

An Empirical Analysis of the Pricing of Collateralized Debt Obligations

FRANCIS A. LONGSTAFF and ARVIND RAJAN*

ABSTRACT

We use the information in collateralized debt obligations (CDO) prices to study market expectations about how corporate defaults cluster. A three-factor portfolio credit model explains virtually all of the time-series and cross-sectional variation in an extensive data set of CDX index tranche prices. Tranches are priced as if losses of 0.4%, 6%, and 35% of the portfolio occur with expected frequencies of 1.2, 41.5, and 763 years, respectively. On average, 65% of the CDX spread is due to firm-specific default risk, 27% to clustered industry or sector default risk, and 8% to catastrophic or systemic default risk.

A COLLATERALIZED DEBT OBLIGATION (CDO) is a financial claim to the cash flows generated by a portfolio of debt securities or, equivalently, a basket of credit default swaps (CDS contracts). Thus, CDOs are the credit market counterparts to the familiar collateralized mortgage obligations (CMOs) actively traded in secondary mortgage markets. Since its inception in the mid-1990s, the market for CDOs has become one of the most rapidly growing financial markets ever. Industry sources estimate the size of the CDO market at the end of 2006 to be nearly \$2 trillion, representing more than a 30% increase over the prior year.¹ Recently, CDOs have been in the spotlight because of the May 2005 credit crisis in which downgrades of Ford's and General Motors' debt triggered a wave of large CDO losses among many credit-oriented hedge funds and Wall

*Francis Longstaff, Allstate Professor of Insurance and Finance, UCLA Anderson School and the NBER. Arvind Rajan, Managing Director, Relative Value Trading, Citigroup. We are grateful for the valuable comments and suggestions received from Vineer Bhansali, Pierre Collin-Dufresne, Keith Crider, Sanjiv Das, Darrell Duffie, Youseff Elouerkhaoui, Bjorn Flesaker, Kay Giesecke, Ben Golub, Mitch Janowski, Holger Kraft, David Lando, L. Sankarasubrahmanian, Alan Shaffran, David Shelton, Ken Singleton, Jure Skarabot, Ryoichi Yamabe, Jing Zhang, and seminar participants at Columbia University, the Federal Reserve Bank of New York, Harvard Business School/Harvard Department of Economics, IXIS Capital Markets, the Journal of Investment Management Conference, the NBER Summer Institute, New York University, Pimco, Stanford, the University of British Columbia, the University of California at Berkeley, the University of California at Los Angeles, the University of Michigan, and Yale University. We are particularly grateful for the comments and suggestions of the editor Campbell Harvey and an anonymous referee and associate editor. We also thank Guarav Bansal, Yuzhao Zhang, and Xiaolong Cheng for capable research assistance. All errors are our responsibility.

¹ A key driver of the growth in the CDO market is the parallel growth in the credit derivatives market, which the International Swaps and Derivatives Association (ISDA) estimates reached \$26 trillion notional in mid-2006.

Street dealers. Despite the importance of this market, however, relatively little research on CDOs has appeared in the academic literature to date.

CDOs are important not only to Wall Street, but also to researchers since they provide a near-ideal “laboratory” for studying a number of fundamental issues in financial economics. For example, CDOs allow us to identify the *joint* distribution of default risk across firms since CDOs are claims against a portfolio of debt, information that cannot be inferred from the *marginal* distributions associated with single-name credit instruments. The joint distribution is crucial to understanding how much credit risk is diversifiable and how much contributes to the systemic risk of “credit crunches” and liquidity crises in financial markets. Furthermore, clustered default risk has implications for the corresponding stocks since default events may map into nondiversifiable event risk in equity markets.

CDO-like structures are emerging as a major new type of financial vehicle and/or “virtual” institution.² In particular, the CDO structure can be viewed as an efficient special purpose vehicle for making illiquid assets tradable, creating new risk-sharing and insurance opportunities in financial markets, and completing markets across credit states of the world. CDO-like structures are now used not only for corporate bonds and loans, but also for less liquid and more private assets such as subprime home equity loans, credit card receivables, commercial mortgages, auto loans, student loans, equipment leases, trade receivables, small business loans, private equity, emerging market local assets, and even the “intellectual” property rights of rock stars.³ Finally, observe that a CDO could also be viewed as a “synthetic bank” in the sense that its assets consist of loans and its liabilities run the gamut from near-riskless senior debt to highly leveraged equity. The key distinction, however, is that the “synthetic” CDO bank may not engage in the same type of monitoring activities as actual banks. Thus, a comparison of CDO equity and bank stocks could provide insights into the delegated-monitoring role of financial intermediaries.⁴

This paper represents a first attempt to understand the economic structure of default risk across firms using information from the CDO market. Specifically, we use the prices of standardized tranches on the CDX credit index to infer the market’s expectations about the way in which default events cluster across firms. Motivated by recent research by Collin-Dufresne, Goldstein, and Martin (2001), Elton et al. (2001), Eom, Helwege, and Huang (2004), Longstaff, Neis, and Mithal (2005), and others who show that corporate credit spreads are driven by firm-specific factors as well as broader economic forces, we develop a simple multifactor portfolio credit model for pricing CDOs. Our framework has some features in common with Duffie and Gârleanu (2001), who allow for three types of default events in their framework: idiosyncratic or firm-specific defaults, industrywide defaults in a specific sector of the economy, and

² For a discussion on the role of tranching in markets with asymmetric information, see Demarzo (2005).

³ For example, see Richardson (2005) for a discussion of the “Bowie” bonds.

⁴ For example, see Diamond (1984).

economywide defaults affecting every industry and sector. However, rather than focusing on the individual “quantum” or “zero-one” states of default for each firm and aggregating up to the portfolio level, our framework takes a “statistical mechanics” approach by modeling portfolio credit losses directly. Specifically, we allow portfolio losses to occur as the realizations of three separate Poisson processes, each with a different jump size and intensity process.⁵ We take the model to the data by fitting it to the CDX index spread and the prices of the 0–3%, 3–7%, 7–10%, 10–15%, and 15–30% CDX index tranches for each date during the sample period.

We first address the issue of how many factors are needed to explain CDO prices. To do this, we estimate one-factor and two-factor versions of the model and use a likelihood-ratio approach to test whether a $N + 1$ -factor model has significant explanatory power relative to a N -factor model. The three-factor model significantly outperforms the two-factor model, which, in turn, significantly outperforms the one-factor model. These results provide the first direct evidence that the market expects defaults for the firms in the CDX index to cluster (correlated defaults).

Focusing on the three-factor results, the estimated jump sizes for the three Poisson processes are about 0.4%, 6%, and 35%, respectively. Since there are 125 firms in the CDX index, the jump size of 0.4% for the first process can be interpreted as the portfolio loss resulting from the default of a single firm, given a 50% recovery rate ($1/125 \times 0.50 = 0.004$). The jump size of 6% for the second process can be viewed as an event in which, say, 15 firms default together. Since this represents roughly 10% of the firms in the portfolio, one possible interpretation of this event could be that of a major crisis that pushes an entire industry or sector into financial distress. However, there are many other possible interpretations. For example, this type of event could just as easily involve clustered defaults among firms with similar accounting ratios, currency or raw materials exposures, firm age, firm size, etc.⁶ Finally, the 35% jump size for the third process could be viewed as a catastrophic or systemic event that wipes out the majority of firms in the economy. Our analysis indicates that all three types of credit risk are anticipated by the market.

We also estimate the probabilities or intensities of the three Poisson events (under the risk-neutral pricing measure). On average, the expected time until an idiosyncratic or firm-specific default is 1.2 years, the expected time until a clustered industry default crisis is 41.5 years, and the expected time until a catastrophic economywide default event is 763 years.⁷

⁵ In independent work, Giesecke and Goldberg (2005) put forward an interesting approach to modeling multiname credit risk that also has many similarities to ours. Their approach is called the top–down approach.

⁶ I am grateful to the associate editor for this insight.

⁷ An expected time of 763 years may seem unrealistically long, but it is important to observe that there has never been a credit event in the U.S. history—not even during the U.S. Civil War or the Great Depression—in which more than 50% of the firms in the economy defaulted or went bankrupt. On the other hand, there are numerous documented economic collapses and sovereign defaults in erstwhile safe countries over the past centuries, suggesting that a nonzero probability is appropriate to attach to such an event (see Kindleberger (2005)).

In an effort to understand whether clustering in default risk is in fact linked to industry, we perform a principal components analysis of changes in the CDS spreads for the individual firms in the CDX index. We find that there is a dominant first factor driving spreads across all industries. This is consistent with there being a pervasive economywide component to credit. Moving beyond this first factor, however, we find that the second, third, fourth, etc. principal components are significantly related to specific industries or groups of industries. Thus, there is some evidence that default clustering occurs in ways that have some relation to industry. On the other hand, when we repeat the principal components analysis using stock returns for the individual firms in the CDX index, we find that the second, third, fourth, etc. principal components for stock returns are much more strongly related to industry than is the case for the CDS spreads. Thus, there are intriguing differences in the cross-sectional structure of stock returns and credit spreads for the firms in the CDX index.

Using the intensity estimates, we decompose the level of the CDX index spread into its three components. We find that on average, firm-specific default risk represents only 64.6% of the total CDX index spread, while clustered industry or sector and economywide default risks represent 27.1% and 8.3% of the index spread, respectively. Thus, the risk of industry or economywide financial distress accounts for more than one-third of the default risk in the CDX portfolio. Recently, however, idiosyncratic default risk has played a larger role.

Next we examine how well the model captures the pricing of individual index tranches. Even though tranche spreads are often measured in hundreds or even thousands of basis points, the root-mean-squared error (RMSE) of the three-factor model is typically on the order of only two to three basis points, which is well within the typical bid-ask spreads in the market. Thus, virtually all of the time-series and cross-sectional variation in index tranche prices is captured by the model. We find that the largest pricing errors occur shortly after the inception of the CDX index and tranche market, but decrease rapidly after several weeks. Thus, despite some early mispricing, the evidence suggests that the CDX index tranche market quickly evolved.

