Kamakura Corporation

2222 Kalakaua Avenue
Suite 1400
Honolulu, Hawaii 96815, USA
phone 1 808 791 9888
facsimile 1 808 791 9898
www.kamakuraco.com

Advanced Financial Risk Management:

Tools and Techniques for Integrated Credit Risk, Market Risk, Liquidity Risk and Interest Rate Risk Management

April 2012

Donald R. van Deventer
Kenji Imai
Mark Mesler

Part 4: Risk Management Applications, Instrument by Instrument

Chapter 20
Credit derivatives and collateralized debt obligations

Introduction

The credit crisis of 2007 to 2010 represented one of the greatest episodes of financial hysteria seen since the Nikkei stock price index (which traded at 8,734.62 on June 25, 2012) reached an intra-day peak of 38,957.44 on December 29,1989. The introduction to this book summarizes much of this mass hysteria, and it is very difficult to resist the temptation to make this a "what were they thinking?" chapter when recounting the tens of billions of dollars lost in the credit markets during this period. As best we can, we will try to re-emphasize some common risk management fallacies that we raised in the introduction to this book. We reprise those fallacies in this chapter with specific reference to credit default swaps ("CDS") and collateralized debt obligations ("CDOs") and how Wall Street responded to market participants who clung to these fallacies as if they were true.

- "If it hasn't happened to me yet, It won't happen to me, even if it's happened to someone else." Wall Street response: if everyone thinks home prices won't go down, sell them 100% loan to value mortgage loans until they learn otherwise.
- "Silo risk management allows my firm to choose the 'best of breed' risk model for our silo." Wall Street response: if the credit silo can't model the value of securities whose default is tied to macro factors like home prices, sell them mortgage-related CDOs.
- "I don't care what's wrong with the model. Everyone else is using it." Wall Street response: after convincing many to use the flawed copula method, Wall Street stuffed investors with CDOs until they finally figured out the model didn't work.
- "I don't care what's wrong with the assumptions. Everyone else is using them." Wall Street response: if everyone is using historical value at risk during a period of rising home prices, work fast to stuff them with mortgage-backed securities after home prices begin to fall.
- "Mathematical models are superior to computer simulations." Wall Street response: try to stuff the firms relying on "F9 Model Monkeys" (see the Felix Salmon 2012 reference below) before they realize that only the kind of sophisticated simulation we describe below can accurately measure the risk.
- "Big North American and European banks are more sophisticated than other banks around the world and we want to manage risk like they do." Wall Street has long used the sales strategy "real men buy [insert the name of an over-priced derivative]" with great success. Once "buy side" firms stopped buying CDOs, traders used this strategy on their own management teams, loading up CDO structuring firms like Citigroup and Merrill Lynch with unsalable CDO tranches in the hopes that bonuses would be paid before their strategy was discovered.
- "Goldman says they do it this way and that must be right." Naïve investors have long regarded their investment bankers as someone with whom they need "a relationship." Too many look at a Wall Street salesman wearing a suit and envision that person as a saintly combination of wise finance professor and Mother Theresa. Sadly, the investment bankers have proven in the 2007 to 2010 credit crisis that their view of a "client relationship" is the same type of relationship that they have with a cheeseburger at McDonalds (where they rarely have to eat).

Too many journalistic classics have emerged from this human comedy to summarize them all here. The authors recommend a few classics, however, both because of their timely warning and their "plain English" warning of what can and did go wrong in the CDS and CDO markets. We list them here in chronological order:

Mark Whitehouse, "Slices of Risk: How a Formula Ignited Market That Burned Some Big Investors," Wall Street Journal, page 1, August 12, 2005.
Felix Salmon, "Recipe for Disaster: The Formula that Killed Wall Street," Wired Magazine, February 23, 2009
Kevin Drum, "The Gaussian Copula," Mother Jones, February 24, 2009

Sam Jones, "Of Couples and Copulas," Financial Times, April 24, 2009
Lisa Pollack, “On the clairvoyance of sovereign CDS “, Financial Times, October 10, 2011
Robert A. Jarrow, “Problems with Using CDS to Infer Default Probabilities,” Cornell University and Kamakura Corporation memorandum, www.kamakuraco.com, January 25, 2012
Donald MacKenzie and Taylor Spears, “‘The Formula That Killed Wall Street’?
The Gaussian Copula and the Material Cultures of Modeling,” University of Edinburgh, June, 2012
Felix Salmon, “The dangerous Gaussian copula function,” www.reuters.com, June 21, 2012

With this real world background out of the way, we turn to the theory and practice of the credit default swap and collateralized debt obligation markets. As we often do during this volume we highlight flawed approaches with the “MODEL RISK ALERT” label.

Credit default swaps: theory

The focus of this chapter is on pricing, valuation, hedging and simulation CDS and CDOs with the same ultimate goal of this volume: being able to measure and hedge risk in a fully integrated way, as captured by the Jarrow-Merton put option that we discussed beginning in Chapter 1. We will not focus on the nitty-gritty details of transaction terms and structures, which are covered in great depth by a number of other others including Tavakoli [2001], Ranson [2003], Meissner [2008], and Rutledge and Raynes [2003]. Our focus here is on analyzing credit derivatives in the context of a highly accurate enterprise wide risk management framework with a complete integration of credit risk, market risk, liquidity risk, and interest rate risk.

Credit default swaps have been highly touted as a credit risk management tool because of the difficulty in hedging single-name credit risk by shorting the common stock of the issuer, which we discussed in Chapter 18, or the publicly traded bonds of the issuer. Clearly, if shorting the bonds were easy to do, one would either sell the debt issue being hedged outright or sell it short. Short sales of risky debt are rarely an option because of the illiquidity of the bond market (which is explicitly incorporated in the Jarrow model of risky debt discussed in Chapter 16). For that reason, the development of the credit default swap market was initially greeted with enthusiasm by market participants over the last decade. This enthusiasm has faded, however, as we explain in the next section.

