Understanding the Risk of Synthetic CDOs

Michael S. Gibson[*]

Revised, July 2004

Understanding the Risk of Synthetic CDOs

Abstract: Synthetic collateralized debt obligations, or synthetic CDOs, are popular vehicles for trading the credit risk of a portfolio of assets. Following a brief summary of the development of the synthetic CDO market, I draw on recent innovations in modeling to present a pricing model for CDO tranches that does not require Monte Carlo simulation. I use the model to analyze the risk characteristics of the tranches of synthetic CDOs. The analysis shows that although the more junior CDO tranches – equity and mezzanine tranches – typically contain a small fraction of the notional amount of the CDO's reference portfolio, they bear a majority of the credit risk. One implication is that credit risk disclosures relying on notional amounts are especially inadequate for firms that invest in CDOs. I show how the equity and mezzanine tranches can be viewed as leveraged exposures to the underlying credit risk of the CDO's reference portfolio. Even though mezzanine tranches are typically rated investment-grade, the leverage they possess implies their risk (and expected return) can be many times that of an investment-grade corporate bond.

The paper goes on to show how CDO tranches and other innovative credit products, such as single-tranche CDOs and first-to-default basket swaps, are sensitive to the correlation of defaults among the credits in the reference portfolio. Differences of opinion among market participants as to the correct default correlation can create trading opportunities. Finally, the paper shows how the dependence of CDO tranches on default correlation can also be characterized and measured as an exposure to the business cycle, or as "business cycle risk." A mezzanine tranche, in particular, is highly sensitive to business cycle risk.

1 Introduction

A synthetic collateralized debt obligation, or synthetic CDO, is a transaction that transfers the credit risk on a reference portfolio of assets. The reference portfolio in a synthetic CDO is made up of credit default swaps. Thus, a synthetic CDO is classified as a credit derivative. Much of the risk transfer that occurs in the credit derivatives market is in the form of synthetic CDOs. Understanding the risk characteristics of synthetic CDOs is important for understanding the nature and magnitude of credit risk transfer.

In this paper, I present a simple model that can be used to understand the pricing and risk characteristics of synthetic CDOs. The simple model in this paper can be solved analytically, unlike models typically used for cash CDOs that require Monte Carlo simulation. Because the model's computations are quick, I can do a more extensive analysis of synthetic CDO tranches along various dimensions of risk. However, the the model in this paper does require simplifying assumptions that limit how broadly the model can be applied.

2 Some background on synthetic CDOs

Synthetic CDOs developed as an outgrowth of cash CDOs. Cash CDOs have a reference portfolio made up of cash assets such as corporate bonds or loans. Cash CDOs developed when traditional securitization techniques used to create mortgage-backed or asset-backed securities were applied to transfer credit risk. Unlike cash CDOs, where pricing and risk management draw on traditional securitization techniques, the pricing and risk management of synthetic CDOs as discussed in this paper draw on derivatives techniques, similar to those used to price and risk-manage interest rate swaps and credit default swaps.

A defining feature of both cash and synthetic CDOs is the tranching of credit risk. The risk of loss on the reference portfolio is divided into tranches of increasing seniority. Losses will first affect the "equity" or "first loss" tranche, next the "mezzanine" tranche(s), and finally the "senior" and "super-senior" tranches.[1] CDO "investors" take on exposure to a particular tranche, effectively selling credit protection to the CDO "issuer." The CDO "issuer," in turn, hedges its risk by selling credit protection on the reference portfolio in the form of single-name credit default swaps. Parties on the other side of these hedging transactions are the ultimate "sellers" of credit risk to the CDO "investor," with the CDO "issuer" acting as intermediary.

Within the class of corporate credit risk, synthetic CDOs tend to focus on either the investment-grade sector or the high-yield sector. Recently, the range of collateral used in synthetic CDO transactions has expanded beyond corporate credit risk to include residential mortgage-backed securities, commercial mortgage-backed securities, CDO tranches, and other structured finance securities.

---

[1]A "super-senior" tranche is often defined as one that is senior to a AAA-rated tranche.

