaults. For example, the default of Texas International Air on 9/24/1983 was accompanied by the defaults of the Texas International Air Finance and Texas International Air Capital subsidiaries on the same date. The formulation of Eq. (4) aggregates these defaults when modeling their impact on the default intensity. Alternatively, we could treat each default separately. For example, if there are J_n defaults on day T_n, and u_{n_j} denotes the face value of defaulted debt of

Table 15
Comparison to Das et al. (2007).
Moments and p-values of different tests of the binned defaults P_i^w and times t_i^w for the base model and the model of Das et al. (2007). We compute tests statistics and p-values as in Das et al. (2007) after fitting the base model to the data subsample that mirrors the data used by Das et al. (2007). We copy the p-values from Section III of Das et al. (2007). The number of samples are as follows (ordered by increasing bin size). Base model: 499, 250, 167, 125, 100. Das et al. (2007): 230, 116, 77, 58, 46. Our samples are larger than the ones in Das et al. (2007) because we do not need to eliminate data points for which no data is available on the CRSP data set given that we do not include firm-specific variables. * indicates rejection with 95% confidence, ** rejection with 99% confidence, and *** rejection with 99.9% confidence.

Panel A: Poisson distribution tests

Fisher dispersion test			Upper tail test (mean)			Upper tail test (median)		
w	Base	DDKS	w	Base	DDKS	w	Base	DDKS
2	0.982	*** 0.000	2	0.781	*** 0.000	2	*** 0.000	*** 0.000
4	0.970	*** 0.001	4	0.728	*** 0.000	4	0.274	*** 0.000
6	0.805	*** 0.000	6	0.636	* 0.020	6	0.056	0.060
8	0.447	*** 0.000	8	0.522	0.080	8	0.321	0.190
10	0.243	*** 0.000	10	0.285	*** 0.000	10	0.630	*** 0.000

Panel B: Distribution of P_i^w

	Mean			Variance		
w	Base	DDKS	Theory	Base	DDKS	Theory
2	2.00	2.12	2.00	1.89	2.92	2.00
4	3.99	4.20	4.00	4.42	5.83	4.00
6	5.97	6.25	6.00	6.94	10.37	6.00
8	7.98	8.33	8.00	11.28	14.93	8.00
10	9.97	10.39	10.00	15.63	20.07	10.00
	Skewness			Kurtosis		
w	Base	DDKS	Theory	Base	DDKS	Theory
2	0.76	1.30	0.71	3.62	6.20	3.50
4	0.65	0.44	0.50	3.49	2.79	3.25
6	0.47	0.62	0.41	3.08	3.16	3.17
8	0.43	0.41	0.35	3.29	2.59	3.12
10	0.48	0.02	0.32	3.49	2.24	3.10

Panel C: Distribution of t_i^w

	Mean		
	Base	DDKS	Theory
Mean	1.00	0.95	1.00
Variance	0.96	1.17	1.00
Skewness	3.06	2.25	2.00
Kurtosis	20.22	10.06	6.00
Prahl's test	** 0.005	*** 0.000	
KS test	*** 0.000	*** 0.000	

the j-th default on one such day, then we could set

$$U_n = \sum_{j=1}^{J_n} \max\left\{0, \log u_{n_j}\right\}. \tag{5}$$

We have analyzed our models using this alternative convention; see Table 17. The only affected models are those that include the contagion channel. We find that the signs, magnitudes, and statistical significance of the parameter estimates of all models are similar to those reported in Table 4, with few exceptions. The contagion sensitivity parameter b decreases by a factor of roughly 0.77. This is primarily driven by the fact that the average U_n under the formulation of Eq. (5) is about 1.3 times higher than under the formulation of Eq. (4). The estimated contagion decay parameter κ is similar under the formulations of Eq. (4) and (5). However, the maximum log-likelihoods are much lower under the specification of Eq. (5) than under the specification of Eq. (4). We conclude that Eq. (5) yields an inferior specification of contagion.