Two credit default swap structures are common in the market place. In the first structure, which we call the traditional structure, the credit protection provider pays the purchaser of credit protection the “recovery amount” associated with the notional principal of the credit default swap with a specified delay after a specified time frame from the carefully defined event of default. This recovery amount is determined by an auction process overseen by ISDA.org. In the second structure, the digital default swap, the provider of credit protection essentially pays $1 upon the occurrence of an event of default. By comparing the pricing on the two structures for the same issuer and the same maturity, one can imply the recovery rate and loss given default implicit in market prices. The latter structure is quite rare but it is of great academic interest because it ignores the thorny issues of recovery.

The conceptual benefits of these derivatives structures to financial institutions are powerful, but we summarize the increasingly dramatic problems in the next section: lack of liquidity, market concentration, and potential market manipulation like that mentioned in Chapter 17 with respect to the London Interbank Offered Rate market.

Initially, the market for credit default swaps called for a periodic payment in advance (the insurance premium) by the purchaser of credit protection. Upfront insurance payments by the purchaser of credit protection were restricted to troubled reference credits. After the credit default swap market "big bang" in 2009, upfront payments with more modest periodic payments became the norm for credits of all quality ranges.

Jarrow [1999] provides a valuation framework for credit default swaps that explicitly allows for the fact that there may be illiquidity in the bond markets which is in return reflected in risky bond prices. We reviewed the essentials of the Jarrow model in Chapter 16. The hazard rate, or "default intensity"

in the Jarrow model is given by a linear combination of three terms:

$$\lambda(t) = \lambda_0 + \lambda_1 r(t) + \lambda_2 Z(t)$$

Recall that the first term, λ_0, can be made time dependent, extending the credit model exactly like we extended term structure models in Chapters 6 through 9 and Chapter 13. The r(t) term is the short rate of interest in a sophisticated term structure model. The term *Z(t)* is the "shock" term with mean zero and standard deviation of one which creates random movements in the macro factor(s) (like home prices, in the case of the 2007 to 2010 credit crisis), which drives default for that particular company.[1] Movements in this macro factor are generally written in this form:

$$dM(t) = M(t)[r(t)dt + \sigma_m dZ(t)]$$

We start our valuation review with the simplest case: a digital default swap with a constant continuous "default intensity." Jarrow shows that the value of a digital default swap when the default intensity λ_0 is constant can be written

$$I_t(T)1_{\{\tau>t\}} = \lambda_0 \int_t^T v(t,s:e)ds = \lambda_0 \int_t^T p(t,s)e^{-\lambda_0(s-t)}ds$$

In this expression, the expression v (t,s, i) refers to the time t zero coupon bond of the risky issuer with maturity s and seniority i. The expression p(t,s) is a risk-free zero coupon bond as of time with maturity s. The integrals in this expression recognize that default can occur any time between time t and the maturity of the credit protection at time T. Jarrow shows how this expression changes as interest rates and macro factors are allowed to drive the default intensities. The first of the two

[1] In more mathematical terms, Z(t) is standard Brownian motion under a risk neutral probability distribution Q with initial value 0 that drives the movements[1] of the market index *M(t)*

equivalent expressions makes it clear that the value of the digital default swap is the continuous time equivalent of the promise to pay $1 at each point in time that default could occur, weighted by the probability can occur at that time. When the default swap is the more conventional type, where the amount of recovery δ comes into play, the formula needs to be modified. The nature of the modification depends on the assumptions about recovery: both timing and sensitivity of recovery to the macro-economic factors that drive default. For a detailed discussion, see Jarrow [1999] and Chapter 16. In the Jarrow-Turnbull model, we discussed the valuation of a zero coupon bond in which the holder received either $1 at maturity or a recovery δ. The Jarrow Turnbull model can be thought of as a "credit default swap" zero coupon bond where the credit protection provider pays $0 (instead of $1) if there is no default and δ if there is a default. The value of this simple credit default swap is a slight variation on the Jarrow-Turnbull defaultable zero coupon bond valuation formula from Chapter 16:

$$v(t,T) = \left[e^{-\lambda\mu(T-t)} + (1 - e^{-\lambda\mu(T-t)})\delta\right]P(t,T)$$

The first term in brackets is the probability of no default during the term of the zero coupon bond. The second term in brackets is probability of default during the term of the bond. If there is no payoff in the event of no default, the value of a promise to pay a credit insurance payment of δ at the scheduled maturity date (instead of the date of the event of default) if default occurs during the life of the bond is

$$v(t,T) = \left[(1 - e^{-\lambda\mu(T-t)})\delta\right]P(t,T)$$

The Jarrow formulation discussed earlier allows for payment of $1, not δ, and the credit payment is made at the time of default, not the scheduled maturity date.

In modeling the valuation, loss distribution, and cost of the Jarrow-Merton put option of a credit portfolio that consists of many underlying borrowers or "reference names," market participants frequently confuse the meaning of the term correlation. Correlation can mean three things, and in two of the three cases the term has a model-specific meaning. Jarrow and van Deventer (2005) review the links between these correlation concepts:

Correlation in the value of company assets in the Merton framework, which is definitely not the same as the other two meanings of the term correlation. Note that the correlation in the value of company assets between firms A and B can be 100% but the correlation in the events of default will not be 100% except when the two firms A and B have identical default probabilities.

Correlation in default intensities due to dependence on common macro factors like interest rates and the S&P 500 in a U.S. context, the r(t) and Z(t) terms given in the Jarrow formulas above. The correlation between the default intensities of companies A and B in the reduced form modeling context depends on the magnitude of the parameters λ_0, λ_1, and λ_2 for each of the two companies A and B.

Correlation in the events of default, not the default probabilities themselves. This is the correlation between the two vectors for Companies A and B which have a zero in a period when the company did not default and a 1 when the company does default. For two solvent companies, obviously, these historical time series vectors have nothing but zeros so the historical correlation in the events of default cannot be implied by the historical data on events of default for Companies A and B.

We will try to make these distinctions clear in what follows.