FIGURE 1. SYNTHETIC CDO ISSUANCE, BY TYPE OF COLLATERAL
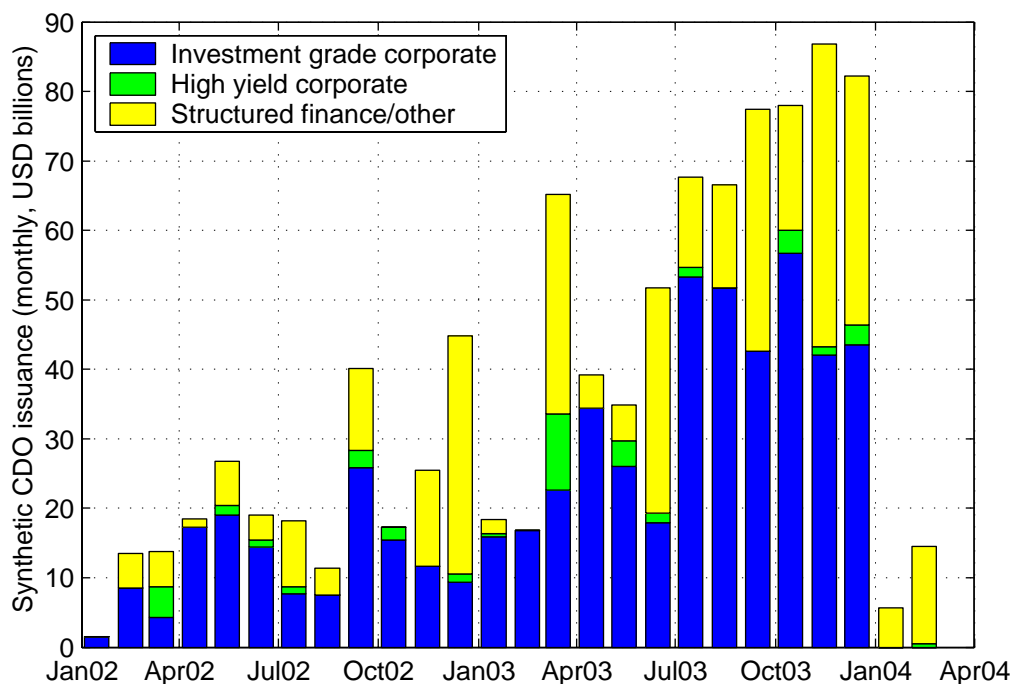
Note: "Structured finance/other" includes asset-backed securities (including CDO tranches), residential mortgage-backed securities, commercial mortgage-backed securities, and emerging market debt. Deals not denominated in US dollars were converted at March 2004 market exchange rates. Source: www.creditflux.com.

To give some background on the development of the synthetic CDO market, Figure 1 shows the notional amount of synthetic CDO tranches issued over January 2002 to February 2004, broken down by collateral type. Over this time period, total issuance volume has grown rapidly and has averaged $37 billion per month. Investment grade corporate debt is the most common collateral type. Structured finance is the second-largest collateral type and has been growing in importance in recent months, as declining corporate credit spreads have reduced the attractiveness of corporate credit risk to investors. The relatively small amount of high-yield synthetic CDOs reflects the lack of liquidity in the high-yield segment of the single-name credit default swap market.

The first synthetic CDO deals were done by banks in 1997.[2] These "bank balance sheet" deals were motivated by either a desire to hedge credit risk, a desire to reduce regulatory capital, or both. Following these early deals, the same synthetic CDO technology has been used to create CDO tranches with risk-return profiles that investors find attractive. These

---

[2]According to Goodman (2002), the first synthetic CDOs in 1997 were conducted by Swiss Bank Corporation ("Glacier Finance Ltd.") and JP Morgan ("BISTRO").

later deals, driven by the needs of credit investors rather than banks, are termed "arbitrage" deals.[3] In the early days of the synthetic CDO market (1997-1999), nearly all synthetic CDOs were motivated by balance sheet considerations. However, arbitrage deals have been growing rapidly and now account for a significant majority of market activity.