There is a rapidly growing literature on credit derivatives and correlated defaults.⁸ This paper contributes to this primarily theoretical literature by presenting a new approach to modeling portfolio default losses, conducting the first extensive empirical analysis of pricing in the CDO markets, and providing the first direct estimates of the nature and degree of default clustering across firms expected by market participants.

The remainder of this paper is organized as follows. Section I provides an introduction to the CDO market. Section II describes the data used in the study.

⁸ Important recent contributions in this area include Duffie and Gârleanu (2001), Hull and White (2003), Giesecke (2004), Jorion and Zhang (2005), Longstaff et al. (2005), Saita (2005), Yu (2005, 2007), Giesecke and Goldberg (2005), Frey and Backhaus (2005), Schönbucher (2005), Das et al. (2006), Bakshi, Madan, and Zhang (2006), Duffie et al. (2006), Errais, Giesecke, and Goldberg (2006), Duffie, Saita, and Wang (2007), Das et al. (2007), and many others.

Section III presents the three-factor portfolio credit model. Section IV applies the model to the valuation of index tranches. Section V reports the results from the empirical analysis. Section VI summarizes the results and makes concluding remarks.

I. An Introduction to CDOs

CDOs have become one of the most important new financial innovations of the past decade. It is easiest to think of a CDO as a portfolio containing certain debt securities as assets, and multiple claims in the form of issued notes of varying seniority. The liabilities are serviced using the cash flows from the assets, as in a corporation. Although CDOs existed in various forms previously, it was only in the mid-1990s that they began to be popular. Over subsequent years, issuance experienced rapid growth. For example, during the first three quarters of 2006, issuance was \$322 billion, representing nearly a 102% increase over the same period during 2005.⁹ The assets securitized by cash CDOs have broadened to include investment-grade bonds, high yield bonds, emerging market securities, leveraged loans, middle market loans, trust preferred securities, asset-backed securities, commercial mortgages, and even previously issued CDO tranches.¹⁰

Over the past few years, the technology of cash CDOs has merged with the technology of credit derivatives to create the so-called synthetic CDO, which is the main focus of this paper. Synthetic CDOs differ from cash CDOs in that the portfolios that provide the cash flow to service their liabilities consist of credit default swaps rather than bonds or other cash securities. The majority of synthetic securities are based on corporate credit derivatives, and tend to be simpler to model.

A. An Example of a Stylized CDO

To build up understanding of a full-fledged synthetic CDO, we consider a simple example based on a \$100 million investment in a diversified portfolio of 5-year par corporate bonds. Imagine that a financial institution (CDO issuer) sets up this portfolio, which consists of 100 separate bonds, each with a market value of \$1 million, and each issued by a different firm. Imagine also that each bond is rated BBB and has a coupon spread over Treasuries of 100 basis points. The CDO issuer can now sell 5-year claims against the cash flows generated by the portfolio. These claims are termed CDO tranches and are constructed to vary in credit risk from very low (senior tranches) to low (junior or mezzanine tranches) to very high (the “equity” tranche).

⁹ To put these numbers in perspective, we note that according to the Securities Industry and Financial Markets Association, the total issuance of corporate bonds and agency mortgage-backed securities during 2005 was \$703.2 and \$966.1 billion, respectively.

¹⁰ For additional insights into the CDO market, see the excellent discussions provided by Duffie and Gârleanu (2001), Duffie and Singleton (2003), Roy and Shelton (2007), and Rajan, McDermott, and Roy (2007).

Let us illustrate a typical CDO structure by continuing the example. First, imagine that the CDO issuer creates a so-called equity tranche with a total notional amount of 3% of the total value of the portfolio (\$3 million). By definition, this tranche absorbs the first 3% of any defaults on the entire portfolio. In exchange, this tranche may receive a coupon rate of, say, 2,500 basis points above Treasuries. If there are no defaults, the holder of the equity tranche earns a high coupon rate for 5 years and then receives back his \$3 million notional investment. Now assume that one of the 100 firms represented in the portfolio defaults (and that there is zero recovery in the event of default). In this case, the equity tranche absorbs the \$1 million loss to the portfolio and the notional amount of the equity tranche is reduced to \$2 million. Thus, the equity tranche holder has lost one-third of his investment. Going forward, the equity tranche investor receives the 2,500 basis point coupon spread as before, but now only on his \$2 million notional position. Now assume that another two firms default. In this case, the equity tranche absorbs the additional losses of \$2 million, the notional amount of the equity tranche investor's position is completely wiped out, and the investor receives neither coupons nor principal going forward. Because a 3% loss in the portfolio translates into a 100% loss for the equity tranche investor, we can view the equity tranche investor as being leveraged 33 1/3 to 1. However, unlike an investor who leverages by borrowing, the equity tranche investor has no liability beyond a 3% portfolio loss, a condition referred to as "nonrecourse" leverage.

Now imagine that the CDO issuer also creates a junior mezzanine tranche with a total notional amount of 4% of the total value of the portfolio (\$4 million). This tranche absorbs up to 4% of the total losses on the entire portfolio after the equity tranche has absorbed the first 3% of losses. For this reason, this tranche is designated the 3–7% tranche. In exchange for absorbing these losses, this tranche may receive a coupon rate of, say, 300 basis points above Treasuries. If total credit losses are less than 3% during the 5-year horizon of the portfolio, then the 3–7% investor earns the coupon rate for 5 years and then receives back his \$4 million notional investment. If total credit losses are greater than or equal to 7% of the portfolio, the total notional amount for the 3–7% investor is wiped out.

The CDO issuer follows a similar process in creating additional mezzanine, senior mezzanine, and even super-senior mezzanine tranches. A typical set of index CDO tranches might include the 0–3% equity tranche, and 3–7%, 7–10%, 10–15%, 15–30%, and 30–100% tranches. The initial levels 3%, 7%, 10%, 15%, and 30% at which losses begin to accrue for the respective tranches are called attachment points or subordination levels. Note that the total notional valuation of all the tranches equals the \$100 million notional of the original portfolio of corporate bonds.

An interesting aspect of the CDO creation process is that since each tranche has a different degree of credit exposure, each tranche could have its own credit rating. For example, the super-senior 30–100% tranche can only suffer credit losses if total losses on the underlying portfolio exceed 30% of the total notional amount. Since this is highly unlikely, this super-senior tranche would typically

have a AAA rating, even if all the underlying bonds were below investment grade. This example illustrates that the tranching process allows securities of any credit rating to be created. Thus, the CDO process can serve to complete the financial market by creating high credit quality securities that might not otherwise exist in the market.

B. Synthetic CDOs

To take advantage of the wide availability of credit derivatives, credit markets have recently introduced CDO structures known as synthetic CDOs. This type of structure has become very popular and the total notional amount of synthetic CDO tranches is growing rapidly. A synthetic CDO is economically similar to a cash CDO in most respects. The principal difference is that rather than there being an underlying portfolio of corporate bonds on which tranches are based, the underlying portfolio is actually a basket of credit default swap contracts. Recall that a CDS contract functions as an insurance contract in which a buyer of credit protection makes a fixed payment each quarter for some given horizon such as 5 years.¹¹ If there is a default on the underlying reference bond during that period, however, then the buyer of protection is able to give the defaulted bond to the protection seller and receive par (the full face value of the bond).¹² Thus, the first step in creating a synthetic CDO is to define the underlying basket of CDS contracts.

C. Credit Default Indexes and Index Tranches

In this study, we focus on CDOs with cash flows tied to the most liquid U.S. corporate credit derivative index, the DJ CDX North American Investment Grade Index. This index is managed by Dow Jones and is based on a liquid basket of CDS contracts for 125 U.S. firms with investment grade corporate debt. The CDX index itself trades just like a single-name CDS contract, with a defined premium based on the equally weighted basket of its 125 constituents. The individual firms included in the CDX basket are updated and revised ("rolled") every 6 months in March and September, with a few downgraded and illiquid names being dropped and new ones taking their place. CDX indexes are numbered sequentially. Thus, the index for the first basket of 125 firms was designated the CDX NA IG 1 index in 2003, the index for the second basket of 125 firms the CDX NA IG 2 index, etc., and so on up to CDX NA IG 7 in September 2006, of which the first five series comprise the data set analyzed in this paper. While there is considerable overlap between successive CDX NA IG indexes, occasionally there are significant changes across index rolls.

¹¹ As with any swap contract, however, CDS contracts carry the small additional risk of a counterparty default. In reality, this risk can be largely mitigated by the posting of collateral between swap counterparties.

¹² This aspect of the contract design means that the protection buyer can be compensated for his losses relatively quickly; the protection buyer does not need to wait until the end of the bankruptcy and recovery process.

For example, the CDX NA IG 4 index (beginning in March 2005) includes Ford and General Motors while the CDX NA IG 5 index (beginning in September 2005) does not since the debt for these firms dropped below investment grade in May 2005.

Index CDO tranches have also been issued, each tied to a specific CDX index. The attachment points of these CDO tranches are standardized at 3%, 7%, 10%, 15%, and 30%, exactly as in the example above. Since these instruments are structured as credit default swaps, when investors “buy” a synthetic index tranche from a counterparty, they are selling protection on that tranche. Their counterparty has bought protection on the same tranche from them. This highlights a convenient feature of these index tranches—that is, a dealer need not create and sell the entire capital structure of tranches to investors; rather, investors are free to synthetically create and trade (sell or buy) individual index tranches (single-tranche index CDOs) according to their needs. As observed earlier, the losses on an N – M % tranche are zero if the total losses on the underlying portfolio are less than N . On the other hand, the total losses on the tranche are 1.00 or 100% if the total losses on the underlying portfolio equals or exceeds M . For underlying portfolio losses between N and M , tranche losses are linearly interpolated between zero and one. Thus, the losses on a N – M % tranche can be viewed intuitively as a call spread on the total losses of the underlying portfolio. This intuition will be formalized in a later section. Just as an option has a “delta,” that is, an equivalent exposure to the underlying, the tranche has a delta with respect to its underlying index.