8.5. Industry-specific contagion

The formulations of Eqs. (4) and (5) do not distinguish among defaults in different industries. However, industry characteristics may be important determinants of the contagion effects induced by a default; see Hertzel et al. (2008), Hertzel and Officer (2012), Jorion and Zhang (2007), Jorion and Zhang (2009), and Lang and Stulz (1992), among others. Omitted industry-specific contagion

Table 16
Parameter estimates for each of the four model alternatives described in Table 3 after excluding defaults of financial firms. Under a financial firm default we understand any default by a firm in our data set with the Moody's 11 Sector Code "banking" or "finance, insurance, real estate." While unreported, the asymptotic standard errors are computed using the Hessian matrix of the log-likelihood at the parameter estimates reported. "." indicates significance at the 90% confidence level, "*" indicates significance at the 95% level, "**" significance at the 99% level, and "***" significance at the 99.9% level.

	Complete	Base	Base + Contagion	Base + Frailty
Constant a_0	* −1.8939	*** 0.9501	** −1.7729	0.6876
IP coefficient a_1	−0.0113	*** −0.0502	−0.0097	−0.0938
GDP a_3	** −0.1189	*** −0.1086	*** −0.1182	0.0403
RET coefficient a_4	0.0015	** −0.0050	0.0016	−0.0190
VOL coefficient a_5	0.0105	** 0.0115	0.0080	0.0074
LEVEL coefficient a_6	0.0319	** −0.0483	0.0321	. −0.6211
SLOPE coefficient a_8	−0.0005	* 0.0849	0.0078	. 0.5893
CRED coefficient a_9	0.0067	*** −0.7385	−0.0230	−0.1573
Frailty volatility σ	*** 0.2287			*** 0.1176
Frailty mean reversion rate k	*** 5.9771			* 0.0213
Frailty mean reversion level z	*** 0.0044			** 0.3278
Contagion sensitivity b	*** 0.0094		*** 0.0094	
Contagion decay rate κ	*** 0.0521		*** 0.0523	
Log-likelihood	617.45	466.55	617.36	601.61
Likehood ratio vs. Base model	*** 0.000		*** 0.000	*** 0.000
Likehood ratio vs. Base + Frailty model	*** 0.000			

Table 17
Parameter estimates for the complete model and the Base + Contagion submodel after accounting for multiple defaults in one day separately and for industry-specific contagion effects. While unreported, the asymptotic standard errors are computed using the Hessian matrix of the log-likelihood at the parameter estimates reported. We also report the p-values of likelihood ratio tests; see Table 8. "·" indicates significance at the 90% confidence level, "*" indicates significance at the 95% level, "**" significance at the 99% level, and "***" significance at the 99.9% level.

	Complete model		Base + Contagion submodel	
	Formulation of Eq. (5)	Formulation of Eq. (6)	Formulation of Eq. (5)	Formulation of Eq. (6)
Constant a_0	· −0.8879	** −1.6931	−0.8902	** −1.7084
IP coefficient a_1	0.0081	−0.0152	0.0090	−0.0146
GDP coefficient a_2	** −0.1432	** −0.0827	*** −0.1413	* −0.0805
RET coefficient a_4	0.0068	−0.0037	0.0076	−0.0040
VOL coefficient a_5	−0.0107	−0.0052	−0.0110	−0.0054
LEVEL coefficient a_6	0.0302	−0.0207	0.0335	−0.0159
SLOPE coefficient a_7	−0.0449	0.0261	−0.0555	0.0009
CRED coefficient a_8	−0.2329	0.3316	−0.2512	0.3069
Frailty volatility σ	0.2114	· 0.2222		
Frailty mean reversion rate k	6.0130	6.0374		
Frailty mean reversion level z	0.0037	0.0041		
Contagion sensitivity b	*** 0.0072		*** 0.0079	
Contagion decay rate κ	*** 0.0567	0.0325	*** 0.0618	· 0.0320
Banking sensitivity b		* 0.0088		*** 0.0088
Capital industries sensitivity $b + b_1$		0.0030		0.0031
Consumer industries sensitivity $b + b_2$		*** 0.0071		· 0.0070
Energy & environment sensitivity $b + b_3$		** 0.0064		** 0.0065
F.I.R.E sensitivity $b + b_4$		0.0000		0.0000
Media & publishing sensitivity $b + b_5$		0.0001		0.0001
Retail & distribution sensitivity $b + b_6$		*** 0.0116		* 0.0112
Technology sensitivity $b + b_7$		0.0006		0.0006
Transportation sensitivity $b + b_8$		0.0000		0.0000
Utilities sensitivity $b + b_9$		0.0001		0.0000
Log-likelihood	599.63	615.92	599.31	615.76
Likehood ratio vs. Base model	*** 0.000	*** 0.000	*** 0.000	*** 0.000
Likehood ratio vs. Base + Frailty model	1.000	*** 0.000		