Also in the early days of the synthetic CDO market, deals included a complete set of tranches (equity, mezzanine, senior) whose notional amounts summed to the notional amount of the reference portfolio. Now, most synthetic CDO activity consists of so-called single-tranche CDOs, where only one tranche of the CDO's capital structure is sold. Single-tranche CDOs are discussed in detail in section 6.1 below.

Synthetic CDO tranches can also be either "funded" or "unfunded." If a tranche is funded, the CDO investor pays the notional amount of the tranche at the beginning of the deal and any defaults cause a writedown of principal. Throughout the deal, the investor receives LIBOR plus a spread that reflects the riskiness of the tranche. The investor's funds are put into a collateral account and invested in low-risk securities (government or AAA-rated debt). "Unfunded" tranches are similar to swaps. No money changes hands at the beginning of the deal. The investor receives a spread and pays when defaults in the reference portfolio affect the investor's tranche (after any subordinate tranches have been eaten away by previous defaults). Because unfunded tranches rely on the investor's future ability and willingness to pay into CDO, they create counterparty credit risk that must be managed.

How do CDOs make economic sense when they merely rearrange the payment priority of other credit-risk-sensitive instruments such as credit default swaps, bonds and loans? In a recent paper, Duffie and Garleanu (1999) give two answers. First, regulatory capital requirements on credit risk can at times exceed the economic capital the market requires to bear the risk. Banks cite reducing regulatory capital as a motive for their participation in credit derivatives markets, more so in the early days of the market than today.

Second, CDOs help investors overcome market imperfections associated with the illiquidity of bonds and loans. Few corporate bonds trade more than twice a day.[4] Loans trade even less frequently. Illiquidity makes it costly for credit investors to assemble a portfolio that meets their diversification and risk-return targets. Efficiently managing portfolio credit risk is a dominant motive for commercial banks' participation in credit derivative markets.

The economic value of a CDO is apparent in the fact that the CDO's spread income from the reference portfolio can compensate investors in the CDO tranches and also cover transactions costs. The rapid adoption of CDO technology by credit investors suggests that the cost of creating a CDO is less than the cost a credit investor would incur to assemble a portfolio of bonds and/or loans to meet the investor's diversification and risk-return targets. The high

---

[3]The terminology has taken hold despite the absence of a true "arbitrage," which can be loosely defined as a risk-free investment with a positive excess return.

[4]Among the roughly 5,000 corporate bonds for which secondary market transaction data are disseminated on the NASD's trade reporting system (TRACE), only about 250, or 5 percent, traded fifteen times or more in the seven days ending January 29, 2004.

cost of investing directly in a portfolio of bonds or loans is presumably driven by the high bid-ask spreads an investor would pay, reflecting the illiquidity of bond and loan markets.

3 The distribution of default losses on a portfolio

To price or measure the risk of a synthetic CDO tranche, the probability distribution of default losses on the reference portfolio is a key input. In this section, I show how to compute the loss distribution in an analytically tractable way. To do this, I assume that the correlation of defaults within the reference portfolio is driven by a common factor. I compute the loss distribution conditional on the factor, then integrate out the factor to get the unconditional loss distribution. The approach described here synthesizes a number of recent papers that tackle this problem in a similar way.[5]

Each credit in the reference portfolio of $i = 1, \ldots, N$ credits is described by the following parameters:

A_i notional amount of credit i
$q_i(t)$ risk-neutral probability that credit i defaults before t
R_i recovery rate of credit i

Notional amounts are known quantities, risk-neutral default probabilities can be estimated using each credit's single-name credit default swap (CDS) spreads, and recovery rates are assumed known and constant. Note that all analysis in this paper uses risk-neutral default probabilities and recovery rates derived from market spreads.[6]

Creditworthiness is assumed to depend on a reference credit's normalized asset value x_i. Default occurs when x_i falls below a threshold $\bar{x}_i$. The asset value x_i is assumed to depend on a single common factor M:

$$x_i = a_i M + \sqrt{1 - a_i^2} Z_i \tag{1}$$

where x_i, M, and Z_i are mean-zero, unit-variance random variables with distribution functions F_i, G, and H_i. The random variables $M, Z_1, \ldots, Z_N$ are assumed to be independently distributed. The factor loading a_i is constrained to lie between zero and one.