Although index tranches are the most liquid synthetic tranches, a synthetic tranche can be based on any portfolio. A tranche created with a specific nonindex portfolio, and with customized attachment points, e.g., 5–8%, is called a bespoke tranche. While the results in this paper are based on index tranche data, the analysis can also be applied to most bespoke CDO tranches. Finally, there are also full capital structure synthetic CDOs, created when demand exists for the entire capital structure. Provided a CDO observes the simple type of structure we specified in the example, a model such as the one in this paper may be used to price its tranches.¹³

II. The Data

CDOs are a relatively new financial innovation and have only recently begun to trade actively in the markets. As a result, it has been difficult for researchers to obtain reliable CDO pricing data. We were fortunate, however, to be given access by Citigroup to one of the most extensive proprietary data sets of CDO index and tranche pricing data in existence.¹⁴

¹³ The analysis in this paper, however, may not apply directly to certain other types of portfolio derivative products, for example, N th-to-default baskets, CDO-squareds, and cash CDOs, which have more granular compositions, more complex structures, or more difficult-to-model cash flows and rules, respectively.

¹⁴ Although the data set we were given access to is proprietary, data for standardized CDX index tranches are now available on the Bloomberg system and other commercial sources.

The data consist of daily closing values for the 5-year CDX NA IG index (CDX index for short) for the period from October 2003 to October 2005. As discussed earlier, the underlying basket of 125 firms in the index is revised every March and September. Thus, the index data correspond to the five individual indexes denoted CDX i , $i = 1, 2, 3, 4$, and 5. CDX 1 covers October 20, 2003 to March 19, 2004; CDX 2 covers March 22, 2004 to September 22, 2004; CDX 3 covers September 23, 2004 to March 18, 2005; CDX 4 covers March 21, 2005 to September 19, 2005; and CDX 5 covers September 20, 2005 to October 18, 2005. This data set covers virtually the entire history of the CDX index through 2005. Data are missing for some days during the earlier part of the sample. We omit these days from the sample, leaving us with a total of 435 usable daily observations for the 2-year sample period. For the primarily descriptive purposes of this section, we report summary statistics based on the continuous series of the on-the-run CDX index (rather than reporting statistics separately for the individual CDX series).

In addition to the index data, we also have daily closing quotation data for the 0–3%, 3–7%, 7–10%, 10–15%, and 15–30% tranches on the CDX index. The pricing data for most tranches are in terms of the basis point premium paid to the CDO investor for absorbing the losses associated with the individual tranches. Thus, a price of 300 for the 3–7% tranche implies that the tranche investor would receive a premium of 300 basis points per year paid quarterly on the remaining balance in exchange for absorbing the default losses from 3% to 7% on the CDX index. The exception is the market convention for the equity tranche (the 0–3% tranche), which is generally quoted in terms of points up front. A price of 50 for this tranche means that an investor would need to receive \$50 up front per \$100 notional amount, plus a premium of 500 basis points per year paid quarterly on the remaining balance, to absorb the first 3% of losses on the CDX index. Rather than using this market convention, however, we convert the points up front into spread equivalents to facilitate comparison with the pricing data for the other tranches.

In addition to the CDX index and tranche data, we also collect daily New York closing data on 3-month, 6-month, 12-month Libor rates, and on 2-year, 3-year, 5-year, 7-year, and 10-year swap rates. The Libor data are obtained from the Bloomberg system. The swap data are obtained from the Federal Reserve Board's web site. From this Libor spot rate and swap par rate data, we use a standard cubic spline approach to bootstrap zero-coupon curves that will be used throughout the paper to discount cash flows.¹⁵ Since the same zero-coupon curve is used to discount both legs of the CDO contract, however, the results are largely insensitive to the decision to discount using the Libor-swap curve; the results are virtually identical when the bootstrapped Treasury curve is used for discounting cash flows.

Table I provides summary statistics for index and tranche data. As shown, the average values of the spreads are monotone decreasing in seniority (attachment point). The average spread for the 0–3% equity tranche is 1,758.87

¹⁵ See Longstaff et al. (2005) for a more detailed discussion of this bootstrapping algorithm.

Table I
Summary Statistics for the Levels and First Differences of the CDX North American Investment Grade Index and Index Tranche Spreads

This table reports summary statistics for the market spreads and the daily change in the spreads (spreads measured in basis points) for the indicated time series. Correlations shown in the top panel are correlations of levels; correlations shown in the bottom panel are correlations of first differences. Results are reported for the combined on-the-run time series. The sample period is from October 2003 to October 2005.

	Correlations					Mean	SD	Min.	Med.	Max.	Skew.	Kurt.	Serial corr.	N
	0-3	3-7	7-10	10-15	15-30									
CDX Index	0.809	0.668	0.607	0.672	0.880	54.52	6.95	41.00	54.50	78.00	0.25	-0.44	0.977	435
0-3 Tranche		0.183	0.068	0.217	0.711	1758.87	245.30	1282.52	1780.28	2914.55	1.05	3.28	0.978	435
3-7 Tranche			0.960	0.968	0.610	240.07	85.66	103.50	232.00	411.00	0.13	-1.33	0.992	435
7-10 Tranche				0.933	0.540	82.27	36.73	24.75	74.75	158.00	0.22	-1.23	0.992	435
10-15 Tranche					0.674	34.43	14.49	11.88	29.50	66.50	0.39	-1.16	0.989	435
15-30 Tranche						11.54	2.89	5.00	11.88	20.75	0.03	-0.42	0.962	435
Δ CDX Index	0.896	0.872	0.761	0.750	0.730	-0.02	1.51	-7.00	0.00	7.00	0.02	3.86	0.013	434
Δ 0-3 Tranche		0.762	0.648	0.662	0.651	0.48	51.63	-221.86	0.27	285.15	0.52	6.38	0.081	434
Δ 3-7 Tranche			0.879	0.840	0.773	-0.55	10.88	-54.00	0.00	40.50	-0.44	4.24	-0.028	434
Δ 7-10 Tranche				0.861	0.751	-0.20	4.68	-24.50	0.00	19.50	0.07	3.30	-0.079	434
Δ 10-15 Tranche					0.764	-0.09	2.14	-8.50	0.00	9.50	0.17	3.00	-0.067	434
Δ 15-30 Tranche						-0.02	0.79	-3.00	0.00	4.50	0.39	4.10	-0.043	434

basis points (which translates into an average number of points up front of 39.34). This spread is many times larger than the average spread for the junior mezzanine 3–7% tranche, indicating that the expected losses for the equity tranche are much higher than those for more senior tranches. Similar comparisons hold for all the other tranches. Figure 1 plots the time series of tranche spreads for the various attachment points. The correlations indicate that while these spreads have a high level of correlation with each other, there is also considerable independent variation.




Figure 1. CDX index and tranche spreads. This figure graphs the time series of the CDX index and its tranche spreads for the October 2003 to October 2005 sample period. Spreads are in basis points. The vertical division lines denote the roll from one CDX index to the next.




Figure 1. *Continued*

III. The Model

Motivated by these aspects of the data, as well as by the mounting evidence in the literature that credit spreads are driven by idiosyncratic as well as broader market factors, we develop a simple multifactor portfolio credit model for valuing CDO index tranches in this section.¹⁶ Although developed independently, our framework complements important recent theoretical work on “top down” portfolio credit modeling by Giesecke and Goldberg (2005) and others.¹⁷

¹⁶ Evidence about the multifactor nature of credit risk is provided by Collin-Dufresne et al. (2001), Elton et al. (2001), Eom et al. (2004), Longstaff et al. (2005), and many others.

¹⁷ Also see recent papers by Giesecke (2004), Schönbucher (2005), and Sidenius, Piterbarg, and Andersen (2005).

To date, most modeling of CDOs has been done at the firm level by modeling individual losses and then aggregating over the portfolio. However, losses on the tranches are simple functions of the total losses on the underlying portfolio. Thus, the distribution of total portfolio losses represents a “sufficient statistic” for valuing tranches. Accordingly, rather than modeling individual defaults, we model the distribution of total portfolio losses directly.

We stress that we are not implying that individual firm-level information about default status is unimportant. For many types of credit derivatives (such as credit default swaps or first-to-default swaps on small baskets of firms), individual firm default status is essential in defining the cash payoffs. Rather, we suggest that for many other types of credit-related contracts that are tied to larger portfolios, the “reduced-form” approach of modeling portfolio-level losses directly may provide important advantages with little loss in our ability to capture the underlying economics. In general, the smaller the single-name risk concentration in a portfolio, the more applicable is the aggregate loss approach taken here.

Let L_t denote the total portfolio losses on the CDX portfolio per \$1 notional amount. By definition, $L_0 = 0$. To model the dynamic evolution of L_t we assume

$$\frac{dL_t}{1 - L_t} = \bar{\gamma}_1 dN_{1t} + \bar{\gamma}_2 dN_{2t} + \bar{\gamma}_3 dN_{3t}, \quad (1)$$

where $\bar{\gamma}_i = 1 - e^{-\gamma_i}$; $i = 1, 2, 3$; γ_1 , γ_2 , and γ_3 are nonnegative constants defining jump sizes; and N_{1t} , N_{2t} , and N_{3t} are independent Poisson processes. Note that for small values of γ_i , the jump size $\bar{\gamma}_i$ is essentially just γ_i . Thus, for expositional simplicity, we will take a slight liberty and generally refer to the parameters γ_1 , γ_2 , and γ_3 simply as jump sizes. Integrating equation (1) and conditioning on time-zero values (a convention we adopt throughout the paper) gives the general solution for L_t

$$L_t = 1 - e^{-\gamma_1 N_{1t}} e^{-\gamma_2 N_{2t}} e^{-\gamma_3 N_{3t}}. \quad (2)$$

From this equation, it can be seen that the economic condition $0 \leq L_t \leq 1$ is satisfied for all t . Furthermore, since N_{1t} , N_{2t} , and N_{3t} are nondecreasing processes, the intuitive requirement that total losses be a nondecreasing function of time is also satisfied. These dynamics imply that there are three factors at work in generating portfolio losses, each of which could be a firm-specific default event or a multifirm default event. Thus, this approach explicitly allows for the possibility of default correlation.

