Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.

This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

http://www.elsevier.com/copyright

Author's personal copy

Journal of Economic Dynamics & Control 35 (2011) 1340–1357

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

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

Premia for correlated default risk

Shahriar Azizpour [a], Kay Giesecke [b,*], Baeho Kim [c]

[a] Deutsche Bank, London, United Kingdom
[b] Department of Management Science and Engineering, Stanford University, Stanford, United States
[c] Korea University Business School, Seoul, South Korea

ARTICLE INFO

Article history:
Received 30 November 2008
Accepted 16 March 2011
Available online 30 March 2011

Keywords:
Correlated defaults
Risk premium
Measure change
Maximum likelihood

JEL classification:
G12
C58

ABSTRACT

Using data on corporate default experience in the U.S. and market rates of CDX index and tranche swaps of various maturities, we estimate reduced-form models of correlated default timing in the CDX High Yield and Investment Grade portfolios under actual and risk-neutral probabilities. The striking contrast between the estimated processes followed by the actual and risk-neutral arrival intensities of defaults, and between the parameters governing the actual and risk-neutral dynamics of the risk-neutral intensities, indicates the presence of substantial default risk premia in CDX swap market rates. The effects of risk premia on swap rates covary strongly across maturities, and depend on general stock market volatility and several measures of credit spreads. Large moves in the effects of these premia on swap rates have natural interpretations in terms of economic and financial market developments during the sample period, April 2004 to October 2007. Our results suggest that a large portion of the movements in CDX swap market rates observed during the sample period may be caused by changing attitudes toward correlated default risk rather than changes in the economic factors affecting the actual risk of clustered defaults, which ultimately governs swap payoffs.

© 2011 Elsevier B.V. All rights reserved.

1. Introduction

Does the credit index derivatives market price correlated corporate default risk? If so, how do risk premia behave over time, across index credit quality, and across maturities? What are their components? These questions are of interest to investors allocating assets to credit index and tranche markets, and investors seeking relative value opportunities across the index capital structure, universe, and term structure. They are also relevant to policymakers, who often look at signals from credit markets to inform their decisions. Finally, the answers to these questions are of substantial academic interest, as they yield important insights into the movements of the prices of assets exposed to clustered default risk relative to the variation of the economic factors affecting payoffs.

In order to address these issues, we estimate reduced-form models of correlated default timing under actual and risk-neutral probabilities, using data on corporate default experience in the U.S. and time series of market rates of index and tranche swaps referenced on the CDX North American High Yield and Investment Grade portfolios. The square-root jump-diffusion models of risk-neutral CDX portfolio default intensities we propose incorporate the influence of a Brownian risk

* Corresponding author. Tel.: +1 650 723 9265; fax: +1 650 723 1614.
E-mail address: giesecke@stanford.edu (K. Giesecke).
URL: http://www.stanford.edu/~giesecke (K. Giesecke).

0165-1889/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.jedc.2011.03.010

Author's personal copy

factor and past default timing on risk-neutral arrival rates of events. They also address the negative correlation between recovery rates and default rates, and are found to capture most of the time-series variation of market swap rates between April 2004 and October 2007.

Our maximum likelihood estimates provide strong evidence that the CDX index and tranche swap market prices correlated corporate default risk. The contrast we uncover between the actual and risk-neutral distributions of the risk-neutral portfolio intensities indicates the presence of substantial premia for the diffusive and jump volatility of index and tranche swap mark-to-market values. This volatility is associated with systematic and default-induced movements of risk-neutral default rates, and is of concern to mark-to-market investors required to mark their positions to the market. The distinction we find between the stochastic processes followed by the actual and risk-neutral portfolio intensities suggests the presence of substantial premia for uncertainty regarding the timing of defaults of CDX index constituents. This jump-to-default risk plays an important role for investors selling index or tranche protection, who bear exposure to actual default losses in addition to any volatility in the mark-to-market values of their positions. Jump-to-default premia are found to be significantly higher for the Investment Grade portfolio than for the High Yield portfolio, providing evidence that the correlated default risk of high-quality issuers is much harder to diversify and to hedge than the correlated default risk of low-quality names.

The estimated effects of risk premia on CDX swap rates covary strongly across maturities and exhibit substantial time-series variation. They are strongly dependent on general stock market volatility, as measured by the VIX, the corporate credit yield spread, and the LIBOR-OIS spread. The large moves in these premia have natural interpretations in terms of economic and financial market developments, including the May 2005 "correlation crisis" triggered by the downgrades of Ford and General Motors and the defaults of several auto parts suppliers, as well as the financial crisis of 2007–2009. Our results suggest that a large portion of the movements in CDX swap market rates between 2004 and 2007 was caused by changing attitudes toward correlated default risk rather than changes in the economic factors affecting the actual risk of clustered defaults, which ultimately governs CDX swap payoffs. This finding complements the results of Berndt et al. (2008), who show that the majority of the variation in single-name credit swap rates during 2002–2004 was due to variation in risk premia and not expected loss rates on the underlying reference bonds.

Prior research has examined default risk premia. Using corporate bond price data and average historical default frequencies by credit rating, Driessen (2005) estimates the relationship between actual and risk-neutral default intensities of individual issuers. Berndt et al. (2008) infer this relationship from single-name credit swap market rates and estimated Moody's KMV Expected Default Frequencies. Pan and Singleton (2008) estimate the dynamics of risk-neutral intensities of sovereign issuers under actual and risk-neutral probabilities from sovereign credit swap rates. Carr and Wu (2007) estimate these dynamics from currency option prices and sovereign credit swap rates. These prior studies focus on the behavior of entity-level default risk premia.[1] They do not address the implications of default clustering. In this paper, we focus on portfolio-level premia for correlated corporate default risk, as implied by market rates of index and tranche swaps referenced on the various CDX portfolios, and actual default timing in the U.S.

Incorporating the effects of default correlation, Eckner (2007) estimates from market rates of CDX Investment Grade index, tranche and constituent credit swaps a doubly stochastic model of actual and risk-neutral intensities of individual firms. In this model, firms default independently of one another given a set of common risk factors influencing different firms. In contrast to this study, we formulate and estimate top-down models of the actual and risk-neutral intensities of a CDX portfolio as a whole. We also relax the doubly stochastic assumption by incorporating the impact of defaults on a portfolio intensity, i.e., self-exciting effects. Das et al. (2007) and Azizpour et al. (2008) find strong evidence of the violation of the doubly stochastic hypothesis under the actual measure, and Collin-Dufresne et al. (2009) present related evidence under the risk-neutral measure. Because of the impact of defaults on arrival rates, in our model correlated default risk is not conditionally diversifiable in the sense of Jarrow et al. (2005). Therefore, a default event may command a premium even in well-diversified portfolios of credit-sensitive positions. Our estimates indicate the presence of substantial premia for default event (jump-to-default) risk.

Addressing the implications of default clustering in an economy with infinitely many firms, Coval et al. (2009) calibrate from market rates of CDX Investment Grade index and tranche swaps and prices of S&P 500 index options a model of the risk-neutral and actual default probabilities of a representative CDX constituent. In contrast to this paper, we develop likelihood methods to exploit the full time-series dynamics of the CDX market rates when fitting risk-neutral CDX portfolio intensities, and to utilize the rich historical default experience in the U.S. to accurately estimate the actual CDX intensity. Moreover, we incorporate the negative correlation between default rates and recovery rates documented by Altman et al. (2005) and others when estimating risk premia, we disentangle various components of these premia, and we provide a term structure perspective on the premia. The fits to the CDX market data we obtain are an improvement over those obtained by Coval et al. (2009).

Longstaff and Rajan (2008) and Bhansali et al. (2008) formulate risk-neutral intensity models of the CDX Investment Grade portfolio that incorporate the effects of default correlation. They calibrate these models from CDX market index and

[1] Theoretical models of entity-level default risk premia are provided by Blanchet-Scalliet et al. (2005), Campi et al. (2009) and Giesecke and Goldberg (2003).

Author's personal copy

tranche rates using least-squares methods. In contrast to these studies, we develop and implement likelihood estimators for the risk-neutral and actual CDX portfolio intensities, and we extract estimates of risk premia for correlated default risk.

We begin in Section 2 by estimating models of CDX portfolio intensities under actual probabilities. In Section 3, we identify the equivalent change of measure from actual to risk-neutral probabilities by estimating models of risk-neutral CDX portfolio intensities. The components of the measure change reflect the premia for correlated default risk. They are analyzed in Section 4. Section 5 concludes.

2. Actual intensity from default events

This section specifies and estimates from U.S. default experience reduced-form models of actual default timing in the CDX Investment Grade and High Yield portfolios. We begin by explaining the features of the data.

2.1. Corporate default data

Data on U.S. default timing are from Moody's Default Risk Service, which provides detailed issue and issuer information on industry, rating, date and type of default, and other items. The data cover the period January 1970 to October 2007. An issuer is included in our data set if it is not a sovereign and has a senior rating, which is an issuer-level rating generated by Moody's from ratings of particular debt obligations using the Senior Rating Algorithm described in Hamilton (2005).[2] As of October 2007, the data set includes a total of 5057 firms, of which 3263 are investment-grade rated issuers.

For our purposes, a "default" is a credit event in any of the following Moody's default 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; (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. Following Moody's, 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.

Our data set includes a total of 1398 defaults. These occur on 930 distinct dates. The distribution of these defaults, shown in Fig. 1, indicates several major default clusters. We note the cluster following the 1987 crash, and the cluster associated with the burst of the internet bubble in 2001.

2.2. Actual portfolio intensity

In order to formulate our statistical model of default timing, we fix a complete probability space $(\Omega,\mathcal{F},P)$ and an information filtration $(\mathcal{F}_t)_{t\geq 0}$ satisfying the usual conditions (see Protter, 2004). The probability P is the actual data-generating measure. Defaults in a portfolio of interest arrive at stopping times $T_1<T_2<\cdots$ with intensity λ satisfying $\int_0^t \lambda_s\,ds<\infty$ for all t. The intensity represents the conditional mean arrival rate of defaults in the portfolio, relative to P. We have $E(N_{t+h}-N_t|\mathcal{F}_t)\approx h\lambda_t$ for "small" h, where N is the process counting defaults in the portfolio, defined by $N_t=\sum_n 1_{\{T_n\leq t\}}$.

The portfolio intensity λ is intended to capture the clustering in the default timing data displayed in Fig. 1. This is achieved through a simple auto-regressive specification, in which the current intensity is a function of past default timing. More precisely, we follow Giesecke and Kim (2011) and suppose that λ is governed by

$$d\lambda_t=\kappa_t(c_t-\lambda_t)\,dt+dJ_t \tag{1}$$

where $\lambda_0>0$ is the initial intensity, $\kappa_t=\kappa\lambda_{T_{N_t}}$ is the decay rate, $c_t=c\lambda_{T_{N_t}}$ is the reversion level, J is a jump process given by $J_t=\sum_{n=1}^{N_t}\max(\gamma,\delta\lambda_{T_n-})$, and $\kappa>0$, $c\in(0,1)$, $\delta>0$ and $\gamma\geq 0$ are parameters. We suppose that $c(1+\delta)<1$ ensuring that N does not explode. Let $\Theta=(\kappa,c,\delta,\gamma,\lambda_0)$ be the parameter vector governing λ.