Three implications of this setup are immediate and will be useful later on:

- The correlation of asset value between credits i and j is equal to $a_i a_j$.

- The default threshold $\bar{x}_i$ is equal to $F_i^{-1}(q_i(t))$.

---

[5]See Li (2000), Boscher and Ward (2002), Andersen, Sidenius and Basu (2003), and Hull and White (2003). Generalization to more than one common factor is straightforward but comes at a cost of greater computational time; see Andersen, Sidenius and Basu (2003) and Hull and White (2003).

[6]Using historical, rather than risk-neutral, default probabilities or rating transition matrices to assess the risk of CDO tranches is a useful approach but not the one pursued in this paper.

- The conditional default probability $q_i(t|M)$, which in this setup can also be written as $\text{Prob}\{x_i < \bar{x}_i | M\}$, is equal to

$$q_i(t|M) = H_i\left(\frac{\bar{x}_i - a_i M}{\sqrt{1 - a_i^2}}\right) \tag{2}$$

using (1).[7]

The distribution of losses on the reference portfolio takes on a discrete number of values. Each credit either takes no loss or a loss of $A_i(1 - R_i)$. If A_i and R_i are equal across credits, the number of discrete values will be equal to $N + 1$, corresponding to no defaults, one default, two defaults, and so on up to N defaults. From this point on, I will assume that A_i and R_i are equal across credits and will drop the i subscript. With this assumption, the distribution of losses and the distribution of the number of defaults become interchangeable; simply multiply the number of defaults by $A(1 - R)$ to get losses.

If A_i and R_i are not equal across credits, the method I describe can still be used. However, the computational burden increases with the number of discrete values the portfolio's losses are distributed over. If the number of discrete values grows too large, it may be necessary to coarsen the loss distribution as described by Andersen, Sidenius and Basu (2003) to keep the number of discrete values down to a manageable size.

The distribution of the number of defaults conditional on the common factor M can be computed with the following recursion.[8] Let $p^K(l, t|M)$ denote the probability that exactly l defaults occur by t, conditional on the common factor M, in a reference portfolio of size K. Assume we know the default distribution for a set of K credits:

$$p^K(l, t|M) \qquad l = 0, \ldots, K \tag{3}$$

Add one credit, with a conditional default probability $q_{K+1}(t|M)$. The default distribution for the new reference portfolio of $K + 1$ credits is:

$$p^{K+1}(0, t|M) = p^K(0, t|M)\left(1 - q_{K+1}(t|M)\right) \tag{4}$$
$$p^{K+1}(l, t|M) = p^K(l, t|M)\left(1 - q_{K+1}(t|M)\right) + p^K(l - 1, t|M)q_{K+1}(t|M) \quad l = 1, \ldots, K$$
$$p^{K+1}(K + 1, t|M) = p^K(K, t|M)\, q_{K+1}(t|M)$$

Starting with the degenerate default distribution for $K = 0$, $p^0(0, t|M) = 1$, we can use the recursion (4) to solve for the default distribution for the reference portfolio of N credits:

$$p^N(l, t|M) \qquad l = 0, \ldots, N \tag{5}$$

---

[7]Vasicek (1987) shows that this model, with the distributions F_i, G, and H_i assumed to be normal, is isomorphic with a Merton-style model where each firm's asset value follows a geometric Brownian motion, firms i and j's asset value processes have correlation $\rho_{ij} = a_i a_j$, and a firm defaults if its asset value at time t falls below a threshold.

[8]The recursion is discussed in Andersen, Sidenius and Basu (2003).