The intensities of the three Poisson processes are designated λ_{1t} , λ_{2t} , and λ_{3t} , respectively. To complete the specification of the general model, we assume that the dynamics for the intensity processes are given by

$$d\lambda_{1t} = (\alpha_1 - \beta_1 \lambda_{1t}) dt + \sigma_1 \sqrt{\lambda_{1t}} dZ_{1t}, \quad (3)$$

$$d\lambda_{2t} = (\alpha_2 - \beta_2 \lambda_{2t}) dt + \sigma_2 \sqrt{\lambda_{2t}} dZ_{2t}, \quad (4)$$

$$d\lambda_{3t} = (\alpha_3 - \beta_3 \lambda_{3t}) dt + \sigma_3 \sqrt{\lambda_{3t}} dZ_{3t}, \quad (5)$$

where Z_{1t} , Z_{2t} , and Z_{3t} are standard independent Brownian motion processes. These dynamics ensure that the intensities for the three Poisson processes are always nonnegative. Furthermore, the mean-reverting nature of the intensities allows the model to potentially capture expected migrations in the credit quality of the underlying portfolio. Specifically, we would anticipate that over time, the lowest credit quality firms would tend to exit the portfolio sooner, resulting in an expected downward trend in the value of λ . This trend could be reflected in the model in the situation in which the initial value of λ is above its long-run mean value of α/β .¹⁸ Since these intensities are stochastic, it is clear from the previous discussion that this framework allows default correlations to vary over time. Although we present analytical results for the general case implied by equations (3) through (5) in this section, the empirical results to be presented later are based on the special case in which $\alpha_i = \beta_i = 0$ for all i .

To value claims that depend on the realized losses on a portfolio, we first need to determine the distribution of L_t . From equation (2), L_t is a simple function of the values of the three Poisson processes. Thus, it is sufficient to find the distributions for the individual Poisson processes, since expectations of cash flows linked to L_t can be evaluated directly with respect to the distributions of N_{1t} , N_{2t} , and N_{3t} .

Since many of the following results are equally applicable to each of the three Poisson processes, we simplify notation whenever possible by dropping the subscripts 1, 2, and 3 when we present generic results and the interpretation is clear from context. Standard results imply that, conditional on the path of λ_t , the probability of $N_T = i$, $i = 0, 1, 2, \dots$ can be expressed as

$$\frac{\exp\left(-\int_0^T \lambda_t dt\right)\left(\int_0^T \lambda_t dt\right)^i}{i!}. \quad (6)$$

Let $P_i(\lambda, T)$ denote $i!$ times the probability that $N_T = i$, conditional on the current (the time-zero unsubscripted) value of λ . Thus,

$$P_i(\lambda, T) = E \left[\exp\left(-\int_0^T \lambda_t dt\right) \left(\int_0^T \lambda_t dt\right)^i \right]. \quad (7)$$

For $i = 0$, the Appendix shows that this expression is easily solved in closed form from results in Cox, Ingersoll, and Ross (1985). For $i > 0$, the results in Karlin and Taylor (1981, pp. 202–204) can be used to show that $P_i(\lambda, T)$ satisfies the recursive partial differential equation

$$\frac{\sigma^2 \lambda}{2} \frac{\partial^2 P_i}{\partial \lambda^2} + (\alpha - \beta \lambda) \frac{\partial P_i}{\partial \lambda} - \lambda P_i + i \lambda P_{i-1} = \frac{\partial P_i}{\partial T}. \quad (8)$$

¹⁸ We are very grateful to the referee for pointing this out.

The Appendix shows that this partial differential equation for $P_i(\lambda, T)$ has the following (poly-affine) closed-form solution

$$P_i(\lambda, T) = A(T) e^{-B(T)\lambda} \sum_{j=0}^i C_{i,j}(T) \lambda^j, \quad (9)$$

where

$$A(T) = \exp\left(\frac{\alpha(\beta - \xi)T}{\sigma^2}\right) \left(\frac{2\xi}{\beta + \xi - (\beta - \xi)e^{-\xi T}}\right)^{\frac{2\alpha}{\sigma^2}}, \quad (10)$$

$$B(T) = \frac{2\xi(\beta + \xi)}{\sigma^2(\beta + \xi - (\beta - \xi)e^{-\xi T})} - \frac{\beta + \xi}{\sigma^2}, \quad (11)$$

and $\xi = \sqrt{\beta^2 + 2\sigma^2}$. The first $C_{i,j}(T)$ function is $C_{0,0}(T) = 1$. The remaining $C_{i,j}(T)$ functions are given as solutions of the recursive system of first-order ordinary differential equations,

$$\frac{dC_{i,i}}{dt} = i C_{i-1,i-1} - (\sigma^2 B(t) + \beta) i C_{i,i}, \quad (12)$$

$$\frac{dC_{i,j}}{dt} = i C_{i-1,j-1} - (\sigma^2 B(t) + \beta) j C_{i,j} + (j+1)(\alpha + j\sigma^2/2) C_{i,j+1}, \quad (13)$$

$$\frac{dC_{i,0}}{dt} = \alpha C_{i,1}, \quad (14)$$

where $1 \leq j \leq i-1$. These differential equations are easily solved numerically subject to the initial condition that $C_{i,j}(0) = 0$ for all $i > 0$.

With these solutions, the expectation of an arbitrary function $F(L_t)$ of the portfolio losses (satisfying appropriate regularity conditions of course) can be calculated directly by the expression

$$E[F(L_t)] = \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} \sum_{k=0}^{\infty} \frac{P_{1,i}(\lambda_1, t)}{i!} \frac{P_{2,j}(\lambda_2, t)}{j!} \frac{P_{3,k}(\lambda_3, t)}{k!} F(L_t). \quad (15)$$

Although the summations range from zero to infinity, only the first few terms generally need to be evaluated since the remainder are negligible.

IV. Valuing Tranches

Given the solutions for the Poisson probabilities, it is now straightforward to value securities with cash flows tied to the realized credit losses of an underlying portfolio such as the CDX index. Let $D(t)$ denote the present value (as of time zero) of a zero-coupon riskless bond with maturity t . For simplicity, we assume that the riskless rate r is independent of the Poisson and intensity processes.

The total losses on an individual N - $M\%$ tranche can be modeled as a call spread on the underlying state variable L_t . Specifically, the total losses V_t on a

N - M tranche (assuming quarterly cash flows and abstracting from day count considerations) can be expressed as

$$V_t = \frac{1}{M - N} (\max(0, L_t - N) - \max(0, L_t - M)), \quad (16)$$

where N and M are denoted in decimal form. This expression indicates that if the total loss on the underlying portfolio L_t is less than N , then the loss on the tranche V_t is zero. If L_t is midway between N and M , the total loss on the tranche V_t is 0.50 or 50%. If L_t equals or exceeds M , the total loss on the tranche V_t equals 1.00 or 100%. As with the total losses on the underlying portfolio, V_t is a nondecreasing function of time.

An investor in an index tranche receives a fixed annuity of h on the remaining balance $1 - V_t$ of the tranche, in exchange for compensating the protection buyer for the losses dV on the tranche. Thus, the value of the premium leg of a $N - M\%$ tranche is given by

$$\frac{h}{4} \sum_{i=1}^{4T} D(i/4) E[1 - V_{i/4}]. \quad (17)$$

Similarly, the value of the protection leg of the $N - M\%$ tranche is given by

$$\sum_{i=1}^{4T} D(i/4) E[V_{i/4} - V_{(i-1)/4}]. \quad (18)$$

Setting the value of the two legs equal to each other and solving gives the value of the tranche spread h . The expectations in these expressions are easily evaluated by substituting the closed-form solutions for the Poisson probabilities into equation (15).

V. Empirical Analysis

In this section, we estimate the model using the times series of CDX index values and the associated index tranche prices. We then examine how the model performs and explore the economic implications of the results.

A. The Empirical Approach

To make the intuition behind the results more clear, we focus on a simple special case of the model in which each of the intensity processes follows a martingale. Thus, we assume that the α and β parameters in equations (3) through (5) are zero. As we will show, even this simplified specification allows us to fit the data with a very small RMSE (and only marginal improvements would be possible by estimating the general case of the model).¹⁹

¹⁹ These parameter restrictions imply that the intensity process is absorbed at zero if it reaches zero. Thus, a more robust specification might allow for a small positive value for α . In actuality, however, the implied intensity values are generally many standard deviations away from zero. Thus, this technical consideration likely has little effect on the estimation results.

In this specification, six parameters need to be estimated: the three jump size parameters γ_1 , γ_2 , and γ_3 , and the three volatility parameters σ_1 , σ_2 , and σ_3 . In addition, the values of the three intensity processes need to be estimated for each date. Our approach in estimating the model will be to solve for the parameter and intensity values that best fit the model to the data. In doing so, we estimate the model separately for each of the five CDX indexes. The reason for this is that there are slight differences in the composition of the individual indexes, potentially resulting in minor differences in parameter values.

Let us illustrate the estimation approach with the specific example of the CDX 1 index. The CDX 1 index was the on-the-run index from October 20, 2003 to March 20, 2004. There are 65 observations for this index in the data set. Let h_{it} denote the market spread for the i^{th} tranche on date t , where i ranges from 1 to 5 (tranche 1 is the equity tranche, tranche 2 is the 3–7 junior mezzanine tranche, etc.) Let θ denote the vector of σ and γ parameters to be estimated. Let λ_t denote the vector of intensities for date t , and λ the set of all 65 of these vectors. The estimation process consists of solving for the parameter vector θ and the 65 λ_t vectors that minimize the following sum of squared errors,

$$\min_{\theta, \lambda} \sum_{t=1}^{65} \sum_{i=1}^5 [h_{it} - \hat{h}_{it}(\theta, \lambda_t)]^2, \quad (19)$$

where $\hat{h}$ denotes the model-implied value of the tranche spread, subject to the model-implied value of the CDX index equaling the market value of the CDX index for each date t . This algorithm is essentially nonlinear least squares and has been widely used in the finance literature in similar types of applications.²⁰ The optimization methodology we use is a direct search algorithm that does not use the gradient or Hessian of the objective function (Fortran IMSL routine DBCPOL). While this algorithm displays robust convergence properties for a variety of starting values, its direct search nature (which keeps trying parameter values far removed from the current minimizing parameter vector in order to avoid local minima) is admittedly somewhat slow to converge.²¹ As a result, some of the optimizations for longer time series such as CDX4 take more than 12 hours of CPU time to complete. Clearly, more efficient optimization algorithms could reduce the computational time significantly. We use a similar procedure to estimate the model for the other CDX indexes. Finally, it is important to note that parameter values and intensities are estimated for the risk-neutral pricing measure (not the objective or historical measure).