Table 18
Average parameter estimates of a restricted version of the complete model of Table 3. We generate 50 sample data sets from variants of the complete model in which the frailty processes have dynamics that differ from the CIR dynamics in Eq. (2). In one model variant, the frailty follows an exponential Ornstein-Uhlenbeck process in which the logarithm of the frailty has a speed of reversion of 6.0203, a mean reversion level of −12.1107, and a volatility of 3.6370. In a second model variant, the frailty follows an exponential white noise process in which the logarithm of the frailty has monthly Gaussian innovations of mean −6.0461 and variance 1.0986. In a third model variant, the frailty is resampled every month from a Gamma distribution with scale parameter 2 and shape parameter 0.0021. The parameters specifying the frailty dynamics in these model variants are chosen so as to match the monthly mean and variances estimated for the frailty process in Table 4. For all model variants, the parameters specifying the contagion factor and the exposure to common macro factors are the same as in Table 5. The time horizon is the same as in Table 4. Empirical standard deviations of the estimated parameters are given parenthetically. We compute significance levels from a centralized normal distribution with standard deviations given by the empirical standard deviations of the parameter estimates. * indicates significance at the 95% level, ** significance at the 99% level, and *** significance at the 99.9% level. The *z*-scores of the parameter estimates are given in brackets when applicable.

	Frailty specified as:		
	Exponential OU	Exponential WN	Gamma
Constant a_0	*** −1.4921	*** −1.4781	*** −1.4772
	(0.1724)	(0.0980)	(0.1755)
	[0.6263]	[1.2437]	[0.6997]
GDP coefficient a_3	*** −0.1196	*** −0.1178	*** −0.1214
	(0.0018)	(0.0052)	(0.0048)
	[0.2090]	[0.4180]	[−0.2956]
Frailty volatility σ	*** 0.2149	*** 0.2152	*** 0.2135
	(0.0116)	(0.0063)	(0.0148)
Frailty mean reversion rate k	* 4.9714	* 6.4225	5.4915
	(2.4724)	(2.8907)	(4.2678)
Frailty mean reversion level z	0.0066	* 0.0046	0.0086
	(0.0041)	(0.0018)	(0.0086)
Contagion sensitivity b	*** 0.0088	*** 0.0080	*** 0.0093
	(0.0016)	(0.0015)	(0.0012)
	[−0.8590]	[−1.4523]	[−0.6323]
Contagion decay rate κ	*** 0.0550	*** 0.0497	*** 0.0589
	(0.0079)	(0.0093)	(0.0083)
	[−0.1664]	[−0.7057]	[0.3143]

effects may thus bias our estimates for the contagion channel. To demonstrate that this is not the case, we incorporate industry-specific contagion effects by formulating U_n as follows:

$$U_n = \sum_{j=1}^{J_n} \max\left\{0, \left(1 + \frac{b_{I_{n_j}}}{b}\right) \log u_{n_j}\right\}, \tag{6}$$

where I_{n_j} indicates the industry of the jth default on day T_n. The parameter $b_{I_{n_j}}$ measures the intensity sensitivity to a dummy variable that is equal to one if and only if the jth default on day T_n occurs in industry I_{n_j}.