How does one model total credit risk in a portfolio of companies in a reduced form model? Jarrow [1999] shows that the present value of a contract which pays $1 if both firms 1 and 2 default prior to time T has the following formula, where the subscripts for 1 and 2 denote the firm:

$$V_t 1_{\{\tau_1>t,\tau_2>t\}} = p(t,T) - E_t\left(e^{-\int_t^T [r(u)+\lambda_1(u)]du}\right) - E_t\left(e^{-\int_t^T [r(u)+\lambda_2(u)]du}\right) + E_t\left(e^{-\int_t^T [r(u)+\lambda_1(u)+\lambda_2(u)]du}\right).$$

The expression E_t denotes the risk neutral expected value of the expressions in brackets as of time t. Jarrow does on to show how closed form solutions can be obtained to value basket swaps in the reduced form model context. We can arrive at the same conclusion by simulating the macro factors driving the default intensities in the reduced form models, simulating default/no default in a multi-period way, and taking the correct risk adjusted present value of the cash flows that result. We outline this procedure below in the context of collateralized debt obligations.

The Jarrow formula above shows us how to model counterparty credit risk. Assume that Company A is buying digital default swap protection on Company C with Company B as the insurance provider. We know that this transaction pays

- 0 if Company A defaults before Company B or Company C (since Company A defaults on the credit insurance premiums)
- 0 if Company B defaults before Company C
- $1 if Company C defaults (a) before Companies A and B but (b) before the credit protection contract expires

We can use this knowledge to simulate the counterparty credit risk using either the Merton or Jarrow reduced form model framework discussed above. Implicit in the formulas above are a very key assumption--that markets are perfectly liquid and that no volume of trading in credit default swaps or risky debt will change their price (MODEL RISK ALERT). Recent events involving JPMorgan illustrate the tenuous nature of this assumption. We turn to the implications in the next section.

Credit default swaps: practice

An earlier version of this section appeared in the BankThink column of www.AmericanBanker.com on June 11, 2012 under the title "Credit default swaps: market or mud pit?" The authors wish to thank the editors for many helpful comments. The May 2012 sight of the mighty JPMorgan Chase, wallowing in a credit default swap position from which it cannot exit, raises many questions about the valuation formulas from the previous section. We know from Chapter 17 that the credit spread, both in the bond and CDS "markets," reflects the intersection of supply and demand. None of the complications discussed in Chapter 17 is represented in the digital default swap formula above, which assumed perfectly liquid markets with no impact on the price of credit protection from trade size. One of the key real world questions is a particularly simple one: is the arena for credit default swap trading a market place or a mud pit? In this section, we sadly conclude that the answer is "mud pit" and outline some of the very many steps that are necessary for credit default swap arena to become a market place where the assumptions of the previous section have a better chance of being accurate. Without these reforms, the CDS mud pit is likely to suffer the same fate as collateralized debt obligation-total irrelevance, as we discuss below.

JPMorgan's May 8, 2012 announcement of $2 billion in losses from a trade from with the firm could not easily exit turned the spotlight on the arena for credit default swaps. How can a sophisticated market participant like JPMorgan Chase end up suffering such discomfort and illiquidity in the U.S. financial markets? We explain that in this section. We focus on single name credit default swaps, not the credit default swap index market in which JPMorgan has suffered its losses. The reason for our focus is also simple: the CDS index market is built on the foundation of the single name CDS market. If the single name CDS market is a mud pit, so is the index market.

What do we mean when we ask if the CDS market is a mud pit? We are in essence asking the following questions:

- Is there a transparent and open transmission of traded prices, bids and offereds in real time by reference name like there is for common stocks and options?
- Is there reliable reporting of trade volume in real time like there is for common stocks and options?
- Is there sufficient competition in the market to prevent market manipulation by a concentrated set of dealers, like there is for common stocks?
- If the market is highly concentrated, is there sufficient regulatory scrutiny to prevent rampant market manipulation?
- Is there sufficient regulatory scrutiny to prevent CDS trading on inside information?

Sadly, the answer to all of these questions is no. The unfortunate answer to another key question is "yes"--that question is "Has there been tacit cooperation among market participants and data vendors to preserve the status quo in the CDS mud pit?"

We briefly deal with each of these questions in turn.

Is there a transparent and open transmission of traded prices, bids and offereds in real time by reference name like there is for common stocks and options?

In spite of changes in reporting by the Depository Trust & Clearing Corporation (www.dtcc.com) in recent years, the single name CDS bids, offereds, and traded prices are tightly controlled by a small cadre of dealers, market data vendors, and DTCC itself. The data is not freely available on a website like Yahoo or Google in the way that common stock prices or even bond prices are. While one broker who distributes CDS data does time stamp bids, offereds and trades in its data service, they are the exception, not the rule, and they see a relatively small market share of total activity. The primary product offered by the largest CDS data vendor does not distinguish between bids, offereds, and traded prices. Moreover, the data service typically lists CDS spreads for 2,000 corporates and sovereigns even though there have never been more than 1,000 reference names traded in a given week since DTCC began reporting trade volume (but not bids, offereds, or traded prices) since the week ended July 16, 2010.

Is there reliable reporting of trade volume in real time like there is for common stocks and options?

No. DTCC has this information, but it makes trade volume available on a weekly basis with a 2 or 3 business day lag, starting with the week ended July 16, 2010. Through January, 2012, the terms of use agreement for that volume data made it illegal for the authors to tell readers how many single name CDS were traded in JPMorgan because of this phrase:

> "You agree to treat any Report containing data on a specific entity, rather than aggregate position or transaction activity or other aggregate data, as confidential, or, if you are a regulator or governmental entity, in accordance with any statutory confidentiality requirements applicable to you."

Fortunately, since January, 2012, this phrase has been removed.

Is there sufficient competition in the market to prevent market manipulation by a concentrated set of dealers, like there is for common stocks?

No. In a recent blog, using the standard measures of market concentration by the Department of Justice and volume data from the Office of the Comptroller of the Currency, we showed that the market for credit default swaps has been "highly concentrated" since its inception. For the full text, see this link:

http://www.kamakuraco.com/Blog/tabid/231/EntryId/371/The-Credit-Default-Swap-Market-and-Anti-Trust-Considerations.aspx

The European Commission announced on April 29, 2011 that it was launching two anti-trust investigations into the CDS market:

http://europa.eu/rapid/pressReleasesAction.do?reference=IP/11/509&format=HTML&aged=0&language=EN&guiLanguage=en

If the market is highly concentrated, is there sufficient regulatory scrutiny to prevent rampant market manipulation?