²⁰ See Longstaff, Mithal, and Neis (2005), Liu, Longstaff, and Mandell (2006), and many others.

²¹ As robustness checks for the results, we use a variety of starting values for the optimization. For example, we use starting values ranging from 0.0001 to 0.05 for γ_1 , from 0.0001 to 0.20 for γ_2 , and 0.0001 to 0.75 for γ_3 (and similarly for the σ_1 , σ_2 , and σ_3 parameters). The convergence results are robust to the choice of starting parameters.

Table II
Root Mean Squared Errors (RMSE) from Model Fitting and Tests of
the Number of Factors

This table reports the RMSEs for the individual CDX index tranches resulting from fitting the indicated models, where the RMSE is calculated from the pricing errors for the individual tranche. The table also reports the overall RMSE, which is calculated from the pricing errors for all five of the tranches. All RMSEs are measured in basis points. The p -value is for the test of n versus $n - 1$ factors. N denotes the number of observations for the indicated CDX index. The sample period is from October 2003 to October 2005.

Number of factors	Index	Tranche RMSE					Overall		R^2	N
		0-3	3-7	7-10	10-15	15-30	RMSE	p -Value		
One Factor	CDX 1	13.13	78.55	17.78	40.15	13.44	41.11	—	0.9924	65
	CDX 2	26.90	54.02	47.08	47.50	13.37	40.73	—	0.9931	108
	CDX 3	39.21	65.48	41.22	24.75	9.29	40.55	—	0.9887	118
	CDX 4	31.81	39.76	47.65	23.28	12.38	33.35	—	0.9966	127
	CDX 5	49.58	29.92	26.37	13.83	6.21	29.25	—	0.9965	17
Two Factors	CDX 1	9.89	20.24	8.46	6.73	9.82	12.01	0.000	0.9993	65
	CDX 2	5.60	17.17	5.12	5.89	3.02	8.90	0.000	0.9997	108
	CDX 3	3.78	5.77	3.28	1.68	3.07	4.74	0.000	0.9998	118
	CDX 4	22.40	14.30	9.95	7.39	10.85	13.98	0.000	0.9994	127
	CDX 5	2.24	1.95	7.13	6.99	0.66	4.67	0.000	0.9999	17
Three Factors	CDX 1	11.01	4.62	8.68	10.37	8.31	8.88	0.997	0.9996	65
	CDX 2	2.52	1.31	3.04	6.08	4.68	3.90	0.000	0.9999	108
	CDX 3	0.99	0.76	3.36	2.80	2.06	2.23	0.000	1.0000	118
	CDX 4	1.18	1.19	4.52	2.27	2.46	2.63	0.000	1.0000	127
	CDX 5	0.29	0.28	1.10	0.61	0.46	0.63	0.000	1.0000	17

B. Testing for the Number of Factors

One of the key issues to address at the outset is the question of how many factors are actually needed in pricing CDOs. In this section, we explore this issue by testing whether a two-factor version of the model has incremental explanatory power relative to the one-factor version, and similarly for the three-factor version. These tests for the number of factors needed to price tranches also provide insight into an issue that is of fundamental importance in credit markets, namely, default correlation. This follows since if defaults were uncorrelated, then default losses on a portfolio could be modeled using a single-factor Poisson with intensity equal to the sum of intensities for the individual firms in the portfolio (since the sum of independent Poissons is itself a Poisson). Thus, rejecting a single-factor Poisson version of the model would provide direct evidence that the market expects correlation or clustering in the defaults of CDX firms.

Table II presents summary statistics for the pricing errors obtained by estimating one-factor, two-factor, and three-factor versions of the model. In each case, the values of the intensity processes are chosen to match the CDX index spread exactly. In the two-factor and three-factor models, the

RMSE of the difference between market- and model-implied spreads for the five index tranches is also minimized. The table reports the RMSEs for each of the 0–3%, 3–7%, 7–10%, 10–15%, and 15–30% tranches individually, as well as the RMSE computed over all tranches. Table II also reports the p -values for the chi-square tests of the two-factor versus one-factor and three-factor versus two-factor specifications. In the one-factor specification, we estimate the two parameters γ_1 and σ_1 , as well as N values of λ_1 , where N is the number of days in the sample. In the two-factor specification, we estimate the four parameters γ_1 , γ_2 , σ_1 , and σ_2 as well as N values each for λ_1 and λ_2 . Thus, the one-factor specification is nested within the two-factor specification by imposing $N + 2$ restrictions; the chi-square statistic has $N + 2$ degrees of freedom. Similarly, the two-factor specification is nested within the three-factor specification by imposing $N + 2$ restrictions.

As shown in Table II, the RMSEs for the one-factor version of the model are very large across all of the tranches. The overall RMSEs range from about 30 to 41 basis points. Increasing the number of factors to two results in a significant reduction in the RMSEs, both overall and across tranches. Typically, the overall RMSE for the two-factor version of the model is between about 5 and 14 basis points. For each CDX index, the incremental explanatory power of the two-factor version relative to the one-factor version is highly statistically significant. Note that the chi-square likelihood-ratio test is much more relevant than a simple comparison of the nonlinear least squares R^2 s, which are all high (since even the 40bp RMSE of the single-factor model is very small relative to the huge cross-sectional variation in the tranche spreads, which range from fewer than 10 to over 1,000 basis points).

The three-factor version of the model results in very small RMSEs. With the exception of the CDX 1 index, the overall RMSEs are all on the order of two to three basis points. In fact, the RMSE for CDX 5 is actually less than one basis point. Again, with the exception of the CDX 1 index, the incremental explanatory power of the three-factor version relative to the two-factor model is highly significant. Thus, the three-factor model provides a very close fit to the data. Accordingly, we report results based on the three-factor version of the model in the remainder of the paper.

C. The Parameter Estimates

Table III reports the parameter estimates obtained from the three-factor model along with their asymptotic standard errors (Gallant (1975)). Focusing first on the estimates of the jump sizes, the table shows that there is strong uniformity across the different CDX indexes. In particular, the jump sizes associated with the first Poisson process are in a tight range from 0.00387 to 0.00469. Since each firm in the CDX index has a weight of $1/125 = 0.008$ in the index, a jump size of, say, 0.004 is consistent with the interpretation that a jump in the first Poisson process represents the idiosyncratic default of an individual firm, where the implicit recovery rate for the firm's debt is 50%. If

Table III
Parameter Estimates

This table reports the parameter estimates for the indicated CDX indexes. The jump size parameters are the parameters γ_1 , γ_2 , and γ_3 in the model. The volatility parameters are the σ_1 , σ_2 , and σ_3 parameters in the model. Asymptotic standard errors are in parentheses and are computed as in Gallant (1975). The sample period is from October 2003 to October 2005.

Index	Jump size parameters			Volatility parameters			<i>N</i>
	First	Second	Third	First	Second	Third	
CDX 1	0.00453 (0.00820)	0.06093 (0.00215)	0.17091 (0.00925)	0.03016 (3.98870)	0.10847 (0.11198)	0.00098 (2.18805)	65
CDX 2	0.00411 (0.00064)	0.06498 (0.00020)	0.35104 (0.11342)	0.20854 (0.04624)	0.19569 (0.00542)	0.14246 (0.13661)	108
CDX 3	0.00402 (0.00165)	0.06621 (0.00027)	0.35347 (0.10936)	0.11955 (0.20686)	0.16863 (0.00729)	0.18110 (0.10325)	118
CDX 4	0.00387 (0.00007)	0.05260 (0.00081)	0.51615 (0.26476)	0.14003 (0.01756)	0.25083 (0.00530)	0.16539 (0.23667)	127
CDX 5	0.00469 (0.00072)	0.05628 (0.00509)	0.33801 (0.10590)	0.17315 (0.5473)	0.27763 (0.06554)	0.29674 (0.06784)	17

we adopt this interpretation, then the recovery rates implied by the estimated jump sizes are 56.6%, 51.4%, 50.3%, 48.4%, and 58.6% for the individual CDX indexes, respectively.²²

The jump sizes for the second Poisson process are also very uniform across the CDX indexes, ranging from roughly 0.052 to 0.066. These values are consistent with the interpretation of the second Poisson process reflecting a major event in a specific sector or industry. As one way of seeing this, observe that virtually every broad industry classification is represented in the CDX index. In particular, the CDX index includes firms in the consumer durables, nondurables, manufacturing, energy, chemicals, business equipment, telecommunications, wholesale and retail, finance and insurance, health care, utilities, and construction industries. If we place the CDX firms into these 12 broad industry categories, then there are $125/12 = 10.42$ firms per category. Assuming a 50% recovery rate, a major event that resulted in the loss of an entire industry would lead to a total loss for the index of $10.42/125 \times 0.50 = 0.042$, which is on the order of magnitude of the jump size estimated for the second Poisson process. We note again, however, that a number of alternatives to this industry-event-risk interpretation could be equally valid.

The estimated jump sizes for the third Poisson process display somewhat more variation than for the other two Poisson processes, with values ranging

²² Historical recovery rates on corporate debt vary based on macroeconomic conditions, the seniority of the debt, the nature of the default, the rating of the issuer, and many other factors. For the senior unsecured debt referenced by the CDX indexes, the normal range of recovery between 1981 and the present has ranged from 20% to 70% according to Moody's (for example, see Gupton (2005)).

from about 0.17 to 0.52. The average value across all five indexes is about 0.35. Again assuming a 50% recovery rate, a jump size of 0.35 associated with a realization of the third Poisson process can be interpreted as a major economic shock to the entire economy in which as many as 70% of all firms default on their debt. This is a nightmare scenario that is difficult to imagine. Potential examples of such a scenario might include nuclear war, a worldwide pandemic, or a severe and sustained economic depression. The latter would need to be much more severe than any the United States has yet experienced, but has been observed elsewhere in a number of instances during the two-millennium-long experience of sovereign defaults and collapses in ancient Rome, Germany, Russia, and many other states (see Winker (1999)).