We reestimate the models that include the contagion channel using Eq. (6). We consider ten different sectors as defined by Moody's 11 Code (see Fig. 4). Table 17 summarizes our results. The contagion sensitivity is smaller after accounting for different industries. The average sector contagion sensitivity, weighted by empirical default frequencies, is 0.0044. Significant contagion occurs after a default in the banking, the consumer industries, the energy and environment, and the retail and distribution sectors. An average default in these sectors increases the aggregate default intensity by 1.96, 1.88, 1.81, and 2.87 defaults per year, respectively. Despite the contagion sensitivity being smaller, the estimated contagion decay rate κ under formulation of Eq. (6) implies a much larger persistency of the contagion factor Y. Defaults in the banking, consumer industries, energy and environment, and the retail and distribution sectors cause large jumps up in the aggregate default intensity, and these jumps persist with a half-life of about five months (the half-life under specification of Eq. (4) is only three months). All other parameter estimates and standard errors are roughly the same for the formulation of Eq. (6) as for our original formulation of Eq. (4). Likelihood ratio tests of the complete and nested models under the formulations of Eq. (6) and (4) yield the same results. We conclude that omitted industry-specific contagion effects do not bias any of our results.

Consistent with Hertzel et al. (2008) and Jorion and Zhang (2009), our results on industry effects confirm that contagion can occur across industries. Our results also confirm that contagion does not necessarily originate in the financial sector, extending our findings of Section 8.3. A large industrial default can have a similar impact on aggregate default rates as a default with comparable loss in the financial sector.

8.6. Frailty misspecification

We study the robustness of our results with respect to alternative specifications of the frailty factor. We consider three variants of the complete model with different formulations of the frailty factor. In one formulation, the frailty follows an exponential Ornstein-Uhlenbeck process as in Duffie et al. (2009). In a second formulation, the frailty is given by a monthly resampled gamma random variable, following Chava et al. (2011). In a third formulation, the frailty follows an exponential white noise process with monthly innovations. In each of these formulations, the means and variances of the innovations were selected to match the monthly mean and volatility estimates of the frailty factor in the original model, given in Table 4.

We generate 50 sample data sets from each of these model variants via Monte Carlo simulation, and then estimate the original complete model based on the simulated data. Table 18 provides summary statistics of the resulting parameter estimates. We see that the frailty factor is significant for all variants of the data-generating model. This confirms the ability of our estimation approach to detect the presence of latent variables even if the frailty model is misspecified. Moreover, we see that the role of the contagion channel is measured correctly under frailty misspecification. We conclude that our assessment of the contagion channel is robust to misspecification of the frailty factor.

9. Concluding remarks

The US credit markets have suffered several significant clusters of corporate default events. This paper studies the sources of clustering using data on industrial and financial default timing in the US between 1970 and 2012. It presents strong evidence for the existence of a contagion channel of default clustering, after controlling for firms' joint exposure to observable and latent systematic risk factors.

The results have important implications for the management of portfolio credit risk at financial institutions. The models widely used to estimate risk capital address only the default clustering due to firms' joint exposure to observable risk factors. Our results suggest that these models significantly understate the risk of large losses from defaults. This indicates that institutions using these models might end up holding inadequate capital to withstand the large losses induced by default clusters as those seen during the financial crisis of 2007–2009. Our results suggest to estimate risk capital using an approach that controls for latent systematic factors and also addresses the impact of past failures on default rates.

An understanding of the sources of default clustering is also essential for the pricing, rating, risk analysis, and trading of securities that are exposed to correlated default risk, such as collateralized debt obligations. Our findings suggest to account for a contagion channel when examining such securities.