Once we have the conditional default distribution, the unconditional default distribution $p(l, t)$ can be solved as

$$p(l, t) = \int_{-\infty}^{\infty} p^N(l, t|M) \, g(M) \, dM \tag{6}$$

where g is the probability density of M. The integral is calculated with numerical integration. Because the functions $p(l, t|M)$ and $g(M)$ are typically smooth and well-behaved, numerical integration is straightforward and quite fast. For each t of interest, $N + 1$ numerical integrations are required, each requiring many evaluations of the recursion (4) for different values of M. On an AMD Athlon 2.2 GHz PC with 1GB of RAM, computing the unconditional default distribution for $N = 100$ takes 10 seconds using the Scilab software package.[9] The calculation of the default distribution must be repeated for each date t on which a payment may occur (eg, 20 payment dates for a 5-year CDO with quarterly payments).

At this point, it may be useful to give a numerical example of what a default probability distribution looks like. For the numerical example, let $N = 100$, the factor loadings a_i equal $\sqrt{0.3}$, the risk-neutral default probabilities equal $1 - e^{-0.01t}$ (default hazard of 1 percent per year), $t = 1$ year, and the common and idiosyncratic factors (M and Z_i) be normally distributed.

Under these assumptions, Table 1 shows both conditional and unconditional default probability distributions. As the common factor M increases, moving to the right across the table, the weight of the distribution shifts towards fewer defaults. The unconditional distribution, which is computed by integrating out the common factor as in (6), can be interpreted as a weighted average of the conditional distributions, weighted by the relative probability of the different values of the common factor.

4   Pricing a synthetic CDO tranche

A synthetic CDO tranche on a given reference portfolio is defined by an interval of losses, $[L, H]$, that the CDO tranche "investor" is responsible for. For example, on a reference portfolio of \$1 billion, a mezzanine tranche might be responsible for losses from \$30 million to \$100 million. For this tranche, $L = \$30$ million and $H = \$100$ million. The endpoints of the interval are also referred to as "attachment points." The lower attachment point can also be called the tranche's "credit enhancement," since it measures the size of any other, more subordinate tranches that buffer the tranche against loss. In the above example, the mezzanine tranche has attachment points of 3 and 10 percent of the size of the reference portfolio and its credit enhancement is 3 percent.

A synthetic CDO tranche can be valued as a swap contract between two parties.[10] The "investor" receives periodic spread payments from the "issuer" (the "fee" leg) and makes

---

[9]http://www.scilab.org

[10]This is true whether the actual contract is an unfunded swap or a funded note.

TABLE 1. NUMERICAL EXAMPLE OF DEFAULT PROBABILITY DISTRIBUTION

The table shows probability distributions of the number of defaults in a reference portfolio of 100 entities. The example assumes normally distributed factors, $t = 1$ year, $q_i(t) = 1 - e^{-0.01t}$ and $a_i = \sqrt{0.3}$ for $i = 1, \ldots, 100$.

Number of defaults	Conditional default distribution $p(l, t\|M)$					Unconditional default distribution $p(l, t)$
	$M = -2$	$M = -1$	$M = 0$	$M = 1$	$M = 2$	
0	0.001	0.186	0.763	0.971	0.998	0.644
1	0.005	0.316	0.206	0.029	0.002	0.166
2	0.019	0.265	0.028	*	*	0.072
3	0.048	0.147	0.002	*	*	0.039
4	0.087	0.060	*	*	*	0.023
5	0.127	0.020	*	*	*	0.015
6	0.152	0.005	*	*	*	0.010
7	0.154	0.001	*	*	*	0.007
8	0.136	*	*	*	*	0.005
9	0.105	*	*	*	*	0.004
10	0.072	*	*	*	*	0.003
11	0.045	*	*	*	*	0.002
12	0.025	*	*	*	*	0.002
13	0.013	*	*	*	*	0.001
14	0.006	*	*	*	*	0.001
15	0.003	*	*	*	*	0.001
16	0.001	*	*	*	*	0.001
17	*	*	*	*	*	0.001
18	*	*	*	*	*	*
⋮	⋮	⋮	⋮	⋮	⋮	⋮
100	*	*	*	*	*	*

* = less than 0.001

7

contingent payments to the issuer when defaults affect the tranche (the "contingent" leg). Here, I assume that all payments (fee and contingent legs) occur on periodic payment dates $T_i, i = 1, \ldots, n$. I assume that the investor makes a payment on the next periodic payment date after the default occurs, not on the date of the default itself. The value of the tranche, from the perspective of the investor, is the expected present value of the fee leg less the expected present value of the contingent leg.