The intensity λ jumps at a default. The response of λ is proportional to the value of the intensity just before the default. It is governed by the parameter δ, with minimum response given by γ. The higher the pre-event intensity, the weaker the firms in the portfolio tend to be, and the stronger the response of the arrival rate. The influence of a default fades away with time, at a rate proportional to the value of the intensity at the default. The decay rate is governed by the parameter κ. In the absence of an event, the intensity tends to a level proportional to the value of the intensity at the previous default. The factor of proportionality is given by the parameter c.

The auto-regressive specification (1) of the portfolio intensity λ is relatively parsimonious. Below we perform tests demonstrating that this specification accurately captures the clustering of defaults in the CDX High Yield and Investment Grade portfolios, as implied by the default timing data described in Section 2.1.

---

[2] We follow a common convention and subsume the categories Caa and Ca into C. We also subsume the numerical sub-categories Aa1, Aa2, Aa3 into Aa, and proceed similarly in the case of the other numerical sub-categories. Then the set of rating categories is $\mathcal{R}=\{$Aaa,Aa,A,Baa,Ba,B,C,WR$\}$. The category WR indicates a withdrawn rating.

Author's personal copy

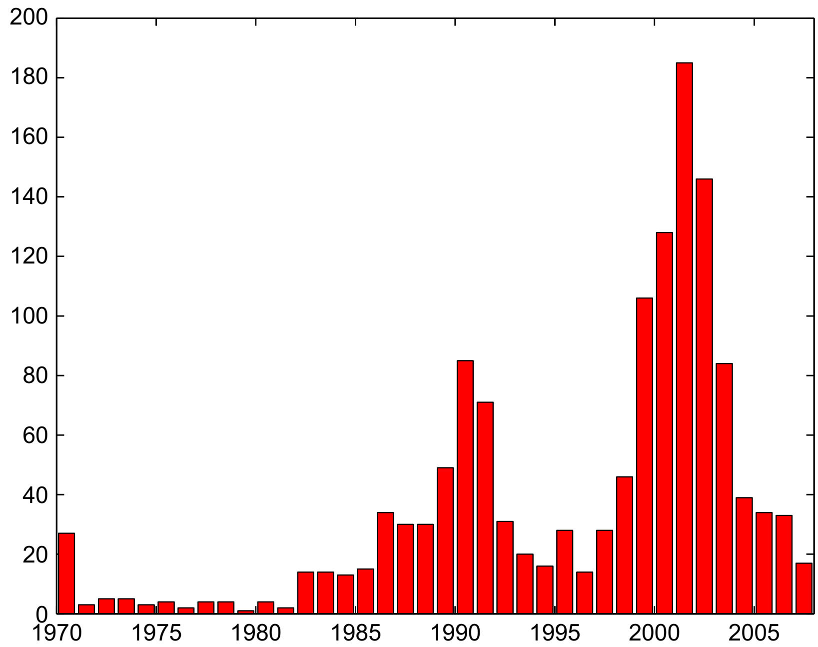

Fig. 1. Annual default data. A bar indicates the number of defaults of Moody's rated corporate issuers, as of 10/31/2007. Source: Moody's Default Risk Service.

2.3. *Likelihood estimators*

In order to estimate the parameter vector Θ governing the actual intensity λ of defaults in the CDX High Yield (HY) and Investment Grade (IG) portfolios, we require a realization of the default process N for each of these portfolios. However, the CDX portfolios were created only in 2004, so the available history is very limited. In fact, between 2004 and 2007 only four HY firms defaulted, and there were no IG defaults at all. To address this problem, we consider the entire history of defaults in the U.S. between 1970 and 2007, providing a realization of the economy-wide default process N^e generated by the default stopping times $T_1^e < T_2^e < \cdots$ in the universe of Moody's rated names.[3] We use a resampling approach developed by Giesecke and Kim (2011) to estimate a CDX intensity λ from the realization of N^e. The basic idea is to generate alternative portfolio default histories from N^e, and to estimate λ from these histories by the method of maximum likelihood.

Rejection sampling is used to generate alternative realizations of N from the available realization of N^e. We randomly select an event time T_n^e as an event time of N with probability $Z_{T_n^e}$, where Z is a predictable thinning process with values in [0,1]. Proposition 2.1 in Giesecke and Kim (2011) implies that the resulting sequence of event times has arrival intensity $Z\lambda^e$, where λ^e is the intensity of N^e. Thus, in order for the sequence of event times to represent a potential portfolio default history governed by the intensity λ, we need to specify Z as the ratio λ/λ^e. With this choice of Z, we can estimate the parameters of λ from a set of alternative portfolio default histories resampled from N^e.

The value λ_t/λ_t^e represents the conditional probability at time t that the next defaulter is a name in the reference portfolio of interest, given that a default occurs in the economy in the next instant. We specify Z nonparametrically, guided by this representation. Intuitively, Z must reflect the relation between the issuer composition of the economy as represented by the collection of Moody's rates issuers in our data base, and the issuer composition of the reference portfolio. We use the ratings of the portfolio constituents to describe the issuer composition.[4] More precisely, let $[0,\tau]$ be the sample period, January 1970 to October 2007. Let $\mathcal{R} = \{\text{Aaa,Aa,A,Baa,Ba,B,C,WR}\}$ be the set of rating categories, $X_\tau(\rho)$ be the number at time τ of constituent firms with rating $\rho \in \mathcal{R}$, and $X_t^e(\rho)$ be the number at time $t \in [0,\tau]$ of ρ-rated firms in the economy. We require that $X_\tau(\rho) \le X_t^e(\rho)$ for all $t \le \tau$ and $\rho \in \mathcal{R}$. For $t \le \tau$, we take

$$Z_t = \sum_{\rho \in R_{t-}} \frac{X_\tau(\rho)}{X_{t-}^e(\rho)} P(\varrho_{N_{t-}^e+1} = \rho | \mathcal{F}_{t-}) \tag{2}$$

[3] Each default event in the data base has a time stamp; the resolution is one day. There are days with multiple events whose exact timing during the day cannot be established. In order to obtain a strictly increasing sequence (T_n^e) of events, we assume that an event date is measured with uniformly distributed noise. We convert a raw economy-wide event date to a real-valued calendar time equal to 12 am on the day of the event, and draw the noise from a uniform distribution on [0,1/365]. With this choice, the randomization does not alter the original time stamp of an observed event. We have experimented with several alternative randomization schemes but have found that the model estimation results are insensitive to the choice of randomization scheme.

[4] Also Driessen (2005) uses ratings to characterize issuers for the purpose of estimating firm-level default intensities under the actual measure.

Author's personal copy

Table 1
Maximum likelihood estimates of the parameters of the portfolio intensity λ, as of October 2007. Estimates of asymptotic standard errors are given parenthetically. 10 K resampling scenarios are used for the HY index and 200 K for the IG index.

Idx	κ	c	δ	γ	λ_0
HY	0.214 (0.0007)	0.029 (0.0009)	0.270 (0.0015)	0.674 (0.0315)	9.268 (0.0681)
IG	1.552 (0.1946)	0.161 (0.0117)	0.025 (0.0143)	0.069 (0.0156)	0.559 (0.0404)

where R_t is the set of rating categories $\rho \in \mathcal{R}$ for which $X_t^e(\rho) > 0$, and where ϱ_n is the rating at the time of default of the nth defaulter in the economy. Formula (2) suggests an interpretation of the value Z_t as the conditional "empirical" probability that the next defaulter is a portfolio constituent. This conditional probability respects the ratings of the reference names, as $P(\varrho_{N^e_{t-}+1} = \rho | \mathcal{F}_{t-})$ is the conditional probability that the next defaulter has rating ρ. Our estimator $\mu_t(\rho)$ of this latter conditional probability is based on the ratings of the defaulters in [0,t). For $\rho \in \mathcal{R}$, it is given by

$$\mu_t(\rho) = \frac{\sum_{n=1}^{N^e_{t-}} 1_{\{\varrho_n = \rho\}} + \zeta}{\sum_{n=1}^{N^e_{t-}} 1_{\{\varrho_n \in R_{t-}\}} + \zeta |R_{t-}|} 1_{\{\rho \in R_{t-}\}} \tag{3}$$

where $\zeta \in (0,1]$ is an additive smoothing parameter guaranteeing that $\mu_t(\rho)$ is well-defined for $t < T_1^e$. For $\zeta = 0$, Eq. (3) defines the empirical rating distribution, which treats the observations $\varrho_1, \ldots, \varrho_{N^e_{t-}}$ as independent samples from a common distribution and ignores all other information contained in $\mathcal{F}_{t-}$. Our implementation assumes $\zeta = 0.5$, a value that can be justified on Bayesian grounds, see Box and Tiao (1992, pp. 34–36).

Based on a collection $\{N(\omega_i)\}$ of portfolio default histories, we estimate an intensity model $\lambda = \lambda^{\Theta}$, where $\Theta = (\kappa, c, \delta, \gamma, \lambda_0)$ is the parameter vector. Using the Nelder–Mead algorithm, we address the log-likelihood problem

$$\sup_{\Theta} \int_0^{\tau} \sum_i (\log \lambda_{s-}^{\Theta}(\omega_i)\, dN_s(\omega_i) - \lambda_s^{\Theta}(\omega_i)\, ds) \tag{4}$$

where the initial parameter set is selected by a grid search on $(0,2) \times (0,1) \times (0,2)^2 \times (0,20)$. The results of Ogata (1978) imply that the associated maximum likelihood estimator of Θ is consistent, asymptotically normal, and efficient.

We remark that the intensity model (1) and the method proposed to estimate it implicitly assume that a defaulted portfolio constituent is replaced with a name that has the same characteristics (as measured by the rating) as the defaulter. This is without loss of generality because we can extend the reach of a fitted intensity model to portfolios without replacement by applying an additional thinning step. This step follows the same approach as the rejection sampling used to generate portfolio default histories. It is detailed in Giesecke and Kim (2011) (Section 2.5).

2.4. Parameter estimates and goodness-of-fit

We apply the estimation method described above to the CDX HY and IG index portfolios. As of October 2007, the on-the-run HY index consists of 2 Baa, 48 Ba, 36 B and 14 C rated firms (100 constituents). The on-the-run IG index consists of 3 Aaa, 3 Aa, 51 A, 64 Baa and 4 Ba rated firms (125 constituents). Table 1 reports the maximum likelihood estimates of Θ and the corresponding standard errors, based on 10 K resampling scenarios for the HY index and 200 K resampling scenarios for the IG index. We use a larger number of IG resampling scenarios to address the rareness of defaults of investment-grade rated firms in the Moody's data base. The parameter δ governing the jump magnitude of the IG intensity is especially hard to pin down, as it must be inferred from relatively few default observations. Nevertheless, the fitted δ is significant at the 92% level.[5]

We use a time-scaling test to evaluate the fit of our intensity model (1). The test is based on a result of Meyer (1971), which implies that a portfolio default history $N(\omega_i)$ can be transformed into a standard Poisson process by a change of time given by the cumulative intensity, $\int_0^t \lambda_s(\omega_i)\, ds$. If the intensity is correctly specified, then the time-scaled event dates form a realization of a standard Poisson process in the time-scaled filtration, which we test using a Kolmogorov–Smirnov (KS) test. The KS test addresses the deviation of the empirical distribution function of the time-scaled inter-arrival times from their theoretical standard exponential distribution function.