Turning now to the estimates of the volatility parameters, Table III shows that the volatility estimates of each of the three intensity processes are generally of the same order of magnitude. Specifically, with the exception of the first CDX index, the volatility parameters range from roughly 0.10 to 0.30 across all three processes and across all the CDX indexes. It is important to stress that these parameters are all estimated in sample, which leaves open the usual issue of how the model would fit out of sample. The fact that many of the estimated parameters are similar across different CDX indexes, however, provides some indirect support that the out-of-sample performance of the model might not be unreasonable. Finally, we note that the standard errors for a few of the parameters are large relative to the parameter estimates, particularly for the CDX 1 results and for the estimates of σ_3 . In general, however, most of the other parameters appear to be reasonably precisely estimated.

D. The Intensity Processes

D.1. The Time Series

Figure 2 plots the time series of the estimated values of the three intensity processes. Again, the estimated intensities are all under the risk-neutral measure. As shown, the first intensity process λ_1 ranges from roughly 0.50 to 1.50 during the sample period. For the majority of the sample period, this process takes values between 0.60 to 0.90 and displays a high level of stability. During the credit crisis of May 2005, however, this intensity process spiked rapidly to a value of 1.52, but then declined to just over 1.00 by the middle of June 2005. Thus, this spike was relatively short lived. The average values of λ_1 for the CDX 1 through CDX 5 indexes are 0.726, 0.854, 0.766, 1.023, and 0.816, respectively. Given the average value of λ_1 during the sample period, the expected (risk-neutral) waiting time until a firm-specific default is 1.16 years.

The second intensity process λ_2 ranges from a high of about 0.04 to a low of about 0.01 during the sample period. The value of this process is generally declining throughout the period. During the credit crisis, the value of this process doubled from about 0.015 to just over 0.030. After the crisis, the value of this process continued to decline. This suggests that the market-implied probability of a major industry or sector crisis declined




Figure 2. Intensity processes. This figure graphs the estimated intensity processes. The vertical division lines denote the roll from one CDX index to the next.

significantly during the past several years. Put another way, the expected waiting time for this type of event declined from roughly 28 years to 125 years during the sample period. The average values of λ_2 for the CDX 1 through CDX 5 indexes are 0.031, 0.035, 0.021, 0.016, and 0.009, respectively. The average (risk-neutral) waiting time for a realization of the second Poisson process is 41.5 years during the sample period.

The third intensity process λ_3 has more apparent variability across CDX indexes than do the other two intensity processes. In particular, the value of this process increases rapidly for the CDX 1 index, but then generally takes




Figure 3. Loss distribution function. This figure graphs the loss distribution implied by the fitted three-factor model for losses ranging from 0.00 to 0.16. The actual loss distribution is discrete, but is approximated by a continuous function in the graph.

lower values for the other four CDX indexes. The apparent discontinuity in this process as it rolls from CDX 1 to 2 is probably related to the higher standard errors of the estimated parameters for CDX 1; the estimated parameters and values of the intensity processes for CDX 1 are likely much noisier than for the other indexes. As with the second intensity process, the third intensity process essentially doubles around the time of the credit crisis. The average values of λ_3 for the CDX 1 through CDX 5 indexes are 0.0026, 0.0009, 0.0009, 0.0014, and 0.0010, respectively. The average value for this intensity process throughout the entire sample period is 0.00131. Thus, the implied risk-neutral probability of a catastrophic meltdown scenario is very small with an expected (risk-neutral) waiting time of about 763 years on average.

To illustrate the implications for the risk-neutral portfolio loss distribution, Figure 3 plots the time series of loss distributions implied by the model. Specifically, the distributions shown are for total portfolio losses at the 5-year horizon and are truncated to show only values ranging from zero to 16% (the probabilities for larger losses are visually difficult to distinguish from zero). As shown, the distribution of portfolio losses is multimodal and displays considerable time-series variation.

D.2. Interpreting the Factors

Although we have referred to the three Poisson processes as being consistent with idiosyncratic, industry or sector, and economywide credit events, respectively, it is important to stress that we have provided no direct evidence

supporting this interpretation. Intuitively, it seems reasonable to think of the first Poisson variable as an idiosyncratic credit event given that its realization maps into a portfolio loss of roughly 0.004. Similarly, it also seems natural to interpret the third Poisson event as a serious credit event affecting a large fraction of firms throughout the economy. In contrast, however, the second Poisson process need not necessarily be an industry or sector event. In fact, it could just as easily represent a default event for a subset of firms related by a variety of other firm attributes.

One possible way to explore the economic role played by the second Poisson process is by examining its implications for the factor structure of credit spreads for individual firms. To see this, imagine that all credit risk was purely idiosyncratic and that the correlation of credit spread changes across firms was zero. This is clearly not the case since the average correlation of daily credit spread changes across firms in the CDX index is 0.245. Similarly, imagine that all credit risk was economywide. In this polar extreme case, all credit spread correlations would be one, which is again easily rejected by the data. Now imagine that credit risk was a blend of both idiosyncratic and economywide risk, where the relative proportion varies across firms. In this case, credit spreads would be cross-sectionally correlated, but factor analysis would reveal that there was one common factor driving credit spreads—the remaining variation in credit spreads would be purely idiosyncratic.

With these preliminaries, now consider the more realistic case corresponding to the model estimated in this paper in which there is idiosyncratic and economywide credit risk, but also clustered default risk for subsets of firms related to the second Poisson process. Assume that this default clustering occurs across firms in a way that has nothing to do with their industry or sector. In this scenario, a principal components analysis would reveal a common economywide component driving individual firm credit spreads, and then a number of other common factors affecting specific subsets of the firms, but unrelated to industry grouping.

To explore this interpretation, we first map each of the 125 firms in the CDX indexes into one of the 12 Fama–French industry categories. Averaging over all five CDX indexes, 5.92% are in consumer nondurables, 3.36% in consumer durables, 10.24% in manufacturing, 5.44% in energy, 3.20% in chemicals, 6.40% in business equipment, 8.00% in telecommunications, 5.60% in utilities, 11.68% in wholesale/retail, 3.20% in healthcare, 22.24% in finance, and 14.72% in “other.” Next, we extract out time series of CDS spreads for the 94 firms that are present in the CDX indexes throughout the sample period and also have traded stock. We then compute the correlation matrix of daily credit spread changes for these firms and perform a principal components analysis. Finally, we regress the principal components (the corresponding eigenvectors) on industry dummy variables for each firm. If default clustering is unrelated to industry, these dummy variables should not have cross-sectional explanatory power for the second, third, fourth, etc. principal components.

The results provide a number of interesting insights into the cross-sectional structure of credit risk. Credit risk is obviously not purely idiosyncratic; the first

principal component explains more than 27% of the variation in credit spreads across firms. On the other hand, idiosyncratic risk appears to be the dominant nature of individual firm credit spreads. Specifically, the next five principal components only explain an incremental 5.1%, 4.5%, 3.5%, 3.1%, and 2.8%, respectively. Furthermore, eight principal components are required before more than 50% of the variation in credit spreads is explained.

Table IV (top panel) reports the results from the cross-sectional regression of the principal component weights on the industry dummy variables. As shown, the first principal component is consistent with the interpretation of an economywide credit variable affecting the majority of firms; all 12 of the industry dummy variables are highly significant. Although not shown, the regression coefficients for the industry dummy variables are remarkably uniform, ranging from about 0.08 to 0.11. Thus, the first factor can be viewed as a “parallel shift” in the credit spreads of all firms.

Moving beyond the first principal component, we can now test whether the default clustering in subsets of firms is related to industry categories. Recall that if the default clustering reflected by the second Poisson process has nothing to do with industry, then these principal components should be orthogonal to the industry dummy variables. In actuality, however, there appears to be a significant relation between many of the principal components and the industry dummies. For example, four of the industry dummy variables are significant for the second principal component and the corresponding adjusted R^2 is 0.419. The four significant industries are the manufacturing, energy, finance, and other industries. Similarly, the energy, telecommunications, and finance industry dummies are significant for the third principal component and the adjusted R^2 is 0.263. Interestingly, for the fourth through eighth principal components, only one or two of the industry dummy variables are significant at the 5% or 10% level. Thus, were the R^2 s for the regressions higher, there would be the possibility of almost a one-to-one mapping between these principal components and a specific industry or pair of industries.

It is important to provide some caveats at this point. For example, a number of the significant industry dummy variables have negative signs. These negative signs muddy the interpretation of the relation between principal components and specific industries. Furthermore, the adjusted R^2 s for many of the principal components are not high, indicating that industry grouping may only be a small part of the total picture in explaining the default clustering being captured by the second Poisson process. Despite these caveats, however, these results provide at least some evidence that industry does play a significant role in explaining the clustering of credit risk in subsets of firms.

To understand better how these results fit into a broader economic perspective, we repeat the same exercise using daily stock return data for the same 94 firms and sample period. The results are reported in the bottom panel of Table IV and indicate that there are many similarities between the results for credit spreads and those for the stock returns for these 94 firms. For example, the first principal component for stock returns explains about 26.9% of the variation in returns, which is almost the same amount explained

by the first principal component for credit spreads. As before, however, the second, third, fourth, etc. principal components explain only small proportions of the total variation, suggesting that much of the variability in stock returns is idiosyncratic.