The expected default loss on the tranche up to payment date T_i, denoted as EL_i, is an important quantity for computing the expected present values of the fee and contingent legs. Recall from the previous section that $p(l, t)$ is the default probability distribution and the loss from any default is $A(1 - R)$. Expected loss is computed as:

$$EL_i = \sum_{l=0}^{N} p(l, T_i) \, \max(\min(lA(1 - R), H) - L, 0) \tag{7}$$

The expected present value of the contingent leg is the discounted sum of the expected payments the tranche investor must make when defaults affect the tranche:

$$\text{Contingent} = \sum_{i=1}^{n} D_i(EL_i - EL_{i-1}) \tag{8}$$

where D_i is the risk-free discount factor for payment date i.

The expected present value of the fee leg is the discounted sum of the spread payments the tranche investor expects to receive:

$$\text{Fee} = s \sum_{i=1}^{n} D_i \Delta_i \big\{ (H - L) - EL_i \big\} \tag{9}$$

where Δ_i is the accrual factor for payment date i $(\Delta_i \approx T_i - T_{i-1})$ and s is the spread per annum paid to the tranche investor. The term $(H - L) - EL_i$ is the expected tranche principal outstanding on payment date T_i. It reflects the decline in principal as defaults affect the tranche.

The mark-to-market value of the tranche, from the perspective of the tranche investor (fee leg receiver) is

$$MTM = \text{Fee} - \text{Contingent} \tag{10}$$

Like other swaps, a synthetic CDO tranche will have the spread set at inception so the swap's mark-to-market value is zero. Setting $MTM = 0$ in (10) and using (8) and (9) to solve for s, the "par spread," or s_{Par}, is equal to

$$s_{\text{Par}} = \frac{\text{Contingent}}{\sum_{i=1}^{n} D_i \Delta_i \{ (H - L) - EL_i \}} \tag{11}$$

8

5   The risk of CDO tranches

Using the model of CDO tranche pricing outlined above, the risk of CDO tranches can be measured. To illustrate the risk of CDO tranches, I use a stylized example of a hypothetical $1 billion synthetic CDO with three tranches. The hypothetical CDO tranches are described in Table 2. The equity tranche bears the first $30 million of losses, the mezzanine tranche bears the next $70 million, and the senior tranche bears any losses above $100 million. Although determining credit ratings on synthetic CDO tranches is outside the scope of this paper, market observation suggests that the senior tranche of such a CDO would typically be rated AAA, the mezzanine tranche would typically be rated A/BBB, and the equity tranche would typically be unrated. In this example, all three tranches are unfunded.

The $1 billion reference portfolio for the hypothetical CDO consists of 100 reference credits with notional amount $10 million each. All 100 reference credits share the following characteristics:

Single-name CDS spread	60 basis points
Notional amount	$10 million
Recovery rate	40 percent
Default hazard rate	1 percent per year
Asset correlation (a_i^2)	30 percent

The hypothetical CDO has a maturity of five years with quarterly payment dates. The risk-free interest rate is assumed to be a constant 5 percent (continuously compounded). The common risk factor and idiosyncratic risk factors are assumed to be normally distributed.[11] The results presented below of course depend on these assumptions. Later, I will present results that vary the correlation parameter a_i.

The three tranches would typically be sold to different investors. In fact, some of the economic value of a CDO stems from its ability to create credit risk-sensitive assets whose risk-return characteristics are tailored to certain investor classes. The equity tranche could be sold to a professional asset manager managing money on behalf of institutional clients, or to a hedge fund. The mezzanine tranche could be sold to a regional bank looking to diversify its credit exposures. The senior tranche could be sold to a reinsurer or other investor looking for low-risk, low return assets.