We consider the p-values of the KS tests across all resampling scenarios ω_i. For the HY portfolio, the mean p-value is 0.297 and the standard deviation is 0.0071, with all p-values exceeding 0.1. For the IG portfolio, the mean p-value is 0.253 and the standard deviation is 0.1507, with 91% of the p-values exceeding 0.1. We conclude that the null hypotheses of exponentially distributed inter-arrival times after time change cannot be rejected for any HY resampling scenario and the

[5] As an alternative to increasing the number of resampling scenarios, one can apply an importance sampling scheme in order to improve the accuracy of the estimators. Here the thinning process Z is scaled up, leading to (artificial) resampled histories with a larger number of events. The likelihood function used to estimate the intensity parameters must be adjusted to account for the scaling.

Author's personal copy

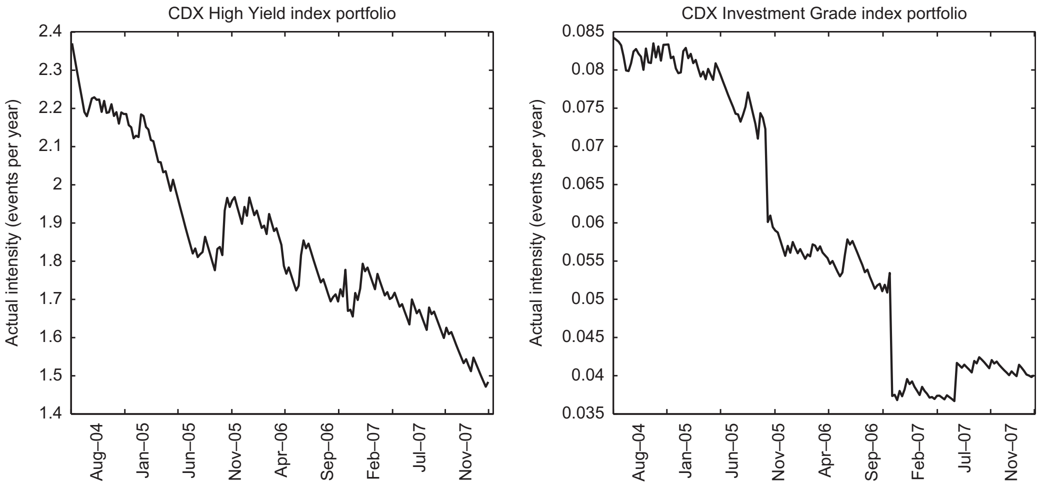

Fig. 2. Fitted mean portfolio default intensities, measured in events per year.

majority of IG resampling scenarios, at a 10% confidence level. This, in turn, indicates the goodness-of-fit of our intensity model (1) to the alternative CDX portfolio default histories.

2.5. Fitted intensity and loss distribution

Fig. 2 graphs the fitted mean intensities λ, averaged over resampling scenarios, from August 2004, when the indexes were established, to October 2007. The fitted λ account for revisions of the portfolio composition ("index rolls"), which take place every 6 months.[6]

There is a clear downward trend in the fitted intensities for both portfolios. This trend echoes the decline of default rates of Moody's rated firms between 2004 and 2007, as indicated in Fig. 1. Default rates reach historically low levels during October 2007, and so do the fitted portfolio intensities. The intermediate run-up in the HY portfolio intensity in the fourth quarter of 2005 is associated with the defaults of three Caa-rated industrial firms during August and the defaults of Delta Air Lines (C-rated) and Northwest Airlines (Caa-rated) on September 14, 2005. The large downward jumps in the IG portfolio intensity are due to index rolls. In the 9/20/2005 roll, two Ba-rated names were replaced by Baa-rated names, and in the 9/20/2006 roll, three Ba-rated names were replaced by Baa-rated names. The replacement of lower-quality names led to a reduction of risk.

The fitted portfolio intensities shown in Fig. 2 are associated with portfolio loss distributions that imply relatively little potential for significant default clustering. Fig. 3 graphs the model-implied conditional distribution of future loss in the CDX HY portfolio for a 5 year horizon, for conditioning times between April 2004 and October 2007. The fitted probability of a loss of 30% or more of the total portfolio notional during the 5 year period ahead is essentially zero at each conditioning date.

3. Risk-neutral intensity from CDX market rates

Having estimated models of CDX portfolio default timing under the actual probability P, in this section we use market prices of index and tranche swaps referenced on the CDX to estimate models of CDX portfolio default timing under a risk-neutral pricing measure. The fitted risk-neutral models reflect investors' assessment of correlated default risk. We are interested in how this assessment differs from actual default risk.

We assume that there are no arbitrage opportunities or market frictions in the CDX market. Then, under mild technical conditions, there exists a risk-neutral probability that is equivalent to the actual probability P. We fix a risk-neutral measure P^* with respect to a constant risk-free interest rate $r>0$.[7] The change of measure will be made precise as we proceed.

3.1. Index and tranche swaps

We consider contingent claims on the CDX portfolio default process N or the portfolio loss process L defined by $L_t = \ell_1 + \cdots + \ell_{N_t}$, where ℓ_n is the random loss at the nth default. An *index swap* is based on a portfolio whose C constituent single-name credit swaps have common notional that we normalize to 1, common maturity date T and common quarterly

[6] For the roll dates and the list of constituents that were exchanged at each roll see the website www.markit.com of Markit, the calculation agent for the CDX.

[7] The assumption of a constant r is common in the empirical credit literature, see Driessen (2005), Pan and Singleton (2008) and others.

Author's personal copy

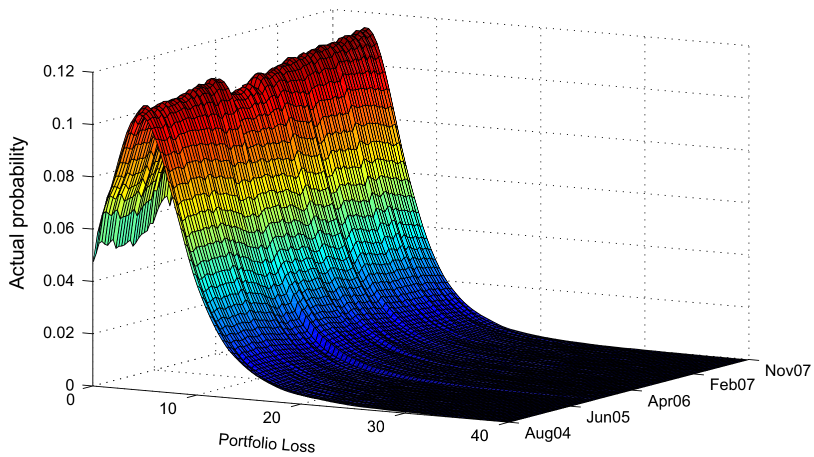

Fig. 3. Kernel-smoothed conditional distribution of future loss in the CDX High Yield portfolio for a 5 year horizon, computed weekly during the sample period based on the fitted intensity models. The loss at a default was drawn from a uniform distribution over {0.34,0.84}, with an expected position loss of roughly 60%.

premium payment dates (t_m). The protection seller agrees to cover portfolio losses as they occur. The value D_t at time $t \le T$ of these payments is given by the discounted cumulative losses. By integration by parts,

$$D_t = e^{-r(T-t)}E_t^*(L_T) - L_t + r\int_t^T e^{-r(s-t)}E_t^*(L_s)\,ds \tag{5}$$

where E_t^* denotes conditional expectation under the risk-neutral measure P^* with respect to the information set $\mathcal{F}_t$. The protection buyer agrees to make a stream of premium payments at dates (t_m). The cash flow at t_m is a fraction I of the total notional on the names that have survived until t_m. Neglecting premium accruals, the value at time $t \le T$ of the premium payments is given by

$$P_t(I) = I\sum_{t_m \ge t} e^{-r(t_m-t)}c_m E_t^*(C - N_{t_m}) \tag{6}$$

where c_m is the appropriate day count fraction for the period m. The fair index swap spread at time t is the solution $I = I_t$ to the equation $D_t = P_t(I)$. The spread depends only on expected defaults and losses for horizons in $(t,T]$.

Investors seeking narrower risk profiles can trade tranches referenced on the CDX. A *tranche swap* is specified by a lower attachment point $\underline{K} \in [0,1]$ and an upper attachment point $\overline{K} \in (\underline{K},1]$. The product of the difference $K = \overline{K} - \underline{K}$ and the portfolio notional C is the tranche notional. The tranche protection seller agrees to cover portfolio losses as they occur, given that the cumulative losses are larger than $\underline{K}C$ but do not exceed $\overline{K}C$. The cumulative payments at time t, denoted U_t, are

$$U_t = (L_t - \underline{K}C)^+ - (L_t - \overline{K}C)^+$$

The value of these payments at time t is

$$D_t = e^{-r(T-t)}E_t^*(U_T) - U_t + r\int_t^T e^{-r(s-t)}E_t^*(U_s)\,ds \tag{7}$$

This formula is analogous to formula (5) for the value of an index swap default leg. The latter can be viewed as the default leg of a tranche swap for which $\underline{K} = 0$ and $\overline{K} = 1$. The premium payments of the tranche protection buyer consist of two parts. The first part is an upfront payment, which is expressed as a fraction R of the tranche notional KC. The second part is a stream of payments at dates (t_m). For a tranche with $\overline{K} < 1$, the cash flow at t_m is a fraction S of the difference between the tranche notional and the tranche loss at t_m. Neglecting accruals, the value of the premium leg is given by

$$P_t(\underline{K},\overline{K},R,S) = RKC + S\sum_m e^{-r(t_m-t)}c_m(KC - E_t^*(U_{t_m})) \tag{8}$$

For a fixed upfront payment rate R, the fair tranche spread S is the solution $S = S_t(\underline{K},\overline{K},R)$ to the equation $D_t(\underline{K},\overline{K}) = P_t(\underline{K},\overline{K},R,S)$. Similarly, for a fixed tranche spread S, the fair tranche upfront rate R is the solution $R = R_t(\underline{K},\overline{K},S)$ to the equation $D_t(\underline{K},\overline{K}) = P_t(\underline{K},\overline{K},R,S)$. The fair spread and upfront rate depend only on the value of call spreads on the portfolio loss L_s with strikes $\underline{K}$ and $\overline{K}$ and maturities $s \in (t,T]$.

From the perspective of the protection seller, the mark-to-market value of an index or tranche swap is given by the difference between the value of the premium leg and the value of the default leg. The mark-to-market value to the protection buyer is the negative of the mark-to-market value to the seller. It is 0 at swap inception, when the swap rate is determined. After inception, the mark-to-market value changes in response to constituent defaults and changes in the constituent single-name swap rates.

Author's personal copy

3.2. Risk-neutral portfolio intensity

In order to value CDX index and tranche swaps, we model the default and loss processes under the risk-neutral measure P^*. The default process N is specified in terms of an intensity λ^* satisfying $\int_0^t \lambda_s^* \, ds < \infty$ for all t, which represents the conditional portfolio default rate with respect to P^*. The risk-neutral intensity λ^* is the counterpart to the actual intensity λ of N.[8] Modeling λ^* for a given specification of λ amounts to identifying the Radon–Nikodym derivative that defines the change of measure for N.