The regression results indicate that the industry dummy variables have significant explanatory power for the stock return principal components. Similar to the results for changes in CDS spreads, the first principal component loads on all 12 of the Fama–French industry dummy variables, consistent with the usual view of the first factor in stock returns being related to the market. Where the results differ from those for credit spreads is in the implications for the other principal components. For example, the industry dummy variables explain more than 82% of the variation in the loadings for the second principal component. This is a much higher proportion than in the credit spread results. Furthermore, 8 of the 12 industry dummies have significant explanatory power for the second stock return principal component. Similar results hold for the third, fourth, fifth, etc. principal components: principal components for the stock returns are much more related to industry than is the case for credit spreads. Furthermore, few, if any, of the stock return principal components can be linked to one or two industries; stock return principal components seem to be related to broader subsets of firms in the economy than is the case for credit spreads. These results are intriguing and argue for a more in-depth comparison of the cross-sectional structure of credit spread changes and that of stock returns for the corresponding firms than we are able to provide in this paper.

E. CDX Index Spread Components

We can decompose the CDX index spread into three distinct components to measure the approximate overall economic impact of idiosyncratic, industry, and economywide default risks. In particular, the idiosyncratic component of the CDX index spread can be approximated by $\bar{\gamma}_1 \lambda_1$, the industry component by $\bar{\gamma}_2 \lambda_2$, and the economywide component by $\bar{\gamma}_3 \lambda_3$. The sum of these three components approximates the value of the CDX index spread.

For the CDX 1 through CDX 5 indexes, the idiosyncratic component of the spread represents 58.83%, 57.70%, 64.19%, 71.56%, and 82.17% of the total spread, respectively. Thus, although idiosyncratic default risk averages about two-thirds of the total value of the CDX index over the entire sample period, the percentage due to idiosyncratic default risk has increased steadily over time. In contrast, the portion of the CDX index spread due to industry or sector default risk declined significantly during the sample period, representing 33.34%, 37.02%, 28.88%, 15.92%, and 10.48% of the total CDX spread for the CDX 1 through CDX 5 indexes, respectively. Economywide default risk accounts for 7.83%, 5.28%, 6.93%, 12.52%, and 7.35% of the total spread for the CDX 1 through CDX 5 indexes, respectively. Thus, there is no clear trend in the size of this component. We note, however, that this component takes its largest value for the CDX 4 index, which spans the period during the May 2005 credit crisis.




Figure 4. Time series of RMSEs. This figure graphs the time series of daily root mean squared error (RMSE) from the fitting of the five CDX tranches. RMSEs are measured in basis points. The vertical division lines denote the roll from one CDX index to the next.

In summary, these results indicate that idiosyncratic default risk constitutes the majority of the CDX index spread. The combined effects of industry and economywide risk, however, are also significant and have represented more than 40% of the total at times. For the entire sample period, the average size of the three default components are 64.6%, 27.1%, and 8.3%, respectively, of the total CDX spread.

F. The Time Series of RMSEs

Although Table II reports summary RMSE statistics for the three-factor model, it is also interesting to examine the time-series variation in the ability of the model to capture market tranche spreads more closely. Accordingly, Figure 4 plots the time series of daily RMSEs obtained by fitting the model to the five tranches.

As shown, the ability of the model to match market tranche spreads increased significantly during the sample period. Initially, some of the RMSEs are as large as 19 basis points. The RMSEs decline rapidly, however, and

are on the order of five basis points by early 2004. By mid-2004, the RMSEs decline further and hover around two basis points for most of the sample period. The only exception is around the May 2005 credit crisis when the RMSE increases slightly to about five basis points. After the crisis, however, the RMSEs decline rapidly and reach values below one basis point near the end of the sample period. The small spikes in the RMSEs at the beginning and end of each index series are potentially due to investors rolling positions from tranches based on the previous index to tranches based on the new on-the-run index.

G. Pricing Errors

We turn next to the pricing errors, defined as the difference between the model-implied spreads and the market spreads for the various CDX index tranches, and examine their properties. Table V presents summary statistics and reports *t*-statistics for the significance of the average pricing errors.

As shown, the pricing errors from the three-factor model are surprisingly small across all indexes and tranches. In particular, the average pricing errors are all within 10 basis points of zero and most are within 1 or 2 basis points of zero. Recall from Table I that the average sizes for the 0–3%, 3–7%, 7–10%, 10–15%, and 15–30% tranche spreads are about 1,759, 240, 82, 34, and 12 basis points, respectively. Thus, average pricing errors of only a few basis points are extremely small in percentage terms as well. Except for the CDX 1 errors, the mean errors are not significant.

While the results on the size of the pricing errors are encouraging, it is important to acknowledge that the model is rejectable. For example, most of the first-order serial correlation coefficients are very large, indicating that there is a high degree of persistence in the pricing errors. Thus, in principle, it might be possible to construct a trading rule that exploits model mispricings.

H. Linking Firm-Level and Portfolio-Level Information

In this paper, we have focused on the implications of the data for the distribution of default losses for a large portfolio of credit-sensitive contracts. Ideally, we would like to be able to use portfolio-level information to infer something about the economic nature of credit risk at the individual firm level.

One way to do this is to solve for the default correlation among individual firms implied by the data. The event that firm *i* has defaulted by time *T* can be characterized as a simple binary or Bernoulli variate with probability $\pi_i = 1 - e^{-\xi T}$, where ξ is a firm-specific constant. With this structure, it is straightforward to show that the value of the premium leg for a 5-year

Table V
CDX Index Tranche Pricing Errors

This table reports summary statistics for the pricing errors for the indicated CDX index tranches. The *t*-statistic for the mean is corrected for first-order serial correlation. Pricing errors are measured in basis points. The sample period is from October 2003 to October 2005.

Tranche	Index	Mean	SD	<i>t</i> -Statistic for the mean	Serial correlation	<i>N</i>
0–3 Tranche	CDX 1	–7.00	8.56	–1.66	0.896	65
	CDX 2	–1.27	2.18	–1.33	0.916	108
	CDX 3	–0.19	0.97	–0.26	0.984	118
	CDX 4	–0.20	1.17	–0.55	0.850	127
	CDX 5	0.02	0.29	0.32	–0.019	17
3–7 Tranche	CDX 1	–3.21	3.34	–1.98	0.893	65
	CDX 2	–0.51	1.21	–1.16	0.880	108
	CDX 3	–0.21	–.73	–0.55	0.947	118
	CDX 4	0.00	1.19	0.01	0.894	127
	CDX 5	–0.03	0.28	–0.42	0.110	17
7–10 Tranche	CDX 1	–2.60	8.34	–0.66	0.887	65
	CDX 2	0.03	3.06	0.04	0.801	108
	CDX 3	0.61	3.32	0.30	0.968	118
	CDX 4	–0.27	4.53	–0.15	0.914	127
	CDX 5	0.16	1.12	0.51	0.152	17
10–15 Tranche	CDX 1	–6.62	8.05	–1.67	0.896	65
	CDX 2	–2.07	5.74	0.57	0.967	108
	CDX 3	–0.45	2.77	–0.22	0.984	118
	CDX 4	–0.27	2.26	–0.53	0.752	127
	CDX 5	–0.76	0.62	–0.31	0.486	17
15–30 Tranche	CDX 1	–8.09	1.88	–9.13	0.885	65
	CDX 2	–2.98	3.63	–1.66	0.936	108
	CDX 3	–0.83	1.90	–0.64	0.975	118
	CDX 4	–0.62	2.39	–0.72	0.892	127
	CDX 5	0.06	0.46	0.45	0.238	17

quarterly-pay CDS contract on the firm is

$$\frac{s}{4} \sum_{t=1}^{20} D(t/4) e^{-\xi t/4}, \tag{20}$$

and the value of the protection leg for the CDS contract is

$$\frac{w}{4} \sum_{t=1}^{20} D(t/4) \xi e^{-\xi t/4}, \tag{21}$$

where *s* is the CDS premium and *w* is the write-down fraction on the firm's debt in the event of a default. Thus, the value of ξ is given immediately from the CDS premium for the firm by the relation $\xi = s/w$.

The joint distribution of losses on the 125 firms in the CDX index is a multivariate correlated Bernoulli distribution. As discussed in Marshall and Olkin (1985), Park, Park, and Shin (1996), Lunn and Davies (1998), and many others, this distribution is very difficult to characterize in either closed form or via simulation. To solve for the implied correlation, however, we do not need to evaluate the joint density explicitly. In particular, assume that the event of a default of any firm translates into a loss fraction of 0.004 for the CDX portfolio. Given this structure, the variance of the loss distribution for the CDX portfolio at horizon T is given by

$$0.004^2 \sum_{i=1}^{125} \pi_i(1 - \pi_i) + 0.004^2 \sum_{i=1}^{125} \sum_{j=1, j \neq i}^{125} \rho_{ij} \sqrt{\pi_i \pi_j (1 - \pi_i)(1 - \pi_j)}, \quad (22)$$

where ρ_{ij} is the pairwise correlation coefficient for the Bernoulli variates for firms i and j . To solve for an implied correlation it is necessary to place some additional structure on the correlations. For simplicity, we assume that ρ_{ij} is constant for all i and j , $i \neq j$. With this assumption, it is now straightforward to solve for the variance of the portfolio loss distribution implied by CDO prices, set it equal to the above expression, and then solve for the implied default correlation.²³

Figure 5 plots the time series of the implied correlation estimates. The implied correlation typically ranges from about 0.05 to 0.10 during the sample period. The average value and standard deviation of the implied correlation during the sample period are 0.0835 and 0.0177, respectively. Around the credit crisis of May 2005, however, the implied correlation increases to about 0.13, but then rapidly declines. The lowest values for the implied correlation occur near the end of the sample period and are in the neighborhood of 0.04.

The expression in equation (22) also provides insight about how the cross-sectional structure of credit risk affects the portfolio loss distribution. From this expression, it is immediately clear that the variance of the loss distribution is an increasing function of the individual pairwise correlations. This is completely consistent with the usual portfolio intuition that as correlations increase, the portfolio is less-well diversified, resulting in a higher portfolio variance.

VI. Conclusion

This paper uses the information in the prices of synthetic CDX index tranches to study the market's expectations about how corporate defaults cluster in various economic environments—the cross-sectional structure of default risk. To do this, we first develop a new portfolio credit model in which three types of Poisson events generate portfolio credit losses. Using an extensive data set of CDX index and tranche spreads, we then estimate the model and evaluate its performance.

²³ The data for the CDS levels of the individual firms in the CDX indexes are also provided by Citigroup.