No. In fact, a courageous journalist recently described how to manipulate the credit default swap market in this post from three years ago:

http://www.marketrap.com/article/view_article/91172/did-the-markit-group-a-black-box-company-partially-owned-by-goldman-sachs-and-jp-morgan-devastate-markets

Is there sufficient regulatory scrutiny to prevent CDS trading on inside information?

No. The SEC launched the first action alleging insider trading in 2009, but it was ultimately unsuccessful:

http://www.sec.gov/news/press/2009/2009-102.htm

The fact that insider trading is a serious problem was highlighted by this announcement from the UK's Financial Services Authority in 2011:

http://www.bloomberg.com/news/2011-03-25/fsa-said-to-review-credit-default-swaps-in-market-abuse-probe.html

That leads us to perhaps the most saddening question of those posed above:

Has there been tacit cooperation among market participants and data vendors to preserve the status quo in the CDS mud pit?

Yes. Perhaps the most egregious form of cooperation is the effort to preserve the impression that there is active trading in a large number of reference names when in fact there is not. For a review of trading volume reported by the DTCC, see the following links:

All Reference Names:

http://www.kamakuraco.com/Blog/tabid/231/EntryId/365/CDS-Trading-Volume-for-1-090-Reference-Names.aspx

U.S. Banks:

http://www.kamakuraco.com/Blog/tabid/231/EntryId/366/Credit-Default-Swaps-and-Deposit-Insurance.aspx

Sovereigns:

http://www.kamakuraco.com/Blog/tabid/231/EntryId/368/Sovereign-Credit-Default-Swap-Trading-Volume.aspx

Sub-Sovereigns and Municipals:

http://www.kamakuraco.com/Blog/tabid/231/EntryId/367/Municipal-Credit-Default-Swap-Trading-Volume.aspx

Breathless reporters or rating agencies claim "Dell's CDS widen 42%" when, in fact, there were only 9.6 gross trades per day and 1.75 non-dealer trades per day in Dell during the week ended May 25, 2012, according to DTCC:

http://finance.yahoo.com/news/fitch-solutions-dells-cds-widen-120600373.html

Reporters need a story, and the CDS mud pit provides material. Rating agencies need a product that is not a rating, and the CDS mud pit provides one. CDS data vendors couldn't sell CDS data if it were well known that (a) 81.68% of trades in the DTCC trade warehouse were trades between dealers or (b) more than 50% of the reference names for which data were reported by the vendor had NO trades in a given week, on average.

Existing beneficiaries of the CDS mud pit have no incentive to change things and lots of incentives not to change things. These changes must be forced by laws and regulations. The necessary changes should emulate the information and data flow of the market for common stocks and options. These markets, while not yet perfect, are one million times more transparent and competitive than the CDS mud pit:

- Distribute data through a large number of data vendors as the New York Stock Exchange does
- Display bid levels and bid volumes in real time
- Display offered levels and offered volumes in real time
- Display traded prices and traded volumes in real time
- Require disclosure of all entities with large open positions, both long and short, as is required in the market for common stock
- Aggressively enforce insider trading rules
- Force trading onto legitimate exchanges. Industry resistance to this is long standing but that is irrelevant. The investors' interest is of greater import.
- Use the Department of Justice whenever necessary to establish competition in the market place

Until then, whenever one sees a credit default swap quotation, it should be regarded with an extremely high degree of suspicion. With this sobering information as background, we now turn to the rise and fall of the market for collateralized debt obligations.

Collateralized debt obligations: theory

Collateralized Debt Obligations are generalized versions of first to default swaps, second default swaps, third to default swaps and so on. In a typical CDO structure, there may be 100 to 150 underlying reference names with a (roughly) equal amount of (actual or notional) principal associated with each of the reference names. Synthetic CDOs have single name credit default swaps as the underlying reference instruments. "Cash flow" CDOs have financial instruments like bonds, mortgage backed securities, asset backed securities, commercial mortgage backed securities or other instruments as the underlying collateral. If the underlying collateral is in the form of bonds issued by the 100 reference names, the proceeds of the issuance of senior, subordinated and equity tranches of the CDO are used to buy these bonds. If the underlying "collateral" is in the form of credit default swaps (a "synthetic" CDO), then the proceeds of the issuance of the tranches are used to buy a bond of a single issuer, the proceeds of which are paid to the CDO tranche holders over time depending on how many of the reference names default. The fact that all of the cash in the synthetic structure is invested in one reference name is a huge factor in differentiating between the so called "cash flow" CDOs (the bond structure) and the synthetic CDO structure. This concentration of risk must not be ignored in CDO analysis.

The value and cash flow generated by the underlying reference names in the CDO is identical in its nature to what we discussed in Chapter 19, when we discussed simulating future values of credit risky bonds, and to the first part of this chapter. The only difference is that the cash flow is parceled out according to the priorities of the tranches. If there is only one tranche (i.e. the equity tranche) CDO analysis is no different from credit risk portfolio analysis. If there are two tranches, the first cash flows go to the senior tranche as specified in the CDO indenture ("the waterfall") and the remaining cash flows (if any) go to the equity tranche, and so on. During the heyday of the CDO market, waterfall structures with 5, 10 or 15 layers or tranches were common.

The simulation of which cash flows go to whom is at the heart of CDO analysis, as is the simulation of correlated defaults. The investment banking community and the rating agencies have aggressively advocated a method of analysis that dramatically understates the worst case losses from a CDO portfolio, and users of this analysis should recognize that (if they take the analysis seriously) they will grossly overpay for CDO tranches. See the last section of this chapter, on the copula approach, for the details of this legacy calculation. Felix Salmon (2012) labeled the die-hard users of this approach "F9 Model Models" in honor of the key in common spreadsheet software that would initiate the lumbering spreadsheet calculation of the copula method.