We develop a maximum likelihood approach to estimating the risk-neutral intensity λ^*. This approach requires the specification of the dynamics of λ^* under both P^* and P. We suppose that under actual probabilities, λ^* follows the jump-diffusion process

$$d\lambda_t^* = \overline{\kappa}^*(\overline{c}^* - \lambda_t^*)\, dt + \sigma^* \sqrt{\lambda_t^*}\, dW_t + \delta^* \, dL_t \tag{9}$$

where $\overline{\kappa}^* \geq 0$, $\overline{c}^* > 0$, $\sigma^* \geq 0$, and $\delta^* \geq 0$ are parameters satisfying the Feller condition $2\overline{\kappa}^*\overline{c}^* \geq \sigma^{*2}$, W is a standard Brownian motion relative to P, and L is the portfolio loss process. The market price of risk ξ_t underlying the change of measure from P to P^* for W is assumed to be a linear function of $\sqrt{\lambda_t^*}$:

$$\xi_t = \eta \sqrt{\lambda_t^*} \tag{10}$$

This market price of risk allows the mean-reversion rate of λ^* to differ across P and P^*, while guaranteeing that λ^* follows jump-diffusion processes with similar features under both measures.[9] Specifically, under the risk-neutral measure P^* associated with ξ_t,

$$d\lambda_t^* = \kappa^*(c^* - \lambda_t^*)\, dt + \sigma^* \sqrt{\lambda_t^*}\, dW_t^* + \delta^* \, dL_t \tag{11}$$

where $\kappa^* = \overline{\kappa}^* + \eta\sigma^*$, $c^* = \overline{\kappa}^*\overline{c}^*/\kappa^*$ and W^* is a standard Brownian motion with respect to P^* satisfying $dW_t^* = dW_t + \eta\sqrt{\lambda_t^*}\, dt$. It is the specification (11) of λ^* under P^* that is relevant for the valuation of index and tranche swaps.

The risk-neutral portfolio intensity λ^* follows similar jump-diffusion processes under both measures. The diffusive variation in λ^* is governed by σ^*. The jumps in λ^* occur at defaults. Since the frequency of defaults is different under the two measures, so is the frequency of the jumps in λ^*: they arrive at rate λ under P and rate λ^* under P^*. The magnitude of the jumps is inversely proportional to the realized recovery rate (one minus the loss rate) at an event. The lower the recovery, the larger the jump of the intensity at an event. This specification addresses the negative correlation between recovery rates and default rates found by Altman et al. (2005) and others.

The random loss rate ℓ_n is assumed to be governed by the same distribution under both measures. That is, there is no premium for recovery risk. This simplifying assumption is standard in the literature on default risk premia. In view of the calibration results in Giesecke and Kim (2007), who fit the specification (11) to market index and tranche rates observed on a single day using a least-squares criterion, we take ν to be uniform on $\{0.34, 0.84\}$. While allowing for random variation in recovery rates, this specification implies a mean loss rate of approximately 0.6, consistent with the industry standard assumption of 40% fixed (non-random) recovery at default.

Our square-root jump-diffusion models of risk-neutral default arrival rates extend the square-root diffusion models widely used in the literature on single-name default timing, for example Longstaff et al. (2005). In the multi-name setting, the jump-diffusion formulation addresses two sources of default clustering and several related effects. First, the portfolio constituents are exposed to a common Brownian risk factor. The movements in this factor generate systematic variation in arrival intensities. Second, portfolio intensities are influenced by defaults. The intensity jumps at a default, with a magnitude that depends on the realized loss. This self-exciting specification is motivated by the empirical results of Collin-Dufresne et al. (2009), who find that defaults tend to generate jumps in credit spreads across the board. The response of the risk-neutral portfolio intensity to a default models this effect. It is this response behavior that differentiates our specification from the square-root jump-diffusion intensity models of Duffie and Garleanu (2001), Eckner (2007) and others. In these alternative models, the intensity does not respond to defaults, but jumps according to a Poisson process with deterministic rate.[10]

The jump-diffusion behavior of the risk-neutral intensity generates two important features of index and tranche mark-to-market values. The diffusive variation in λ^* induces the diffusive volatility in mark-to-market values observed between defaults. This volatility reflects the systematic co-movement of the constituent risks. When a constituent name defaults, the risks of the surviving names are re-evaluated as implied by the response jump in λ^*, and this in turn triggers an

[8] Artzner and Delbaen (1992) show that if there is an intensity under P, then there is also an intensity under any measure equivalent to P.

[9] We do not opt for the more comprehensive "extended affine" market price of risk specification proposed by Cheridito et al. (2007), largely because our sample period is relatively short.

[10] Our specification can be extended to include additional jump and diffusion terms, as well as time-dependent coefficient functions. Extensions along these lines do not reduce the computational tractability of index and tranche swap pricing as developed in Section 3.3: the loss process transform is still of the form (12). We have found that additional terms do not substantially improve model fit.

Author's personal copy

adjustment in the mark-to-market value of an index or tranche position. This jump in mark-to-market values is of concern to investors, in addition to any diffusive mark-to-market volatility between defaults.

3.3. Valuation

Before addressing the valuation of index and tranche swaps under the risk-neutral intensity model (11), we remark that (11) specifies a default process N that does not terminate at the Cth default in the portfolio. This feature is, however, innocuous because for the relatively large CDX portfolios of interest, the probability of N exceeding C during standard contract terms tends to be small for our sample. The advantage of the model (11) is its computationally tractability. Since the processes N and L are affine point processes in the sense of Errais et al. (2010), the characteristic function of $(N,L)^\top$ is an exponentially affine function of the state. We have the formula

$$E^*_t(\exp(iv(L_T-L_t)))=\exp(\alpha(t)+\beta(t)\lambda^*_t) \tag{12}$$

where $t\leq T$, i is the imaginary unit, v is a real number and the coefficient functions $\alpha(t)=\alpha(v,t,T)$ and $\beta(t)=\beta(v,t,T)$ solve the ordinary differential equations

$$\partial_t\beta(t)=1+\kappa^*\beta(t)-\tfrac{1}{2}(\sigma^*)^2\beta(t)^2-q(\delta^*\beta(t),v) \tag{13}$$

$$\partial_t\alpha(t)=-c^*\kappa^*\beta(t) \tag{14}$$

with boundary conditions $\beta(T)=\alpha(T)=0$ and jump transform

$$q(u,v)=\int e^{(iv+u)z}\,d\nu(z)$$

where ν is the distribution of the loss at default ℓ_n, and u is any complex number such that $q(u,v)$ is finite for a given real v. The conditional characteristic function of N_T-N_t is given by the right hand side of Eq. (12), with coefficient functions $\alpha(t)$ and $\beta(t)$ satisfying Eqs. (13) and (14), where $q(u,v)$ is given by $\int e^{iv+uz}\,d\nu(z)$.

Conditional expected defaults $E^*_t(N_T)$ and losses $E^*_t(L_T)$ are obtained by differentiating the corresponding characteristic functions. We obtain closed form expressions for these expectations. From Eqs. (5) and (6), we then get a closed form expression for the model index swap spread I_t. This expression is approximate to the extent that it ignores the small probability that the portfolio default process N defined by the intensity model (11) exceeds C. An exact but non-analytical expression for the index spread can be obtained by stopping N and L at the Cth event.

Model tranche rates do not take a closed form. We first obtain the conditional loss distribution by Fourier inversion of the characteristic function (12). The price of an option on L is obtained by integrating the option payoff against the loss distribution. The option price determines the tranche rate through formulae (7) and (8).

3.4. Likelihood estimators

We fix a set $\{0,1,\ldots,\tau\}$ of observation dates. The data at date t consist of (I_t,S_t,N_t,L_t), where I_t is the mid-market index spread and $S_t=(S_{1t},\ldots,S_{nt})$ collects the mid-market rates of n tranches. All rates are with respect to a common maturity. The Markov property of the process (λ^*,N,L) under P^* along with the index swap pricing formulae developed in Section 3.1 imply that there is a function $G:\mathbb{R}_+\times\mathbb{N}\to\mathbb{R}_+$ that is increasing and continuously differentiable in its first argument such that

$$I_t=G(\lambda^*_t,N_t;\theta) \tag{15}$$

where $\theta=(\kappa^*,c^*,\sigma^*,\delta^*,\eta)$ is a vector of parameters to be estimated. The model index spread at t does not depend on L_t, but only on the risk-neutral expected loss at an event. At a given (θ,N_t,I_t), the model-implied risk-neutral intensity at t is thus given by

$$\lambda^*_t=G(\cdot,N_t;\theta)^{-1}(I_t)=:H(I_t,N_t;\theta) \tag{16}$$

where $H:\mathbb{R}_+\times\mathbb{N}\to\mathbb{R}_+$. While the spread on the liquid index is measured without error, we suppose the rates on the less liquid tranches are corrupted with Gaussian noise.[11] More precisely, for functions $F_k:\mathbb{R}_+\times\mathbb{R}_+\to\mathbb{R}_+$ implied by the tranche pricing formulae developed in Section 3.1, we take $S_{kt}=F_k(\lambda^*_t,L_t;\theta)+\epsilon_{kt}$ for tranches $k=1,\ldots,n$. Here, ϵ_{kt} is normally distributed with mean zero and standard deviation $v_k|\mathrm{Bid}_{kt}-\mathrm{Ask}_{kt}|$ for some $v_k\geq 0$, where Bid_{kt} is the bid price of tranche k at t and Ask_{kt} is the corresponding ask price. The variables ϵ_{kt} are independent of one another across tranches and observation dates. The specification of additive Gaussian observation noise with time-varying variances that depend on the bid/ask spread follows Pan and Singleton (2008). It allows for the possibility that the fits deteriorate during periods of market turmoil when bid/ask spreads widen substantially. We let $v=(v_1,\ldots,v_n)$.

[11] While the bid/ask spreads for the index are typically much less than one basis point for most days in the sample, the bid/ask spreads for the tranches tend to be substantially higher.

Author's personal copy

The index spread is the transformed value of the latent risk-neutral intensity, with the index pricing function (15) defining the transformation. Using a change of variable argument, we obtain the log-likelihood function of the data as

$$\begin{aligned}\mathcal{L}(\theta,v) = & \sum_{t=1}^{\tau} \log g_t(H(I_t,N_t;\theta);\theta|H(I_{t-1},N_{t-1};\theta),N_{t-1},L_{t-1},N_t,L_t) \\ & + \log g_0(H(I_0,N_0;\theta);\theta|N_0,L_0) + \sum_{t=1}^{\tau} \log |\partial_1 H(I_t,N_t;\theta)| + \mathcal{L}_{N,L}(\Theta) \\ & + \sum_{t=1}^{\tau}\sum_{k=1}^{n} \log \phi(S_{kt} - F_k(H(I_t,N_t;\theta),L_t;\theta); v_k^2(\text{Bid}_{kt} - \text{Ask}_{kt})^2)\end{aligned}$$

where $g_t(\cdot;\theta|\lambda^*_{t-1},N_{t-1},L_{t-1},N_t,L_t)$ is the conditional density of λ^*_t and $g_0(\cdot;\theta|N_0,L_0)$ is the conditional density of λ^*_0, both under P. Note from Eq. (9) that λ^* follows a P-Feller diffusion between events. The jumps in λ^* arrive with P-intensity λ. Therefore, if $N_{t-1} = N_t$ then $g_t(\cdot;\theta|\lambda^*_{t-1},N_{t-1},L_{t-1},N_t,L_t)$ is the non-central chi-squared density. If $N_{t-1}+1 = N_t$, which occurs four times in the HY sample, we suppose that the default occurs exactly at $t-1$, allowing us to use the non-central chi-squared density once again. This treatment is reasonable given our daily observations. The function $\partial_1 H(\cdot,\cdot;\theta)$ is the partial derivative of $H(\cdot,\cdot;\theta)$ with respect to its first argument. It defines the third term of $\mathcal{L}(\theta,v)$, which reflects the fact that the risk-neutral intensity is extracted from the index spread observation. The fourth term $\mathcal{L}_{N,L}(\Theta)$ represents the log-likelihood function of the vector $(N_0,L_0,\ldots,N_\tau,L_\tau)$, which is determined by the P-intensity model λ and the distribution v of the loss at events. Finally, the last term of the log-likelihood function accounts for the noisy tranche rate observations, where $\phi(\cdot;V)$ is the density of a Gaussian random variable with mean zero and variance V.