Appendix A. Reduced-form model

This appendix provides some technical details regarding our model of default timing (Section 4). The default data is a realization of a marked point process $(T_n, U_n)_{n\geq 1}$ defined on a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with an information filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$ satisfying the usual conditions of right-continuity and completeness (see Protter, 2004). That is, $(T_n)_{n\geq 1}$ is an increasing sequence of stopping times tending to ∞. The times T_n represent the dates in the sample with at least one default. Moreover, $(U_n)_{n\geq 1}$ is a sequence of $\mathcal{F}_{T_n}$-measurable random variables, where U_n represents the total amount of defaulted debt on date T_n. We assume that there exists a positive process λ such that the variables $N_t - \int_0^t \lambda_s ds$ form a local martingale, where $N_t = \sum_{n\geq 1} 1_{\{T_n \leq t\}}$. The process λ is the intensity of N. We have that $\mathbb{E}[N_{t+\Delta} - N_t \mid \mathcal{F}_t] \approx \lambda_t \Delta$ for small $\Delta > 0$, suggesting the interpretation of a conditional mean arrival rate for λ.

Appendix B. Goodness-of-fit tests

This appendix extends a key result obtained by Das et al. (2007) for doubly stochastic models with observable risk factors. The extension allows us to construct goodness-of-fit tests for our default timing model of Eq. (1), which violates the doubly stochastic hypothesis due to the presence of the frailty and self-exciting terms. A key role is played by the posterior mean $h_t = \mathbb{E}[\lambda_t \mid \mathcal{G}_t]$ of the intensity, which is discussed in Internet Appendix B.

Proposition 1. Let H_t be the right-continuous inverse to $C_t = \int_0^t h_u du$, i.e.,

$$H_t = \max\left\{s \geq 0 : \int_0^s h_u du \leq t\right\}. \tag{B.1}$$

The time-changed process J defined by $J_t = N_{H_t}$ is a standard Poisson process on $[0, \infty)$ relative to $\mathbb{P}$ and the minimal right-continuous completion $\mathbb{H}$ of $(\mathcal{G}_{H_t})_{t\geq 0}$.

Proof of Proposition 1. A result of Meyer (1971) implies that a counting process with compensator that is continuous and increasing to ∞ almost surely can be transformed into a standard Poisson process by a change of time given by the compensator. Relative to the filtration $\mathbb{G}$ generated by the observable data, the counting process N has compensator C given by $C_t = \int_0^t h_u du$, where h is the posterior mean of the intensity (see Internet Appendix B). Since each H_t is a stopping time with respect to $\mathbb{G}$, we can define the stopping time σ-algebra $\mathcal{H}_t = \mathcal{G}_{H_t}$, which is the smallest σ-algebra containing all right-continuous left limit processes sampled at H_t. Meyer's theorem implies that the time-scaled process J is a standard Poisson process in the time-changed filtration $\mathbb{H}$, which is the minimal right-continuous completion of $(\mathcal{H}_t)_{t\geq 0}$. □

Proposition 1 states that the default count N_t can be transformed into a standard Poisson process by a change of time given by the cumulative posterior mean of the intensity. The posterior mean is given by Proposition B.1 in Internet Appendix B, and can be computed using Algorithm A.2 in Internet Appendix A. This result allows us to test model fit by testing the Poisson property of the time-scaled default times.

References

Acemoglu, D., Carvalho, V.M., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The network origins of aggregate fluctuations. Econometrica 80 (5), 1977–2016.

Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi, A., 2015. Systemic risk and stability in financial networks. Am. Econ. Rev. 105 (2), 564–608.

Acharya, V.V., Amihud, Y., Bharath, S.T., 2013. Liquidity risk of corporate bond returns: conditional approach. J. Financ. Econ. 110 (2), 358–386.

Aït-Sahalia, Y., Cacho-Diaz, J., Laeven, R., 2015. Modeling financial contagion using mutually exciting jump processes. J. Financ. Econ. 117 (3), 585–606.

Aragon, G.O., Strahan, P.E., 2012. Hedge funds as liquidity providers: evidence from the Lehman bankruptcy. J. Financ. Econ. 103 (3), 570–587.