The correct procedure for valuing credit default swaps (traditional or digital), first to default swaps and CDOs is to follow a slight variation on the process we outlined in a corporate bond context in Chapter 19. For the reduced form model, the steps are as follows:

1. Simulate the risk-free term structure of interest rates as given in Chapters 6 through 9 and summarized in Chapter 10. This should be a simulation with 5 to 10 risk factors driving the term structure of interest rates.
2. Choose a formula and the risk-drivers for the liquidity component of credit spread as discussed in Chapter 17
3. Choose a formula and the risk-drivers for the default intensity process in the Jarrow model as discussed in Chapter 16. The macro factors driving correlated risk can be estimated using historical default data and historical macro factors using logistic regression or from observable market prices of bonds. Twenty-six macro factors are embedded in the instructions for the Comprehensive Capital Analysis and Review 2012 of the Federal Reserve. They provide a good starting point for any serious analyst of risk:
    a. BBB Corporate Yield
    b. CPI Inflation Rate
    c. Developing Asia Inflation Rate
    d. Euro Area Inflation Rate
    e. Japan Inflation Rate
    f. UK Inflation Rate
    g. Commercial Real Estate Price Index
    h. Dow Jones Total Stock Market Index
    i. Developing Asia Bilateral Dollar Exchange Rate
    j. Euro Area Bilateral Dollar Exchange Rate
    k. Japan Bilateral Dollar Exchange Rate
    l. UK Bilateral Dollar Exchange Rate
    m. House Price Index
    n. Mortgage Rate
    o. Nominal Disposable Income Growth
    p. Nominal GDP Growth
    q. Real Disposable Income Growth
    r. Real GDP Growth
    s. Developing Asia Real GDP Growth
    t. Euro Area Real GDP Growth
    u. Japan Real GDP Growth

   v. UK Real GDP Growth
   w. 10-Year Treasury Yield
   x. 3-Month Treasury Yield
   y. Unemployment Rate
   z. Market Volatility Index - VIX
4. Simulate the random values of the drivers of liquidity risk and the default intensity for M time periods over N scenarios, consistent with the risk free term structure. MODEL RISK ALERT: *One should take care not to make common risk management assumptions (i.e. that the macro factors have normally distributed returns and are independent from period to period) when those common assumptions are not true.*
5. Calculate the default intensity and the liquidity component at each of the M time steps and N scenarios. Note again, as in Chapter 19, that these will be changing randomly over time because interest rates and other macro factors are changing at each time step of each scenario. The simulation of these macro factors is essential to capturing cyclicality in bond prices and defaults that create correlated defaults and which have a major impact on the equity tranches and junior tranches of CDOs
6. Simulate the default/no default of each reference name at each of the M time steps and N scenarios
7. Calculate the payments made on the derivative structure given the number and amount of defaults to that date in each scenario using the "waterfall" described in the indenture to that structured product.
8. Apply the Jarrow model as we have done in Chapters 16 and 17 to get the zero coupon bond prices for the risk free curve and for each reference name in the CDO collateral pool or other portfolio. This calculation is very relevant to tranche value. It was fairly common for naïve investors to argue a CDO tranche was still worth par "because no principal or interest payment had been missed" even through the underlying reference collateral was worth far less than par.
9. Calculate the derivative's value at each of the M time steps and N scenarios, conditional the number of defaults among the reference names that have occurred to that point in the simulation.

This is slightly more complicated but not that much more complicated than the credit risky bond simulation discussed in Chapter 19 and the risk-free yield curve simulation discussed in Chapters 6 through 10. It relies on the identical interest rate scenarios and simulation process. Note that if the simulation is done for many companies with common risk factors driving default, we will see correlated default behavior that we have discussed throughout this book.

Collateralized debt obligations: a worked example of reduced form simulation

In this section, we review how to simulate the correlated default risk of the reference names underlying a credit portfolio or a collateralized debt obligation. The process is identical for retail borrowers, small businesses, public firms and sovereign credits. Only the driving risk factors and coefficients differ. For simplicity, we focus on two hypothetical corporations, ABC Company and XYZ Company. We assume we have completed an extensive historical analysis that proves ABC Company's 1 year default probability is sensitive to the 2 year change in home prices and the real growth rate in gross domestic product, as shown in Figure 1:

Figure 1

ABC Company 1 Year Default Probability and Response to Macro Factors					
2 Year Home Price Return	Real Growth Rate of Gross Domestic Product				
	-4.00%	-2.00%	0.00%	2.00%	4.00%
40.00%	2.46%	2.32%	2.19%	2.06%	1.95%
20.00%	3.63%	3.42%	3.23%	3.05%	2.87%
10.00%	4.39%	4.15%	3.92%	3.70%	3.49%
0.00%	5.32%	5.02%	4.74%	4.48%	4.23%
-10.00%	6.42%	6.07%	5.73%	5.42%	5.12%
-20.00%	7.73%	7.31%	6.91%	6.54%	6.18%
-40.00%	11.11%	10.53%	9.98%	9.45%	8.95%

Similarly, XYZ Company is sensitive to a separate set of macroeconomic factors. A combination of historical analysis and business judgment establishes that its one year default probability has the following sensitivity to the 2 year change in home prices and the 10 year yield on government bonds:

Figure 2

XYZ Company 1 Year Default Probability and Response to Macro Factors					
2 Year Home Price Return	10 Year Government Yields				
	1.00%	3.00%	5.00%	7.00%	9.00%
40.00%	1.03%	1.09%	1.15%	1.22%	1.30%
20.00%	2.25%	2.39%	2.53%	2.69%	2.85%
10.00%	3.32%	3.52%	3.73%	3.95%	4.19%
0.00%	4.88%	5.17%	5.47%	5.79%	6.12%
-10.00%	7.11%	7.52%	7.94%	8.39%	8.87%
-20.00%	10.25%	10.81%	11.41%	12.03%	12.68%
-40.00%	20.26%	21.25%	22.27%	23.33%	24.42%

In a reduced form modeling context, theoretical default probabilities are usually modeled in discrete blocks of time, and the default rate in each period is normally given as a logistic function of macroeconomic factors and other risk factors which move randomly. We outlined this procedure in the prior section and in Chapter 16. For simulation purposes, default probabilities are normally generated for discrete daily, weekly, monthly, quarterly or annual time periods. For simulation purposes, as we discussed in Chapter 16, the logistic formula is an attractive choice for modeling because

- It is the maximum likelihood estimator for 0/1 problems like default/no default
- It will never produce simulated default probabilities outside of the range from 0 to 100%

The logistic function takes the form

$$P[t] = \frac{1}{1 + e^{-\alpha - \sum_{i=1}^{n} B_i X_i}}$$

where P[t] is the un-annualized default probability over some discrete interval. The variables X_i are the explanatory variables and the alphas and the betas are the "best fitting" coefficients that produce the maximum likelihood estimates of default probabilities.