The likelihood problem has features that are unique to the multi-name setting. While a credit swap referenced on an individual issuer expires at default, index and tranche swaps continue to trade after the reference portfolio suffers defaults. Model index and tranche rates depend on the number of events and the cumulative loss induced by these events. Therefore, the likelihood function $\mathcal{L}$ contains the density of the portfolio default and loss processes relative to the actual measure P. This density is represented by the term $\mathcal{L}_{N,L}$, which is governed by the specification of the actual intensity λ. Thus, one could estimate λ together with the risk-neutral intensity λ^* by maximizing $\mathcal{L}$ with respect to the parameters (Θ,θ) of the models for λ and λ^*. However, this strategy is not appropriate for our sample due to the low number of CDX defaults during the sample period 2004–2007. The failure of this strategy motivates our two-step procedure, in which the estimation of the actual CDX intensity λ is based on historical default experience between 1970 and 2007. The risk-neutral CDX intensity λ^* is estimated from index and tranche price data in a subsequent step, taking the fitted λ as given.

The maximum likelihood estimator of the parameter vector (θ,v) solves $\sup_{\theta,v}\mathcal{L}(\theta,v)$. The `fmincon` routine of MATLAB is used to numerically solve this problem. The initial parameter set is found through a grid search algorithm. Note that the term $\mathcal{L}_{N,L}(\Theta)$ does not depend on (θ,v), and can therefore be ignored in the maximization.

3.5. Data and parameter estimates

We implement the estimators for the CDX High Yield and Investment Grade portfolios. Our data set consists of daily closing bid and ask quotes of spreads and upfront rates for on-the-run index and tranche contracts of all traded maturities. It covers virtually the entire trading history of the two CDX indexes through November 29, 2007. For the HY index, the available maturities are 3, 5 and 7 years. The sample period for the 5 year contract begins on August 3, 2004, for the 3 year contract it begins on June 1, 2006 and for the 7 year contract it begins on September 12, 2006. For the IG index, the available maturities are 5, 7 and 10 years. The sample period for the 5 year contract begins on April 29, 2004, for the 7 year contract it begins on June 22, 2005 and for the 10 year contract it begins on May 10, 2004. The 5 year contracts are the most liquid.

The HY index has tranche attachment points 0%, 10%, 15%, 25%, 35% and 100%, while the IG index has attachment points 0%, 3%, 7%, 10%, 15%, 30% and 100%. Our data set does not contain prices of the 35–100% and 30-100% super senior tranches. The prices of the 0–10% and 10–15% HY tranches are quoted in terms of an upfront rate, while the price of the 0–3% IG tranche is quoted in terms of an upfront rate plus a spread of 500 basis points paid quarterly. All other tranches are quoted in terms of a running spread that is paid quarterly. We use daily mid-market quotes in the estimation. In order to maximize the amount of data used in the estimation, we ignore index rolls.[12]

In view of the distinct observation periods, and to obtain a term structure perspective, we perform the estimation separately for each available contract maturity. Table 2 reports the corresponding maximum likelihood parameter estimates, along with asymptotic standard errors.[13] The parameter estimates are broadly consistent across maturities. The fitted value of δ^* is positive and highly significant. It indicates that the response of the risk-neutral intensity λ^* to a default is significant, after controlling for the influence of the Brownian risk factor. This finding corroborates and complements the

[12] Index rolls can be addressed by fitting separate risk-neutral intensity models for each of the 6 month roll periods. This, however, may degrade the accuracy of the maximum likelihood estimators.

[13] The estimates are based on a risk-free rate of $r = 5\%$.

Author's personal copy

Table 2
Maximum likelihood estimates of the parameters of the risk-neutral intensity λ^* and the tranche pricing error volatility multipliers ν_k. Estimates of asymptotic standard errors are given parenthetically. The "Mat" column indicates the contract maturity. The "Obs" column reports the number of daily observations used.

Idx	Mat	Obs	κ^*	c^*	σ^*	δ^*	η	ν_1	ν_2	ν_3	ν_4	ν_5
	3Y	242	4.85	0.94	3.02	4.18	0.59	5.59	3.95	3.61	5.43	–
			(0.09)	(0.04)	(0.01)	(0.03)	(0.07)	(0.12)	(0.08)	(0.15)	(0.07)	
HY	5Y	651	4.95	1.50	3.86	4.62	0.77	7.76	7.17	8.60	7.53	–
			(0.08)	(0.01)	(0.02)	(0.03)	(0.05)	(0.10)	(0.05)	(0.47)	(0.09)	
	7Y	183	5.10	1.53	3.95	4.94	0.71	14.21	5.07	10.04	7.48	–
			(0.08)	(0.02)	(0.01)	(0.04)	(0.15)	(0.44)	(0.10)	(0.59)	(0.66)	
	5Y	804	5.37	0.25	1.65	5.48	0.31	18.10	26.46	9.04	3.07	7.75
			(0.02)	(0.02)	(0.04)	(0.02)	(0.03)	(0.64)	(0.25)	(0.18)	(0.07)	(0.08)
IG	7Y	515	4.62	0.41	1.95	4.15	0.42	11.90	29.59	13.72	10.01	12.30
			(0.05)	(0.03)	(0.05)	(0.02)	(0.04)	(0.15)	(0.26)	(0.59)	(0.05)	(0.06)
	10Y	685	4.92	0.49	2.20	4.64	0.45	34.78	15.04	37.66	10.02	10.05
			(0.02)	(0.03)	(0.02)	(0.02)	(0.03)	(0.36)	(0.24)	(0.38)	(0.02)	(0.03)

results of Collin-Dufresne et al. (2009), who show that credit events of large firms generate a market-wide response in credit spreads, consistent with a response of the risk-neutral intensity to defaults. It also suggests that the jumps in CDX index and tranche mark-to-market values that occur at defaults are significant.

3.6. *Tranche pricing errors*

The estimation assumes that index spreads are measured without error. Therefore, index spreads are matched perfectly. Tranche rates are corrupted with noise, so they are not fitted exactly. We compare market tranche rates with model-implied rates $F_k(\lambda^*_t, L_t; \hat{\theta})$, where λ^*_t is the fitted risk-neutral intensity obtained from formula (16), and $\hat{\theta}$ is the maximum likelihood estimator of θ. Fig. 4 shows the time series of observed and fitted 5 year HY index and tranche rates. For the purpose of calculating the model-implied rates, we assume a realized recovery rate of 40% for the four HY defaults in the sample, consistent with standard industry practice and our model for the loss at default. Fig. 5 graphs the time series for the 5 year IG index and tranche rates.

Our square-root jump-diffusion models of λ^* capture most of the time-series variation of tranche rates during 2004–2007. To provide some perspective on model fit, Fig. 6 provides histograms of the standardized pricing errors $(S_{kt} - F_k(\lambda^*_t, L_t; \hat{\theta}))/(\nu_k|\text{Bid}_{kt} - \text{Ask}_{kt}|)$ across all fitted 5 year tranches, along with the theoretical (under P) standard normal density. The HY errors are relatively symmetrically distributed and somewhat leptokurtic. The slightly positive sample mean 0.067 indicates that on average, the model tends to slightly understate HY tranche rates. The distribution of the IG errors is somewhat skewed to the left. This skew is largely due to the systematic underpricing of the 15–30% tranche indicated in Fig. 5.

3.7. *Fitted intensity and loss distribution*

Fig. 7 graphs the fitted risk-neutral default intensities λ^* for the HY and IG portfolios. Given a pair of observations (I_t, N_t), the value λ^*_t is obtained from formula (16) evaluated at the maximum likelihood estimator $\hat{\theta}$ of θ for the 5 year swaps. There were no defaults in the IG portfolio so $N_t = 0$ for all observation dates t. The HY portfolio experienced four defaults during the sample period: Collins & Aikman Products on 05/17/05, Delphi on 10/8/05, Calpine on 12/20/05 and Dana on 03/01/06. The default of Collins & Aikman was accompanied by the downgrades of Ford and General Motors during the spring of 2005 and a broad widening of protection rates going beyond the auto industry. After a relatively quiet period, market rates again sharply increased during 2007 in response to signs of a credit crisis in the U.S. (see Figs. 4 and 5). These events correspond to the large run-ups in fitted risk-neutral intensities across index portfolios.

Fig. 8 shows the risk-neutral distribution of future HY portfolio loss for a 5 year horizon, computed weekly based on the parameter estimates for the 5 year contracts. The time-series variation of the risk-neutral distribution is much larger than that of the actual distribution of HY portfolio loss shown in Fig. 3. This volatility over time reflects investors' changing risk-neutral expectations regarding future default loss during the 5 year period ahead. The risk-neutral distributions have a much fatter tail than the actual distributions. This indicates that the market regarded default clusters as much more likely than what was implied by actual default experience, foreshadowing the importance of risk premia. The tails of the loss distribution were lightest during February 2007, when market rates reached their historical lows. The tails grew heavier during May 2005 and July–August 2007, showing that credit investors anticipated relatively strong default clustering during those periods. A direct comparison of the tails during May 2005 and July–August 2007 crises periods suggests that these fears played a relatively larger role during the summer of 2007. This finding is consistent with the results of Bhansali

Author's personal copy

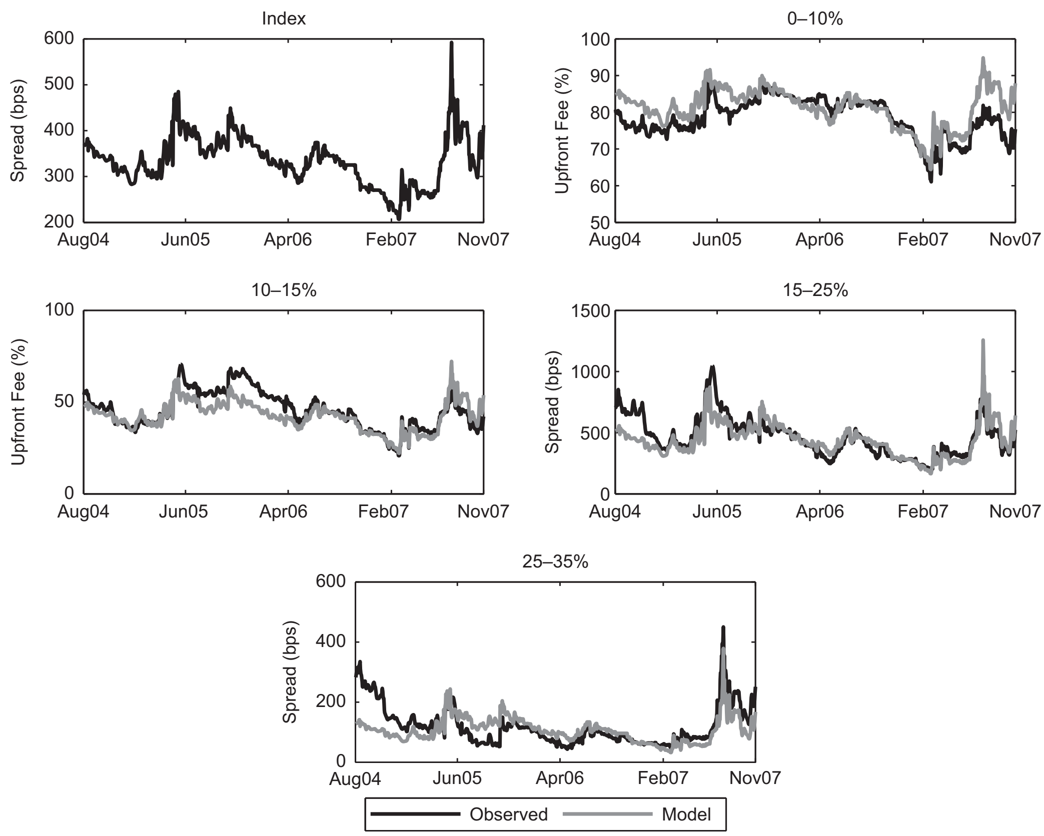

Fig. 4. Time series of observed and model-implied 5 year index and tranche rates for the CDX High Yield index portfolio. The model rates $F_k(\lambda_t^*, L_t; \hat{\theta})$ are based on the parameter estimates $\hat{\theta}$ for the 5 year contracts reported in Table 2.