Atkeson, A.G., Eisfeldt, A.L., Weill, P.-O., 2017. Measuring the financial soundness of US firms, 1926–2012. Res. Econ. 71 (3), 613–635.

Azizpour, S., Giesecke, K., Kim, B., 2011. Premia for correlated default risk. J. Econ. Dyn. Control 35 (8), 1340–1357.

Bai, J., Collin-Dufresne, P., Goldstein, R.S., Helwege, J., 2015. On bounding credit-event risk premia. Rev. Financ. Stud. 28 (9), 2608–2642.
Bams, D., Pisa, M., Wolff, C., 2016. Ripple effects and industry defaults. Unpublished Working Paper. Centre for Economic Policy Research.
Baur, D.G., 2012. Financial contagion and the real economy. J. Bank. Financ. 36 (10), 2680–2692.
Benzoni, L., Collin-Dufresne, P., Goldstein, R., Helwege, J., 2015. Modeling credit contagion via the updating of fragile beliefs. Rev. Financ. Stud. 28 (7), 1960–2008.
Bernanke, B.S., Gertler, M., Gilchrist, S., 1999. The Financial Accelerator in a Quantitative Business Cycle Framework. In: Handbook of Macroeconomics. Elsevier, Amsterdam, pp. 1341–1393.
Berndt, A., Ritchken, P., Sun, Z., 2010. On correlation and default clustering in credit markets. Rev. Financ. Stud. 23 (7), 2680–2729.
Bhansali, V., Gingrich, R., Longstaff, F., 2008. Systemic credit risk: what is the market telling us? Financ. Anal. J. 64 (4), 16–24.
Bharath, S.T., Shumway, T., 2008. Forecasting default with the Merton distance to default model. Rev. Financ. Stud. 21 (3), 1339–1369.
Boone, A.L., Ivanov, V.I., 2012. Bankruptcy spillover effects on strategic alliance partners. J. Financ. Econ. 103 (3), 551–569.
Bowsher, C., 2007. Modelling security market events in continuous time: intensity based, multivariate point process models. J. Econom. 141, 876–912.
Brunnermeier, M.K., 2009. Deciphering the liquidity and credit crunch 2007–2008. J. Econ. Perspect. 23 (1), 77–100.
Cameron, A.C., Trivedi, P.K., 1998. Regression Analysis of Count Data (Econometric Society Monographs). Cambridge University Press..
Campbell, J.Y., Hilscher, J., Szilagyi, J., 2008. In search of distress risk. J. Financ. 63 (6), 2899–2939.
Chakrabarty, B., Zhang, G., 2012. Credit contagion channels: market microstructure evidence from Lehman Brothers' bankruptcy. Financ. Manag. 41 (2), 320–343.
Chava, S., Jarrow, R., 2004. Bankruptcy prediction with industry effects. Rev. Financ. 8, 537–569.
Chava, S., Stefanescu, C., Turnbull, S., 2011. Modeling expected loss. Manag. Sci. 57 (7), 1267–1287.
Christoffersen, P., 1998. Evaluating interval forecasts. Int. Econ. Rev. 39, 842–862.
Collin-Dufresne, P., Goldstein, R.S., Martin, J.S., 2001. The determinants of credit spread changes. J. Financ. 56 (6), 2177–2207.
Collin-Dufresne, P., Goldstein, R.S., Yang, F., 2012. On the relative pricing of long-maturity index options and collateralized debt obligations. J. Financ. 67 (6), 1983–2014.
Cont, R., Minca, A., 2013. Recovering portfolio default intensities implied by CDO quotes. Math. Financ. 23 (1), 94–121.
Coval, J.D., Jurek, J.W., Stafford, E., 2009. Economic catastrophe bonds. Am. Econ. Rev. 99 (3), 628–666.
Covas, F., Den Haan, W.J., 2011. The cyclical behavior of debt and equity finance. Am. Econ. Rev. 101 (2), 877–899.
Cox, D.R., 1972. Regression models and life-tables. J. R. Stat. Soc. Series B (Methodological) 34 (2), 187–220.