For ABC Company, the coefficients which produce the table of default probabilities above are consistent with this logistic function for one year time intervals:

$$P_{ABC}[t] = \frac{1}{1 + e^{-(-3)-(-3X_1 - 2X_2)}}$$

X_1 is the annual growth rate in real gross domestic product and X_2 is the 2 year change in home prices, both expressed as a decimal. These coefficients can be derived either from a historical data base like the sample we showed for Lehman Brothers in Chapter 16 or they can be derived from the numbers in Figure 1 using the 'solver' function in common spreadsheet software.

For XYZ Company, the coefficients which produce the default probabilities above as a function of 10 year government yields and the two year change in home prices are given here:

$$P_{XYZ}[t] = \frac{1}{1 + e^{-(-3)-(3X_2 - 4X_3)}}$$

X_2, as before, is the 2 year change in home prices and X_3 is the 10 year government yield, both expressed as a decimal.

We can simulate these default probabilities forward for M time steps in N scenarios by generating N scenarios for the three macro factors that drive defaults for these two firms:

- The growth in real gross domestic product
- The 2 year change in home prices
- The level of 10 year government yields

ABC Company and XYZ Company have default probabilities which we know are correlated because they have a common dependence on the 2 year change in home prices as given by the tables above. Assuming the current real GDP growth is 0% and that 10 year government yields are 5.00%, a stress test of default probabilities for both firms with respect to the 2 year change in home prices is as follows:

Figure 3

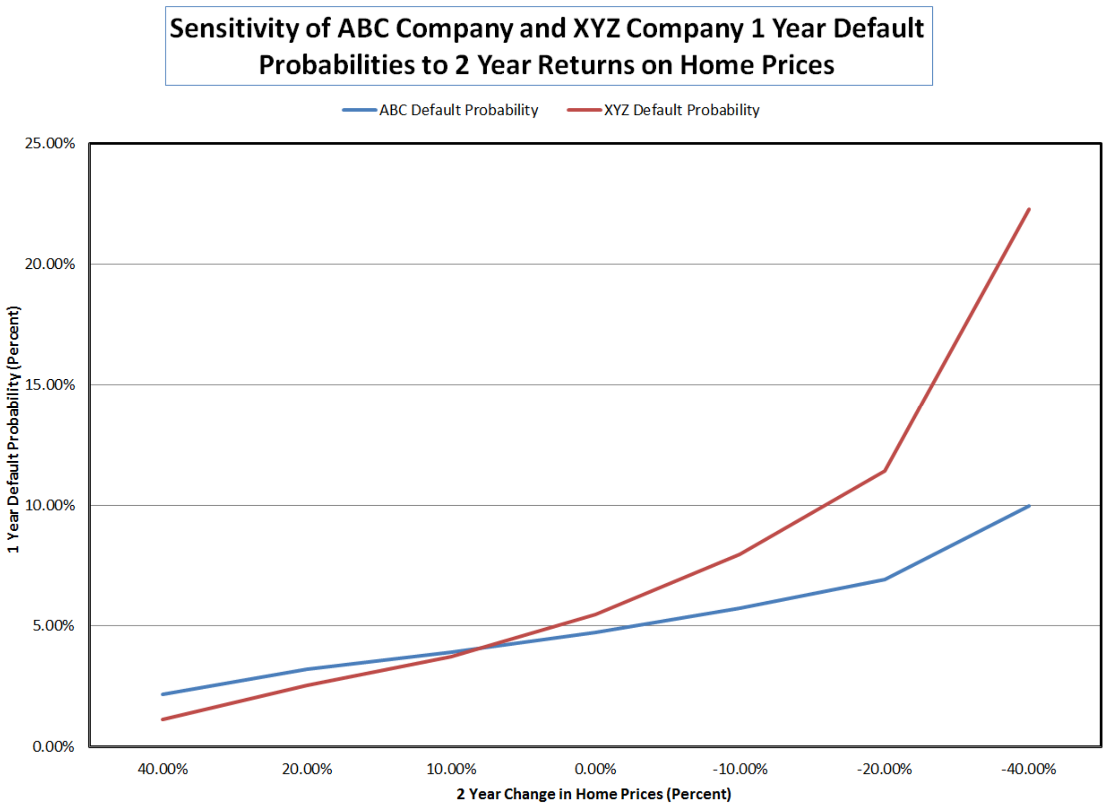

ABC Company and XYZ Company may also have other sources of implicit correlation if the growth in real domestic product (which directly affects only ABC Company) and the 10 year government yield (which directly affects only XYZ Company) are correlated.

Simulating these macro factors forward in an accurate way requires a careful analyst to specify

- The probability distribution for each factor's random movements. MODEL RISK ALERT: Often these distributions are not normally distributed, although that assumption is common.
- Whether or not the distribution of the factor in period J is in fact independent of its value in period J-1. MODEL RISK ALERT: Often, they are NOT independent, although that independence assumption is common.
- The correlations between the random factors themselves

In our N period simulations over M periods, we can value a portfolio of securities for ABC Company and XYZ Company if we know how spreads move, conditional on the level of default probabilities for each firm and the three macro factors above. We discussed this issue in Chapter 17.