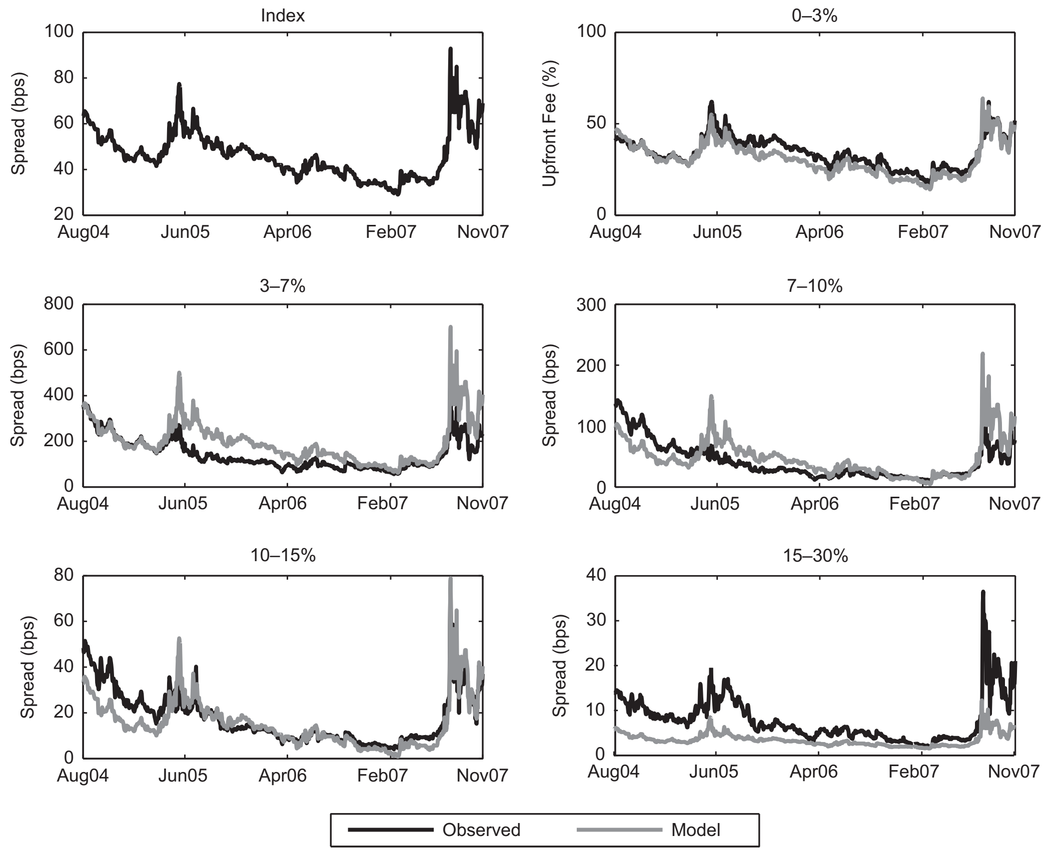

Fig. 5. Time series of observed and model-implied 5 year index and tranche rates for the CDX Investment Grade index portfolio. The model rates $F_k(\lambda_t^*, L_t; \hat{\theta})$ are based on the parameter estimates $\hat{\theta}$ for the 5 year contracts reported in Table 2.

Author's personal copy

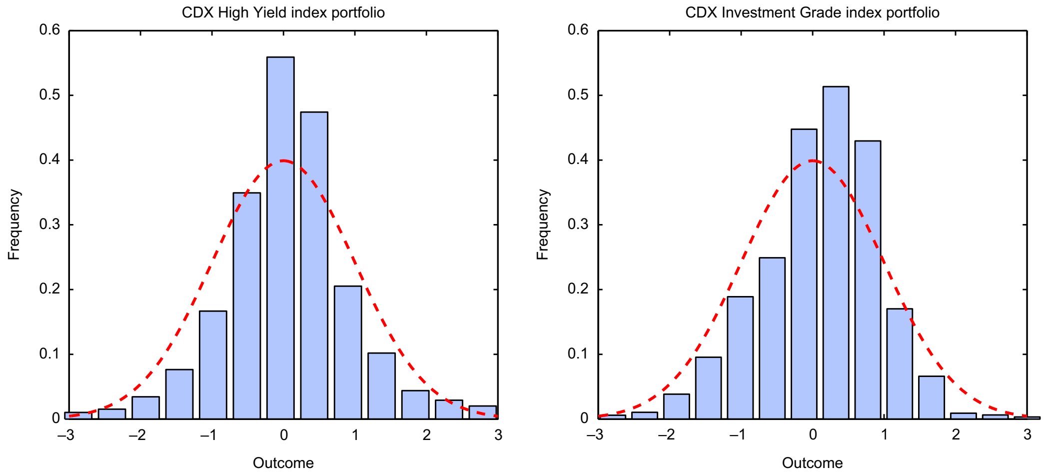

Fig. 6. Histograms of tranche pricing errors $(S_{kt}-F_k(\lambda_t^*,L_t;\hat{\theta}))/(\upsilon_k|\text{Bid}_{kt}-\text{Ask}_{kt}|)$ across all fitted 5 year tranches, and standard normal density (dashed). For the HY tranches, the mean is 0.067 and the standard deviation is 0.996. For the IG tranches, the mean is 0.089 and the standard deviation is 0.901.

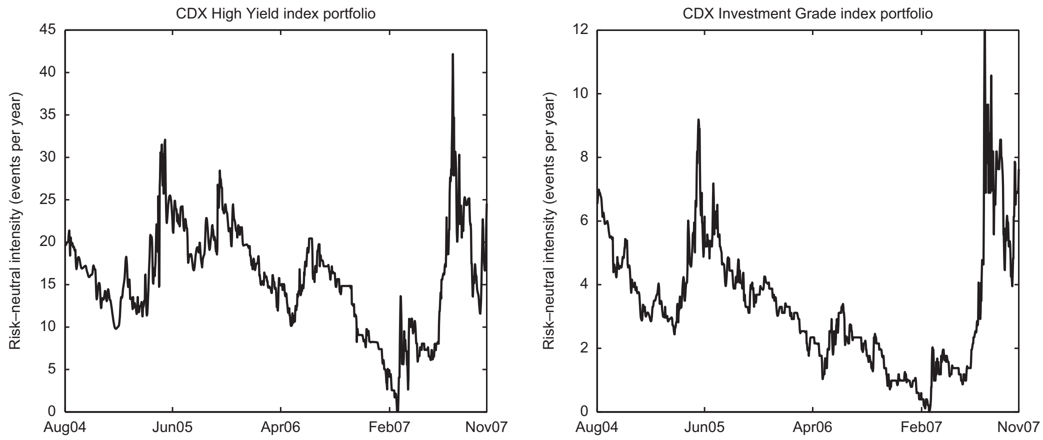

Fig. 7. Fitted risk-neutral portfolio default intensities, measured in events per year, based on the estimates obtained from the 5 year contracts reported in Table 2.

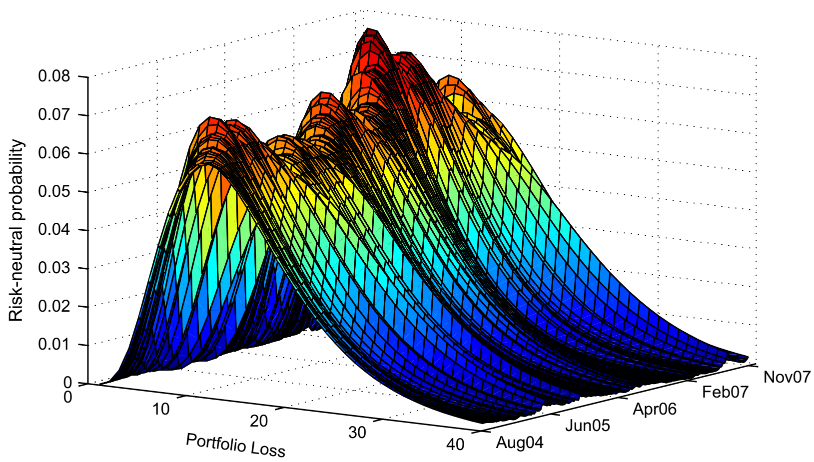

Fig. 8. Kernel-smoothed conditional risk-neutral distribution of future loss in the CDX High Yield portfolio for a 5 year horizon, L_{t+5}, computed weekly based on the parameter estimates obtained from the 5 year contracts reported in Table 2.

et al. (2008), who extract market expectations about systemic default risk from CDX index and tranche market rates. After controlling for idiosyncratic and sector-wide risk, they find that the 2007 crisis had more than twice the systemic default risk of the 2005 auto crisis, relative to risk-neutral probabilities.

Author's personal copy

4. CDX risk premia

The striking contrast between the processes followed by the actual and risk-neutral intensities, and between the parameters governing the actual and risk-neutral dynamics of the risk-neutral intensity, foreshadows the presence of default risk premia. This section analyzes these premia building on the results obtained in the previous sections.

In order to motivate our decomposition of the risk premia, note that an index or tranche investor is exposed to two distinct sources of risk. Swap mark-to-market values vary over time. Between defaults, they change in response to systematic changes in the constituent risks. When a constituent name defaults, they jump because the constituent risks are re-evaluated. This diffusive and jump volatility of mark-to-market values is referred to as *mark-to-market risk*. It is generated by the diffusive and jump volatility of the risk-neutral intensity λ^*. An index or tranche protection seller is also exposed to uncertainty regarding the timing of payouts due to defaults of constituent firms, called *jump-to-default risk*. Risk-averse investors demand compensation for bearing exposure to mark-to-market and jump-to-default risks. Our models of CDX portfolio default timing under actual and risk-neutral measures provide estimates of this compensation. The estimates incorporate the implications of default correlation among the CDX constituents, and the negative correlation between default rates and recovery rates.[14]

The components of the risk premia are related to the components of the Radon–Nikodym derivative $M_\tau = dP^*/dP$ describing the change of measure from actual to risk-neutral probabilities. Under mild technical conditions stated in Brémaud (1980), in our model of correlated default timing the Radon–Nikodym derivative process $(M_t)_{t \le \tau}$ defined by $M_t = E_t(M_\tau)$ takes the form $M_t = M_t(W)M_t(N)$, where

$$M_t(W) = \exp\left(-\eta \int_0^t \sqrt{\lambda_s^*}\, dW_s - \frac{\eta^2}{2}\int_0^t \lambda_s^*\, ds\right)$$

represents the Radon–Nikodym derivative process associated with the change of measure for the Brownian motion W, while the second factor

$$M_t(N) = \exp\left(\int_0^t \log\left(\frac{\lambda_{s-}^*}{\lambda_{s-}}\right) dN_s + \int_0^t (\lambda_s - \lambda_s^*)\, ds\right)$$

represents the Radon–Nikodym derivative process associated with the change of measure for the default counting process N.