Das, S., Duffie, D., Kapadia, N., Saita, L., 2007. Common failings: how corporate defaults are correlated. J. Financ. 62, 93–117.
Duffie, D., 1998. First-to-default valuation. Unpublished Working Paper, Stanford University.
Duffie, D., Eckner, A., Horel, G., Saita, L., 2009. Frailty correlated default. J. Financ. 64, 2089–2123.
Duffie, D., Garleanu, N., 2001. Risk and valuation of collateralized debt obligations. Financ. Anal. J. 57 (1), 41–59.
Duffie, D., Saita, L., Wang, K., 2007. Multi-period corporate default prediction with stochastic covariates. J. Financ. Econ. 83 (3), 635–665.
Dumontaux, N., Pop, A., 2013. Understanding the market reaction to shockwaves: evidence from the failure of Lehman Brothers. J. Financ. Stab. 9 (3), 269–286.
Eisenberg, L., Noe, T.H., 2001. Systemic risk in financial systems. Manag. Sci. 47 (2), 236–249.
Elliott, M., Golub, B., Jackson, M.O., 2014. Financial networks and contagion. Am. Econ. Rev. 104 (10), 3115–3153.
Engle, R., Manganelli, S., 2004. CAViaR: conditional autoregressive value at risk by regression quantiles. J. Bus. Econ. Stat. 22 (4), 367–381.
Errais, E., Giesecke, K., Goldberg, L., 2010. Affine point processes and portfolio credit risk. SIAM J. Financ. Math. 1, 642–665.
Fernando, C.S., May, A.D., Megginson, W.L., 2012. The value of investment banking relationships: evidence from the collapse of Lehman Brothers. J. Financ. 67 (1), 235–270.
Giesecke, K., 2004. Correlated default with incomplete information. J. Bank. Financ. 28, 1521–1545.
Giesecke, K., Longstaff, F., Schaefer, S., Strebulaev, I., 2011. Corporate bond default risk: a 150-year perspective. J. Financ. Econ. 102 (2), 233–250.
Giesecke, K., Longstaff, F., Schaefer, S., Strebulaev, I., 2014. Macroeconomic effect of corporate default crises: a long-term perspective. J. Financ. Econ. 111 (2), 297–310.
Giesecke, K., Schwenkler, G., 2018. Filtered likelihood for point processes. J. Econom. 204 (1), 33–53.
Gilchrist, S., Zakrajšek, E., 2012. Credit spreads and business cycle fluctuations. Am. Econ. Rev. 102 (4), 1692–1720.
Gouriéroux, C., Monfort, A., Renne, J., 2014. Pricing default events: surprise, exogeneity and contagion. J. Econom. 182 (2), 397–411.
Hamilton, D., 2005. Moody's senior ratings algorithm and estimated senior ratings. Moody's Global Credit Research.
Hawkes, A.G., 1971. Spectra of some self-exciting and mutually exciting point processes. Biometrika 58 (1), 83–90.
Helwege, J., 2010. Financial firm bankruptcy and systemic risk. J. Int. Financ. Mark. Inst. Money 20 (1), 1–12.
Helwege, J., Zhang, G., 2016. Financial firm bankruptcy and contagion. Rev. Financ. 20 (4), 1321–1362.
Hendricks, D., 1996. Evaluation of value-at-risk models using historical data. Econ. Policy Rev. 2 (1), 39–69.
Hertzel, M.G., Li, Z., Officer, M.S., Rodgers, K.J., 2008. Inter-firm linkages and the wealth effects of financial distress along the supply chain. J. Financ. Econ. 87 (2), 374–387.
Hertzel, M.G., Officer, M.S., 2012. Industry contagion in loan spreads. J. Financ. Econ. 103, 493–506.
Hilscher, J., Wilson, M., 2017. Credit ratings and credit risk: is one measure enough? Manag. Sci. 63 (10), 3414–3437.
Hull, J.C., White, A.D., 2008. Dynamic models of portfolio credit risk: a simplified approach. J. Deriv. 15 (4), 9–28.