We can value the total collateral in a credit portfolio or collateralized debt obligation by adding the value of the securities for ABC, XYZ, and all other reference names included in the collateral pool. With the default/no default data and valuation data at each of the M time periods in the N scenarios, we can accurately allocate cash flow and value to each tranche of a CDO and calculate its value. In an efficient and competitive market (which the CDO "market" is not), the sum of the value of all tranches (less the structuring fee of the sponsor of the CDO) should each the sum of the value of the reference collateral. Felix Salmon's "F9 Model Monkeys" sought to approximate the correct answer in spreadsheet software using the copula approach that we discuss below. There is no need to settle for that kind of extremely crude approximation when this straightforward procedure gives an exact answer conditional on the macro factors and parameters chosen.

This simple example shows that using reduced form default models as part of a CDO analysis or as part of a comprehensive enterprise wide risk management simulation is powerful and straightforward. It allows the analyst to produce government-mandated stress tests with respect to each macro factor that is important in driving defaults. The macro factor links are explicitly derived from either historical data or current market prices and their history. Any simulation approach that uses ratings, which we discussed in Chapter 18, as part of the simulation will not have the same degree of accuracy or transparency with respect to the macro factors which drive risk. The reason is that the rating agencies have been unable to articulate even the time horizon over which ratings apply, let alone the exact formulas by which macro factors affect ratings. With the reduced form approach, the linkages are clear and based on best practice econometric methodology.

Collateralized debt obligations: practice

We now turn from best practice risk analytics to the horrifying reality of what took place in the CDO market during the credit crisis of 2007 to 2010. The spectacular rise and fall of the collateralized debt obligation market is summarized in the press release "Global CDO Issuance" of the Securities Industry and Financial Markets Association (retrieved 2010-011-21 and cited on www.wikipedia.com).

Global CDO Issuance
Billions of U.S. Dollars

2004	157.4
2005	251.3
2006	520.6
2007	481.6
2008	61.9
2009	4.3
2010	8.0

The 2010 issuance total is for a partial year. It is obvious from this issuance history that there has been a complete rejection of the CDO instrument by the investing public, which experienced tens of billions of dollars of losses that were partially chronicled in the Introduction to this book. The reasons for the market's demise have been richly described by journalists and government studies. Perhaps the most comprehensive U.S. government review of the abuses that led to the market's demise is the Levin report mentioned earlier in this book (Committee on Homeland Security and Governmental Affairs, Carl Levin, Chairman; Tom Coburn, Ranking Minority Member: "Wall Street and the Financial Crisis: Anatomy of a Financial Collapse," Majority and Minority Staff Report, Permanent Subcommittee on Investigations, United States Senate, April 13, 2011).

The authors strongly recommend that this report be reviewed in full, at least three times, by any investor willing to consider the purchase of any security that looks or smells like a CDO. At more than 600 pages in length, the Levin report is the definitive narrative of the greed, corruption and misbehavior by “F9 Model Monkeys” that caused the biggest financial collapse of a generation.

The copula method of CDO valuation: a post mortem

We close this chapter with a eulogy for the copula approach to CDO simulation. The copula approach to valuation of collateralized debt obligations has been blamed for much of the credit crisis. Recent articles (listed above) in Wired ("Recipe for Disaster: The Formula that Killed Wall Street," by Felix Salmon, February 23, 2009), Mother Jones ("The Gaussian Copula," by Kevin Drum, February 24, 2009), and the Financial Times (Sam Jones, "Of Couples and Copulas," April 24, 2009) are three very entertaining examples of this genre. Rather than placing blame, other commentators emphasize the positive, saying "There's plenty of blame to go around, let's not waste time assigning blame, let's move forward." This section takes a different view than either group. We argue "Formulas don't cause losses, people cause losses." There are lots of lessons to be learned from the copula formula's uses and abuses.

This list of lessons isn't definitive, but it highlights many of the reasons why and how financial engineers and management went astray in CDO valuation.

Lesson 1: Senior Management Knew or Should Have Known The Copula Approach Was Flawed, but They Did Nothing

Robert Rubin famously commented that he didn't have "enough experience" on Wall Street to understand the liquidity puts of CDO tranches that were at the heart of Citigroup's CDO related losses. But we are sure Mr. Rubin reads the Wall Street Journal, which on August 12, 2005 ran an excellent story on page 1 by Mark Whitehouse, "Slices of Risk: How a Formula Ignited Market That Burned Some Big Investors." The story described David Li's role in bringing the use of copulas into finance, and it then went on at length on how the technique had failed in the May 2005 period, causing losses of hundreds of millions of dollars. Did the heads of the audit committees of the Boards of Directors of Citigroup, Merrill Lynch, and UBS call in the CEO and chief risk officers and grill them on the use of the copula method in their rapidly growing CDO businesses? Did Charles Prince or Stanley O'Neal read that article and put on the brakes? No, they did nothing. People cause losses, not formulas, even when one of the world's most important business newspapers warned us all in a page 1 story.

Lesson 2: The Copula Approach Assumed Away All But One Risk Factor, but Analysts Used It Anyway

The August 12, 2005 story in the Wall Street Journal was very specific about why the copula approach caused such losses. It implicitly assumes that there is only one common risk factor driving the returns on company assets of the reference names underlying a CDO tranche. Other than this common risk factor, all other risk is assumed to be idiosyncratic and uncorrelated with the risks of other companies. We know this is a gross oversimplification; in fact, a March 23, 2009 press release on www.kamakuraco.com recounts how 40 macro risk factors are used for modeling correlated default on the Kamakura Risk Information Services web-based credit portfolio management tool, KRIS-cpm. For CDO market participants, however, the assumption of one common risk factor was hugely attractive. It made the CDO analysis "Excel friendly" and it made the math tractable to the smartest 2% of Wall Street, Felix Salmon’s “F9 Model Monkeys.” The copula approach allows the kind of mathematical exposition that is attractive to academic journals and that one can find in the very interesting compilation of articles in the book *The Definitive Guide to CDOs* (edited by Gunter Meissner, Risk Publications, 2008). The cost to this is that the results from this highly simplified abstraction to reality are just not accurate. The Wall Street Journal put it plainly. The implication of one risk factor is that you should be able to go long one tranche of a CDO and hedge it with the proper short position in other tranche of the same CDO. This trade caused enormous losses because the implication of the copula model was plain wrong.