4.1. Mark-to-market risk

The parameter η governs the change of measure for the Brownian motion W generating systematic, diffusive movements in the risk-neutral intensity λ^* and the associated diffusive mark-to-market volatility. It links the actual and risk-neutral dynamics of λ^* between default arrivals. A value $\eta = 0$ would imply that λ^* undergoes the same diffusive movements across P and P^*, and that diffusive mark-to-market risk is not priced.

The fitted values of η, reported in Table 2, are positive for all contracts and reference portfolios, indicating that diffusive mark-to-market risk plays an important economic role. While the standard errors of the estimates of η are relatively larger than those of the estimates of the other parameters, all estimates are significant at the 99.9% level. A positive η implies that the market price of risk (10) for W is positive, and that λ^* reverts faster to its long-run mean under P^* than under P. It also implies that the reversion level of λ^* is lower under P^* than under P. The adjustments to the speed and level parameters are more pronounced for the HY intensity than the IG intensity.

It remains to clarify the role of the jump volatility in mark-to-market values, which is caused by the upward jumps in λ^*. These jumps arrive at a much higher rate under P^* than under P. This is because λ^* tends to be much larger than λ, as seen in Figs. 2 and 7. The much higher jump frequency under P^* ensures that λ_t^* tends to be larger and exhibits fatter tails under P^* than under P, despite the increase of the mean-reversion speed and the reduction in the reversion level under P^* caused by the positive value of η. It is this contrast between the actual and risk-neutral distribution of the risk-neutral intensity λ^* that suggests that mark-to-market risk is priced.

4.2. Jump-to-default risk

The premium for jump-to-default risk is measured by the ratio of the risk-neutral portfolio intensity λ^* to the actual intensity λ. This ratio governs the change of measure for the default counting process N. The case $\lambda^*/\lambda = 1$ would imply that CDX index and tranche investors do not demand compensation for jump-to-default risk and jump volatility of mark-to-market values. Jarrow et al. (2005) provide conditions under which the ratio of risk-neutral to actual firm-level intensities is equal to 1. A sufficient condition is that there are infinitely many firms, all exposed to the same risk factor, and all defaulting independently conditional on the risk factor. In our portfolio model, the conditional independence

[14] The distinction between mark-to-market risk and jump-to-default risk is standard in the literature estimating firm-level default risk premia, see Driessen (2005) and Berndt et al. (2008). In contrast to this literature, our portfolio-level estimates address the effects of default clustering, and the premia for jump volatility in mark-to-market values due to defaults of portfolio constituents.

Author's personal copy

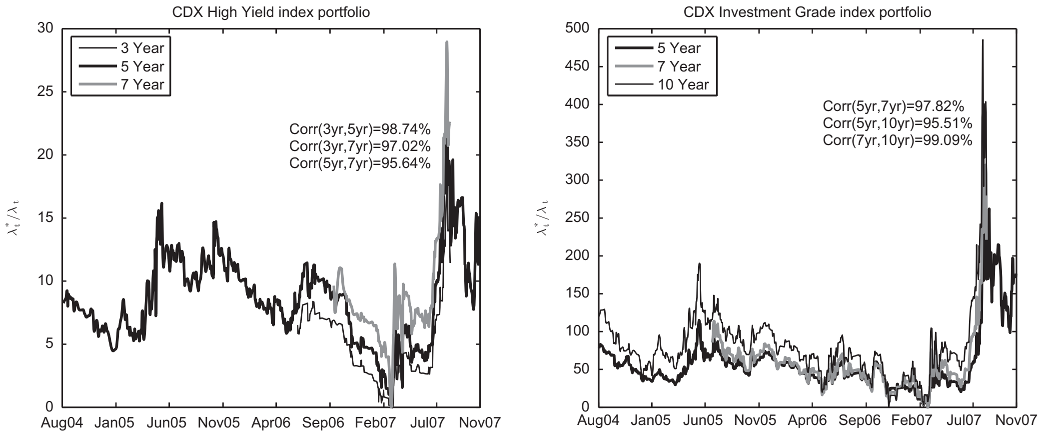

Fig. 9. Fitted ratios of risk-neutral portfolio intensity λ^* to actual portfolio intensity λ, for each of several contract maturities. Also indicated are the correlations between the time series for different maturities.

(doubly stochastic) assumption is violated because the portfolio intensity responds to defaults. Therefore, a default event may command a premium even if the portfolio consists of a large number of firms.

Fig. 9 graphs the fitted λ^*/λ for the HY and IG portfolios.[15] It indicates that CDX investors tend to demand substantial compensation for bearing exposure to correlated jump-to-default risk, with generally much higher jump-to-default premia for the IG portfolio of high-quality firms, for all maturities.[16] This finding complements the results of Driessen (2005), Berndt et al. (2008) and others, who measure single-name jump-to-default premia by the ratio of risk-neutral to actual firm-level default intensities. They find that jump-to-default premia implied by bond prices or market rates of single-name credit swaps are higher for high-quality firms than for low-quality firms.[17] Our portfolio-level results provide evidence that the correlated default risk of high-quality firms is harder to diversify and to hedge than the correlated default risk of low-quality names.

Jump-to-default premia tend to increase with maturity, and they vary dramatically over time. According to the estimates of the correlation coefficients indicated in Fig. 9, there is a high degree of co-movement between the series for different maturities. The time-series behavior of the premia has a natural interpretation in terms of economic events. After an initial downward trend, during the second quarter of 2005 premia increased quickly to peak in May 2005, marking the height of a crisis that is often linked to the downgrades of Ford and General Motors. The premia spiked again in September 2006, and then steadily declined to bottom out in late February 2007. The ratio λ^*/λ stayed below 1 for a few days only. During that period, protection sellers accepted negative jump-to-default premia. From this trough, premia surged rapidly to peak at historically extreme levels in early August 2007, reflecting the appearing signs of a credit crunch in the U.S. due to strings of mortgage defaults. The premia declined somewhat from these levels but remained relatively high through October 2007. This time-series behavior echoes a common view held by market participants that the risk of contagion from the mortgage market was perceived highest during July and August 2007. During these months, signs of a serious mortgage crisis appeared, and investors with exposure to corporate credit risk sought cover in the CDX market at any price, driving index and tranche rates to unprecedented levels. Actual corporate defaults remained at historically low levels, however.

It remains to clarify the role of liquidity for the relatively large jump-to-default premia we estimated during the summer of 2007. While during that period default protection buyers clearly preferred the CDX market to the largely illiquid single-name credit swap market, the number and size of typical CDX trades declined and bid/ask spreads widened, especially for the more senior tranches. However, our analysis of the time series of bid/ask spreads showed that during August 2007, bid/ask spreads for the IG contracts were not significantly wider and bid/ask spreads for the HY contracts were even lower than those measured during the auto crisis in May 2005, when jump-to-default premia stood at roughly one half of their value at the August 2007 peak. This observation suggests that liquidity problems are unlikely to have overly inflated the jump-to-default premia estimates for the summer of 2007.

[15] Our estimates of portfolio-level jump-to-default premia also provide information regarding the firm-level jump-to-default premia for a representative portfolio constituent. This is because the portfolio intensity is the sum of the constituent default intensities (see Giesecke et al. (in press) for a precise statement). For a relatively homogeneous portfolio, the default intensity of a representative constituent is given by the portfolio intensity divided by the number of constituents. Thus, the ratio λ^*/λ also represents the jump-to-default premium for a representative CDX constituent.

[16] The relatively high IG jump-to-default premia we estimate for the second half of 2007 do not appear to be excessive. They are broadly consistent with the estimates of Coval et al. (2009) for a representative CDX IG constituent firm, which are based on a sample period similar to ours. They provide estimates specific to individual index/tranche contracts ranging from 1.7 to 6106.6 (sample period averages).

[17] There are several reasons for this observation, and we thank Darrell Duffie for pointing these out to us. High-quality firms only fail in deep recessions. The risks of selling protection on high-quality names are highly negatively skewed and hard to diversify. Moreover, there are setup costs for OTC derivatives. It is generally not cost effective to trade for a few basis points unless the volume is huge.

Author's personal copy

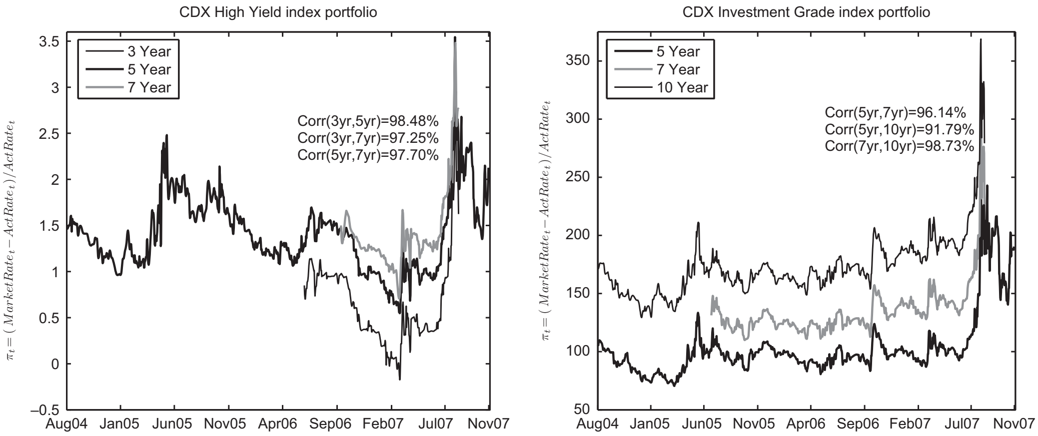

Fig. 10. Fitted effects of risk premia on index swap rates, measured by $\pi_t = (\text{MarketRate}_t - \text{ActRate}_t)/\text{ActRate}_t$, for each of several maturities. Also indicated are the correlations between the time series of π_t for different maturities.