Jarrow, R.A., Lando, D., Yu, F., 2005. Default risk and diversification: theory and applications. Math. Financ. 15, 1–26.
Jarrow, R.A., Yu, F., 2001. Counterparty risk and the pricing of defaultable securities. J. Financ. 56 (5), 555–576.
Jorion, P., Zhang, G., 2007. Good and bad credit contagion: evidence from credit default swaps. J. Financ. Econ. 84 (3), 860–883.
Jorion, P., Zhang, G., 2009. Credit contagion from counterparty risk. J. Financ. 64 (5), 2053–2087.
Karlis, D., Xekalaki, E., 2000. A simulation comparison of several procedures for testing the Poisson assumption. Stat. 49, 355–382.
Koopman, S.J., Lucas, A., Monteiro, A., 2008. The multi-stage latent factor intensity model for credit rating transitions. J. Econom. 142 (1), 399–424.
Koopman, S.J., Lucas, A., Schwaab, B., 2011. Modeling frailty-correlated defaults using many macroeconomic covariates. J. Econom. 162 (2), 312–325.
Koopman, S.J., Lucas, A., Schwaab, B., 2012. Dynamic factor models with macro, frailty, and industry effects for US default counts: the credit crisis of 2008. J. Bus. Econ. Stat. 30 (4), 521–532.
Kupiec, P.H., 1995. Techniques for verifying the accuracy of risk measurement models. J. Deriv. 3 (2), 73–84.
Lando, D., Nielsen, M., 2010. Correlation in corporate defaults: contagion or conditional independence? J. Financ. Intermed. 19 (3), 355–372.
Lang, L., Stulz, R., 1992. Contagion and competitive intra-industry effects of bankruptcy announcements. J. Financ. Econ. 32, 45–60.
Lee, Y., Poon, S.-H., 2014. Forecasting and decomposition of portfolio credit risk using macroeconomic and frailty factors. J. Econ. Dyn. Control 41, 69–92.
Longstaff, F., Piazzesi, M., 2004. Corporate earnings and the equity premium. J. Financ. Econ. 74 (3), 401–421.
Longstaff, F., Rajan, A., 2008. An empirical analysis of the pricing of collateralized debt obligations. J. Financ. 63 (2), 529–563.
Meyer, P.-A., 1971. Démonstration simplifée d'un théorème de Knight. In: Séminaire de Probabilités V, Lecture Notes in Mathematics 191. Springer-Verlag Berlin Heidelberg, Berlin, pp. 191–195.
Mistrulli, P.E., 2011. Assessing financial contagion in the interbank market: maximum entropy versus observed interbank lending patterns. J. Bank. Financ. 35 (5), 1114–1127.
Monfort, A., Renne, J.-P., 2013. Default, liquidity, and crises: an econometric framework. J. Financ. Econom. 11 (2), 221–262.
Mortensen, A., 2006. Semi-analytical valuation of basket credit derivatives in intensity-based models. J. Deriv. 13, 8–26.
Protter, P., 2004. Stochastic Integration and Differential Equations. Springer-Verlag, New York.
Stulz, R.M., 2010. Credit default swaps and the credit crisis. J. Econ. Perspect. 24 (1), 73–92.
Swamy, P.A.V.B., 1970. Efficient inference in a random coefficient regression model. Econometrica 38 (2), 311–323.

Tian, S., Yu, Y., Guo, H., 2015. Variable selection and corporate bankruptcy forecasts. J. Bank. Financ. 52, 89–100.
Winkelmann, R., 1995. Duration dependence and dispersion in count-data models. J. Bus. Econ. Stat. 13 (4), 467–474.
Yang, J., Zhou, Y., 2013. Credit risk spillovers among financial institutions around the global credit crisis: firm-level evidence. Manag. Sci. 59 (10), 2343–2359.