There was another anecdote in the story that seriously called into question the Merton model of risky debt (see Chapter 18), which is closely linked with the copula approach. The 1974 classic by Robert Merton assumes that only (MODEL RISK ALERT) one risk factor, the value of the assets of a company, drives the prices of the firm's debt and equity. If this is true, one should be able to buy the debt of a company and hedge it with a short position in the firm's common stock. Unfortunately, as the Journal recounted, this implication is simply not true and traders in GM and Ford lost tons of money in May 2005 using the Merton-esque trading strategy. Multiple factors drive bond and common stock prices, but people didn't want to hear it. In fact, Bob Jarrow and one of the authors were very proud to receive hate mail from a prominent junk bond analyst when we made exactly this point in a 1998 publication ("Integrating Interest Rate Risk and Credit Risk in Asset and Liability Management," Asset and Liability Management: A Synthesis of New Methodologies, Risk Publications, 1998). For a detailed study of the lack of accuracy of the highly stylized Merton model of risk debt, see the key reference we highlighted in Chapter 18: "In Search of Distress Risk" in the December 2008 Journal of Finance by John Y. Campbell (Harvard University), Jens Hilscher (Brandeis University and Kamakura Corporation), and Jan Szilagyi (Duquesne Capital).

Lesson 3: Lots of Correlations Matter in CDO Valuation, but Copula Users Assumed Only One Correlation Matters

If there are N reference names in a CDO collateral pool, there are N(N-1)/2 pairs of companies in the pool. The correlations between the default probabilities of each pair of companies are of course different. The same is true for the correlation in the returns on the assets of each pair of companies. The two concepts of correlation are mathematically linked (see Jarrow and van Deventer, RISK Magazine, 2005 for the formula). Copula analysts assume (MODEL RISK ALERT) that the pair-wise correlations in asset returns are the same for every one of the N(N-1)/2 pairs of companies, but this is simply not true. The fact that it is not true was well known to traders--the proof was simple. The correlation implied by observable market prices was not the same across tranches of a CDO on the same reference collateral. The assumption was simply wrong, like the assumption in the Black-Scholes model that implied volatility is constant and should be the same for stock options with any maturity and strike price.

Why did traders, analysts and risk managers use such a simple assumption that they knew to be wrong? For many, the reason was because it was the easy thing to do. For the very smartest people on Wall Street, who knew how flawed the copula approach was, they pushed the technique as an easy way to induce investors to buy CDO tranches at the wrong price. The smart ones used a much more realistic valuation technique and they have the profitability to prove it!

Lesson 4: Default Probabilities Vary Randomly, But Copula Users Assumed Away This Randomness

There are two closely related assumptions about the default probabilities used in a copula analysis. The most common assumption is that they are (MODEL RISK ALERT) constant. A rarer but only slightly better assumption is to assume that (MODEL RISK ALERT) they drift over time (non-randomly) in a way that matches an observable term structure of default probabilities, like those in the CDS market. Unfortunately, these assumptions are also dramatically untrue. As the 2007 to 2010 credit crisis showed, default probabilities rise and fall randomly over the business cycle, driven by common macro factors like home prices and interest rates. The upshot of assuming away this randomness is dramatic errors in valuation, particularly from underestimating the "fat tails" of losses when the economy turns bad. This was especially devastating for "super senior" tranches of CDOs.

Lesson 5: Default Can Happen at Any Time, but Copula Users Assumed That Away

A surprisingly large number of copula users employed a (MODEL RISK ALERT) single period simulation, including many third party vendors as explained by Michel Araten of JP Morgan Chase in a presentation at the ICBI Risk Conference in December in Geneva a few years ago. In a single period simulation, defaults can happen only at time zero or at the end of the single period, not in between. For a CDO of bank trust preferred securities, to name one example, maturities can be more than 30 years, so this assumption about timing of defaults is extreme. The right approach is obvious--to simulate default/no default on a multi-period basis, typically monthly. CDO market participants, however, generally didn't do this. This explains the widely held but completely false belief that an increase in the single correlation in a copula model will cause the price of the equity tranche to rise (see "Synthetic CDO Equity: Short or Long Correlation" by Robert Jarrow and Donald R. van Deventer, *Journal of Fixed Income*, Spring 2008). That conclusion about equity tranches is due to the assumption of a single modeling period. A multi-period simulation was essential for accuracy, but most analysts and some vendors clung to a single period model.

Lesson 6: CDO Analysts Can Be Slaves to Fashion, but At High Costs

Gunter Meissner's fascinating compilation of articles in *The Definitive Guide to CDOs* (Risk Publications, 2008) is "definitive" proof of the "trendiness" of the copula approach. Even though most of the chapters in the book were penned as the current credit crisis began to unfold, the overwhelming majority of the chapters endorse or extend the copula method. A reviewer on amazon.com cited only two chapters in the book (Chapter 4 on hedging, and Chapter 16 by Robert Jarrow and Donald R. van Deventer entitled "CDO Valuation: Fact and Fiction") that were critical of the copula approach. Even as events were revealing (again) the flaws of the copula approach that the Wall Street Journal highlighted in 2005, many were still singing its praises. This is certainly not the fault of David Li, who first used the copula approach in finance. It's the fault of all of us who came after David but failed to maintain our scientific objectivity about what works and what doesn't, and why.

A multi-period simulation of CDO values with default probabilities driven up and down by macro-economic factors is an exercise in computer science that's hard to do without sophisticated third party risk management software. It's not fashionable, because there's no closed form solution, so you won't see many academic publications on its implications. Nonetheless, it's the only method that works. That's perhaps the most important thing we have learned from the attempt to use the copula approach for credit portfolio management.

Valuing the Jarrow –Merton Put Option

We have analyzed the reference names in a credit portfolio or in a collateralized debt obligation just as if they were loans in a bank lending portfolio or private placements on the books of a pension fund or insurance company. The next step is to integrate them with the liabilities of the institution so that we can evaluate the Jarrow-Merton put option as an integrated measure of firm risk. We now turn to that task in the remainder of this book.