4.3. The pricing of risks in CDX markets

We examine the joint implications of mark-to-market and jump-to-default risk premia for the pricing of CDX index and tranche swaps. To this end, we construct actuarial index and tranche rates, and then compare these with market rates. An actuarial rate ActRate_t is obtained by evaluating the expected discounted cash flows of a swap under actual probabilities. This amounts to replacing the risk-neutral expectations by actual expectations in the pricing formulae developed in Section 3.1, and assuming that defaults arrive at the actual intensity λ rather than the risk-neutral intensity λ^*. If CDX investors do not require any risk premia, then the actuarial swap rate should replicate the corresponding market rate. If investors do require risk premia, then the market rate should exceed the actuarial rate. In that case, the buyer of index or tranche default protection is willing to pay a premium for holding the contract, while the seller of protection demands a premium. To quantify the role of the premium, we report

$$\pi_t = \frac{\text{MarketRate}_t - \text{ActRate}_t}{\text{ActRate}_t} \tag{17}$$

Fig. 10 shows the fitted π_t for the HY and IG index swaps, for each of several maturities. Monte Carlo simulation is used to compute ActRate_t based on the dynamics of the actual intensity λ, estimated from default data up to time t. The simulation estimators incorporate the effects of index rolls, the corresponding changes in maturity dates, and the term structure of risk-free rates implied by the prevailing swap curve. They are based on 1 million Monte Carlo replications. The fitted π_t is significantly positive during most of the sample period, highlighting the presence of risk premia. It varies substantially across time, maturity, and index portfolio. The large run-ups in risk premia have natural interpretations in terms of financial market developments, as emphasized in the discussion of the jump-to-default premia. While π_t is higher for longer maturity swaps, the corresponding time series are highly correlated. The fitted π_t is larger for the IG swaps, indicating that IG investors demand greater compensation for bearing correlated default risk than HY investors. This finding is consistent with our earlier one regarding the behavior of jump-to-default premia.

We turn to a more in-depth exploration of the relationships between the fitted π_t and various measures of economic and financial market developments. Table 3 reports the results of univariate and multiple OLS regressions of the fitted weekly π_t for the 5-year CDX index swap. As regressors we consider the CBOE VIX volatility index, the credit spread between Moody's seasoned 30-year Baa corporate bond yield and the 10-Year Treasury yield, and the LIBOR-OIS spread, defined as the difference between the 3-month LIBOR and the Overnight Index Swap (OIS) rate. The VIX is a widely watched measure of event risk in credit markets, and highly correlated with investors' appetite for exposure to the high-yield bond credit class. We expect it to be positively correlated with risk premia. The credit spread provides a measure of credit risk in the U.S. The LIBOR-OIS spread is frequently used as an indicator of the health of the banking system. We expect it to be positively correlated with risk premia, especially during the second half of 2007, when fears about the stability of the U.S. banking system appeared. We have considered other variables, including the spread between the U.S. industrial 10 year BB yield and the 6-month Treasury bill yield, the trailing high-yield default rate, the level and trailing return on the S&P 500, the Treasury Constant Maturity Rate of several maturities, the TED spread, and the change of non-farm payrolls, but have found the aforementioned variables to have the most explanatory power if considered jointly.

The three variables explain a significant portion of the fitted effect π_t on CDX index rates. The R^2 for the HY index is 40%, while that for the IG index is 74%. The coefficient of each variable is of the expected positive sign. In the univariate regressions, the VIX has the greatest explanatory power for the HY π_t, while the LIBOR-OIS spread has the greatest explanatory power for the IG π_t. The multiple regressions indicate that, in the presence of the VIX and the credit spread, the influence of the LIBOR-OIS spread is insignificant in the case of the HY index. In the case of the IG index, the influence of the credit spread is insignificant in the presence of the VIX and the LIBOR-OIS spread. The significance of the LIBOR-OIS

Author's personal copy

Table 3
OLS regressions of the effects of risk premia on the 5Y CDX index swap rates, measured weekly by $\pi_t = (\text{MarketRate}_t - \text{ActRate}_t)/\text{ActRate}_t$, on the CBOE VIX index, the spread between the yield on Moody's seasoned 30-year Baa-rated bond and the 10-year Treasury constant maturity rate (CS), and the difference between the 3-month LIBOR and 3-month overnight index swap rate (LOIS), all measured in percentage terms. The t-statistics are shown parenthetically. The variables were obtained from Bloomberg.

CDX High Yield Index				CDX Investment Grade Index			
VIX	CS	LOIS	R^2	VIX	CS	LOIS	R^2
0.067 (9.356)			0.343	6.352 (16.429)			0.616
	1.060 (6.366)		0.194		66.779 (5.496)		0.152
		1.107 (6.029)	0.178			150.617 (17.682)	0.650
0.055 (7.498)	0.609 (3.895)		0.398	3.469 (7.691)		93.837 (9.013)	0.741
0.060 (6.316)	0.635 (3.952)	0.166 (0.722)	0.401	3.440 (7.536)	3.414 (0.456)	92.754 (8.666)	0.742

Table 4
OLS regressions of 5Y CDX index swap rate MarketRate_t on the corresponding actuarial rate ActRate_t. The t-statistics are shown parenthetically.

	Sample	Intercept	Factor loading	R^2
CDX HY Index	Daily	220.832 (5.176)	0.798 (2.604)	0.011
	Weekly	261.724 (3.290)	0.530 (0.930)	0.005
CDX IG Index	Daily	25.819 (15.095)	45.252 (12.252)	0.166
	Weekly	26.529 (7.610)	43.433 (5.858)	0.169

spread for the IG π_t is consistent with the view that the fears regarding the health of the banking system that started to appear in 2007 played an important role for IG risk premia.

In order to better understand the source of the variation of CDX swap market rates, we regress the daily and weekly 5Y CDX index swap MarketRate_t on the corresponding ActRate_t. The results of these regressions, reported in Table 4, suggest that the ActRate_t has no explanatory power for HY swap market rates and very limited explanatory power for IG swap market rates, with an R^2 of roughly 17%. We have also considered the daily and weekly changes of MarketRate_t, but have found the regression coefficients to be insignificant. These results indicate that the (vast) majority of the movements in CDX swap market rates are unrelated to movements in correlated CDX default risk as measured by the actuarial swap rate.

5. Conclusion

We have documented substantial risk premia associated with correlated default events in the CDX High Yield and Investment Grade credit swap portfolios as well as with the unpredictable future variation in the risk-neutral arrival rate of defaults in these portfolios. The effects of these risk premia on CDX index and tranche swap rates covary strongly across swap maturities, and depend on general stock market volatility and various measures of credit spreads. The large moves in these premia during April 2004–October 2007 have natural interpretations in terms of economic and financial market developments. Our results suggest that a large portion of the movements in CDX swap market rates may be caused by changing attitudes toward correlated default risk rather than changes in the economic factors affecting the actual risk of clustered defaults, which ultimately governs swap payoffs.

Acknowledgments

This research was supported by grants from JP Morgan Chase's Academic Outreach Program and Moody's Credit Market Research Fund, for which we are very grateful. We are also grateful for data from Barclays Capital and Moody's Investors Service. We thank Sanjiv Das, Darrell Duffie, Andreas Eckner, Eymen Errais, Francis Longstaff, Ilya Strebulaev, Stefan Weber, Liuren Wu, several anonymous referees, and seminar participants at the University of Chicago's Stevanovich Center for Financial Mathematics Conference on Credit Risk, the Standard and Poor's Credit Risk Summit and the Daiwa International Workshop on Financial Engineering for discussions and comments, and Wilfred Wong for excellent research assistance.

Author's personal copy

References

Altman, E., Brady, B., Resti, A., Sironi, A., 2005. The link between default and recovery rates: theory empirical evidence and implications. Journal of Business 78 (6), 2203–2227.

Artzner, P., Delbaen, F., 1992. Credit risk and prepayment option. ASTIN Bulletin 22, 81–96.

Azizpour, S., Giesecke, K., Schwenkler, G., 2008. Exploring the sources of default clustering. Working Paper, Stanford University.

Berndt, A., Douglas, R., Duffie, D., Ferguson, M., Schranz, D., 2008. Measuring default risk premia from default swap rates and EDFs. Working Paper, Stanford University.

Bhansali, V., Gingrich, R., Longstaff, F., 2008. Systemic credit risk: what is the market telling us. Financial Analysts Journal 64, 16–24.

Blanchet-Scalliet, C., El Karoui, N., Martellini, L., 2005. Dynamic asset pricing theory with uncertain time-horizon. Journal of Economic Dynamics and Control 29 (10), 1737–1764.

Box, G., Tiao, G., 1992. Bayesian Inference in Statistical Analysis. Wiley, New York.

Brémaud, P., 1980. Point Processes and Queues—Martingale Dynamics. Springer-Verlag, New York.

Campi, L., Polbennikov, S., Sbuelz, A., 2009. Systematic equity-based credit risk: a CEV model with jump to default. Journal of Economic Dynamics and Control 33 (1), 93–108.

Carr, P., Wu, L., 2007. Theory and evidence on the dynamic interactions between sovereign credit default swaps and currency options. Journal of Banking and Finance 31, 2383–2403.

Cheridito, P., Filipovic, D., Kimmel, R., 2007. Market price of risk specifications for affine models: theory and evidence. Journal of Financial Economics 83 (1), 123–170.

Collin-Dufresne, P., Goldstein, R., Helwege, J., 2009. How large can jump-to-default risk premia be? Modeling contagion via the updating of beliefs. Working Paper, Columbia University.

Coval, J., Jurek, J., Stafford, E., 2009. Economic catastrophe bonds. American Economic Review 99 (3), 628–666.

Das, S., Duffie, D., Kapadia, N., Saita, L., 2007. Common failings: how corporate defaults are correlated. Journal of Finance 62, 93–117.

Driessen, J., 2005. Is default event risk priced in corporate bonds. Review of Financial Studies 18, 165–195.

Duffie, D., Garleanu, N., 2001. Risk and valuation of collateralized debt obligations. Financial Analysts Journal 57 (1), 41–59.

Eckner, A., 2007. Risk premia in structured credit derivatives. Working Paper, Stanford University.

Errais, E., Giesecke, K., Goldberg, L., 2010. Affine point processes and portfolio credit risk. SIAM Journal on Financial Mathematics 1, 642–665.

Giesecke, K., Kim, B., 2007. Estimating tranche spreads by loss process simulation. In: Henderson, S.G., Biller, B., Hsieh, M.-H., Shortle, J., Tew, J.D., Barton, R.R. (Eds.), Proceedings of the 2007 Winter Simulation Conference. IEEE Press, pp. 967–975.

Giesecke, K., Kim, B., 2011. Risk analysis of collateralized debt obligations. Operations Research 59 (1), 32–49.

Giesecke, K., Goldberg, L., 2003. The market price of credit risk: the impact of asymmetric information. Working Paper, Cornell University.

Giesecke, K., Goldberg, L., Ding, X. A top-down approach to multi-name credit. Operations Research, in press, doi:10.1287/opre.1100.0855.

Hamilton, D., 2005. Moodys senior ratings algorithm and estimated senior ratings. Moodys Investors Service.

Jarrow, R.A., Lando, D., Yu, F., 2005. Default risk and diversification: theory and applications. Mathematical Finance 15, 1–26.

Longstaff, F., Rajan, A., 2008. An empirical analysis of collateralized debt obligations. Journal of Finance 63 (2), 529–563.

Longstaff, F., Mithal, S., Neis, E., 2005. Corporate yield spreads: default risk or liquidity? New evidence from the credit default swap market. Journal of Finance 60, 2213–2253.

Meyer, P.-A., 1971. Démonstration simplifée d'un théorème de Knight. In: Séminaire de Probabilités V, Lecture Notes in Mathematics 191. Springer-Verlag, Berlin, pp. 191–195.

Ogata, Y., 1978. The asymptotic behavior of maximum likelihood estimators of stationary point processes. Annals of the Institute of Statistical Mathematics 30 (A), 243–261.

Pan, J., Singleton, K., 2008. Default and recovery implicit in the term structure of sovereign CDS spreads. Journal of Finance 63, 2345–2384.

Protter, P., 2004. Stochastic Integration and Differential Equations. Springer-Verlag, New York.