Figure 5. Implied default correlations. This figure graphs the implied default correlation among firms in the CDX index. These default correlations are given by solving for the single correlation parameter that sets the variance of the sum of the default losses for the individual firms in the CDX index, where the joint distribution of default events is distributed as a multivariate correlated Bernoulli distribution, equal to the implied variance of the distribution of portfolio losses after 5 years estimated by fitting the three-factor model to the data. The vertical division lines denote the roll from one CDX index to the next.

The results provide a number of insights into the important issue of default clustering. In particular, we find that the market expects significant clustering to occur. We show that roughly one-third of the value of the default spread for the typical firm in the CDX index is due to events in which multiple firms default together.

These results have a number of important economic implications. For example, they suggest that a significant portion of corporate credit risk may not be diversifiable. This has immediate implications for portfolio choice, the cost of corporate debt capital, and the systemic risk of financial institutions. Furthermore, since correlated default risk necessarily translates into correlated shocks to the stock values of the corresponding firms, these results may also have implications for the extreme risks being priced in equity markets.

Appendix

There are several ways in which the partial differential equation for P_i can be derived. For example, the approach outlined in Karlin and Taylor (1981, pp. 202–204) leads directly to the partial differential equation. To provide an alternative approach, recall that

$$P_i = E_t \left[\exp \left(- \int_0^T \lambda_s ds \right) \left(\int_0^T \lambda_s ds \right)^i \right]. \quad (\text{A1})$$

Let

$$H_t = \int_0^t \lambda_s ds. \quad (\text{A2})$$

This implies

$$dH_t = \lambda_t dt. \quad (\text{A3})$$

Now, rewrite P_i as

$$P_i = E_t \left[\exp \left(- \int_0^t \lambda_s ds - \int_t^T \lambda_s ds \right) \left(\int_0^t \lambda_s ds + \int_t^T \lambda_s ds \right)^i \right], \quad (\text{A4})$$

$$= E_t \left[\exp \left(-H_t - \int_t^T \lambda_s ds \right) \left(H_t + \int_t^T \lambda_s ds \right)^i \right]. \quad (\text{A5})$$

From this expression, P_i can be expressed explicitly as a function of λ_t , H_t , and $\tau = T - t$. An application of Itô's Lemma gives

$$\begin{aligned} dP_i &= (\alpha - \beta\lambda) \frac{\partial P_i}{\partial \lambda} dt + \sigma \sqrt{\lambda} \frac{\partial P_i}{\partial \lambda} dZ + \frac{\sigma^2 \lambda}{2} \frac{\partial^2 P_i}{\partial \lambda^2} dt \\ &\quad - \frac{\partial P_i}{\partial \tau} dt - \lambda P_i dt + i \lambda P_{i-1} dt. \end{aligned} \quad (\text{A6})$$

Since P_i is a martingale, however, the expected value of $dP_i = 0$. Thus,

$$\frac{\sigma^2 \lambda}{2} \frac{\partial^2 P_i}{\partial \lambda^2} + (\alpha - \beta\lambda) \frac{\partial P_i}{\partial \lambda} - \lambda P_i + i \lambda P_{i-1} = \frac{\partial P_i}{\partial \tau}, \quad (\text{A7})$$

which is equation (8) (at $t = 0$). For $i = 0$, the boundary condition is $P_0(\lambda, 0) = 1$. For all other i , the boundary condition is $P_i(\lambda, 0) = 0$.

For the case $i = 0$, the solution to the partial differential equation is identical to that provided by Cox et al. (1985) in obtaining zero-coupon bond prices. The solution for this case can be expressed as shown in equation (9) with $i = 0$. For $i > 0$, we conjecture that the solution is of the form shown

in equation (9). Differentiating the conjectured solution for $P_i(\lambda, T)$, substituting into equation (8), and collecting terms in the powers of λ leads to the system of first-order differential equations shown in equations (12), (13), and (14). This system can be solved recursively following in the order $C_{1,1}$, $C_{1,0}$, $C_{2,2}$, $C_{2,1}$, $C_{2,0}$, $C_{3,3}$, $C_{3,2}$, $C_{3,1}$, $C_{3,0}$, etc. Thus, for each i , we solve for $C_{i,j}$, where j runs backwards from i to zero.

REFERENCES

- Bakshi, Gurdip, Dilip Madan, and Frank Zhang, 2006, Investigating the role of systematic and firm-specific factors in default risk: Lessons from empirically evaluating credit risk models, *Journal of Business* 79, 1955–1988.
- Collin-Dufresne, Pierre, Robert S. Goldstein, and J. Spencer Martin, 2001, The determinants of credit spread changes, *Journal of Finance* 56, 2177–2207.
- Cox, John, Jonathan Ingersoll, Jr., and Stephen Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–407.
- Das, Sanjiv, Darrell Duffie, Nikunj Kapadia, and Leandro Saita, 2007, Common failings: How corporate defaults are correlated, *Journal of Finance* 62, 93–117.
- Das, Sanjiv, Laurence Freed, Gary Geng, and Nikunj Kapadia, 2006, Correlated default risk, *Journal of Fixed Income* Fall, 7–32.
- Demarzo, Peter, 2005, The pooling and tranching of securities: A model of informed intermediation, *Review of Financial Studies* 18, 1–35.
- Diamond, Douglas, 1984, Financial intermediation and delegated monitoring, *Review of Economic Studies* 60, 383–414.
- Duffie, Darrell, Andreas Eckner, Guillaume Horel, and Leandro Saita, 2006, Frailty correlated default, Working paper, Stanford University.
- Duffie, Darrell, and Nicolae Gârleanu, 2001, Risk and valuation of collateralized debt obligations, *Financial Analysts Journal* 57, 41–59.
- Duffie, Darrell, Leandro Saita, and Ke Wang, 2007, Multi-period corporate default prediction with stochastic covariates, *Journal of Financial Economics* 83, 635–665.
- Duffie, Darrell, and Kenneth J. Singleton, 2003, *Credit Risk: Pricing, Measurement, and Management* (Princeton University Press, Princeton, NJ).
- Elton, Edwin J., Martin J. Gruber, Deepak Agrawal, and Christopher Mann, 2001, Explaining the rate spread on corporate bonds, *Journal of Finance* 56, 247–277.
- Eom, Young Ho, Jean Helwege, and Jing-Zhi Huang, 2004, Structural models of corporate bond pricing: An empirical analysis, *Review of Financial Studies* 17, 499–544.
- Errais, Eymen, Kay Giesecke, and Lisa Goldberg, 2006, Pricing credit from the top down with affine point processes, Working paper, Stanford University.
- Frey, Rudiger, and Jochen Backhaus, 2005, Portfolio credit risk models with interacting default intensities: A Markovian approach, Working paper, Department of Mathematics, University of Leipzig.
- Gallant, A. Ronald, 1975, Nonlinear regression, *The American Statistician* 29, 73–81.
- Giesecke, Kay, 2004, Correlated default with incomplete information, *Journal of Banking and Finance* 28, 1521–1545.
- Giesecke, Kay, and Lisa Goldberg, 2005, A top down approach to multi-name credit, Working paper, Cornell University.
- Gupton, Greg M., 2005, Advancing loss given default prediction models: How the quiet have quickened, *Economic Notes* 34, 185–230.
- Hull, John, and Alan White, 2003, Valuation of a CDO and a n -th to default CDS without Monte Carlo simulation, Working paper, University of Toronto.
- Jorion, Philippe, and Gaiyan Zhang, 2005, Intra-industry credit contagion: Evidence from credit default swap and equity markets, Working paper, University of California at Irvine.

- Karlin, Samuel, and Howard Taylor, 1981, *A Second Course in Stochastic Processes* (Academic Press Inc, Orlando, Florida).
- Kindleberger, Charles P., 2005, *Manias, Panics, and Crashes: A History of Financial Crises*, 5th edn. (John Wiley & Sons, Hoboken, NJ).
- Liu, Jun, Francis A. Longstaff, and Ravit E. Mandell, 2006, The market price of risk in interest rate swaps: The roles of default and liquidity risks, *Journal of Business* 79, 2337–2359.
- Longstaff, Francis A., Sanjay Mithal, and Eric Neis, 2005, Corporate yield spreads: Default risk or liquidity? New evidence from the credit default swap market, *Journal of Finance* 60, 2213–2253.
- Lunn, A. Daniel, and Stephen J. Davies, 1998, A note on generating correlated binary variables, *Biometrika* 85, 487–490.
- Marshall, Albert, and Ingram Olkin, 1985, A family of bivariate distributions generated by the bivariate Bernoulli distribution, *Journal of the American Statistical Association* 80, 332–338.
- Park, Chul G., Taesun Park, and Dong W. Shin, 1996, A simple method for generating correlated binary variates, *The American Statistician* 50, 306–310.
- Rajan, Arvind, Glen McDermott, and Ratul Roy, 2007, *The Structured Credit Handbook* (John Wiley & Sons, Hoboken, NJ).
- Richardson, Karen, 2005, Bankers hope for a reprise of “Bowie Bonds”, *Wall Street Journal*, August 23.
- Roy, Ratul, and David Shelton, 2007, Trading credit tranches, in Arvind Rajan, Glen McDermott, and Ratul Roy, eds., *The Structured Credit Handbook*, (John Wiley & Sons, Hoboken, NJ).
- Saita, Leandro, 2005, The puzzling price of corporate default risk, Working paper, Stanford University.
- Schönbucher, Philipp, 2005, Portfolio losses and the term structure of loss transition rates: A new methodology for the pricing of portfolio credit derivatives, Working paper, Department of Mathematics, ETH Zürich.
- Sidenius, Jakob, Vladimir Piterbarg, and Leif Andersen, 2005, A new framework for dynamic credit loss modeling, Working paper, Royal Bank of Scotland.
- Winkler, Max, 1999, *Foreign Bonds: An Autopsy* (Frederick, MD: Beard Books).
- Yu, Fan, 2005, Default correlation in a reduced-form model, *Journal of Investment Management* 3, 33–42.
- Yu, Fan, 2007, Correlated defaults in intensity-based models, *Mathematical Finance* 17, 155–173.