Such a synthetic CDO could be issued by a commercial bank choosing the credits for the reference portfolio based on its desire to hedge loans on its balance sheet.[12] Or, the synthetic CDO could be underwritten by an investment bank on behalf of an asset manager who

---

[11]The model can easily accomodate other mean-zero, unit variance distributions for the risk factors, such as Student-t. See Andersen, Sidenius and Basu (2003) and Hull and White (2003) for examples.

[12]Some commercial banks have done such transactions in the past to hedge balance sheet risk. However, for hedging purposes they currently favor single-name credit default swaps over portfolio products such as CDOs due to greater flexibility and more favorable pricing.

TABLE 2. A STYLIZED HYPOTHETICAL CDO

Tranche	Attachment points (percent)	Notional amount ($ millions)	Par spread (basis points)
Equity	0–3	30	1507
Mezzanine	3–10	70	315
Senior	10-100	900	7
Entire portfolio	0–100	1,000	60

chooses the reference portfolio based on fundamental credit analysis and has limited rights to actively manage the reference portfolio.

One fact to note from Table 2 is that the weighted-average spread on the three tranches is 74 basis points, which exceeds the 60 basis point spread on the reference portfolio. If the CDO issuer sells all three tranches at their model-implied par spread and hedges by selling credit protection on the reference portoflio of 100 single-name CDS, the CDO issuer will initially be paying out more than it takes in — its net position is one of negative carry.

Table 3 shows the net carry the CDO issuer expects to pay out on each payment date. Over time, as defaults occur and erode the principal of the high-yielding equity tranche, the negative carry diminishes. It becomes positive about halfway through the CDO's life. By construction, the present discounted value of the carry is zero because all tranches are fairly priced at their par spreads.[13]

One implication of the negative carry is that issuing a complete set of synthetic CDO tranches and selling credit protection on the reference portfolio is not a perfectly hedged position for a CDO issuer. There is residual risk that is related to the timing of defaults. Early defaults benefit the CDO issuer by reducing payments required on the high-spread equity tranche, while late defaults have the opposite effect.

There are two ways the negative carry might in practice be reduced or eliminated. First, actual spreads on CDO tranches would be less than the model-implied par spreads presented here to reflect hedging costs and the profit margin of CDO issuers. Second, equity tranche spread may be deferred to later payment dates through structural means, such as subordination to a reserve account. Note that cash CDOs (and some synthetic CDOs), like traditional securitizations, are structured to pass through cash flows on the reference portfolio to CDO investors according to pre-specified "payment waterfall" rules, thus avoiding negative carry entirely. A typical waterfall might specify that cash flows into the CDO go to pay (in order) deal structuring fees, senior tranche spread, mezzanine tranche spread, equity tranche

---

[13]The negative carry in this situation has been documented by Boscher and Ward (2002) and Schmidt and Ward (2002), among others.

Month	3	6	9	12	15	18	21	24	27	30
Net spread (basis points)	3.1	2.6	2.2	1.8	1.4	1.1	0.7	0.4	0.1	-0.2
Month	33	36	39	42	45	48	51	54	57	60
Net spread (basis points)	-0.5	-0.7	-1.	-1.2	-1.5	-1.7	-1.9	-2.2	-2.4	-2.6

spread, and asset manager fees. Because this paper's model does not include such a water-fall, caution is advised in applying the model in this paper to such deals. However, some synthetic CDO tranches, especially single-tranche CDOs, are structured like swaps and can be accurately valued using the model described in this paper.

5.1 Risk measures for the hypothetical CDO tranches

Clearly the equity, mezzanine, and senior tranches of CDOs have quite different risk characteristics. The risk of CDO tranches can be measured in various ways. In this section, I present three risk measures:

- The sensitivity of the tranche's value to a change in credit spreads on the names in the reference portfolio.

- The tranche's expected loss (EL) from defaults in the reference portfolio occurring up to the maturity of the CDO.

- A level of loss due to default that is one standard deviation above the tranche's expected loss. This is one measure of unexpected loss (UL).[14]

These risk measures are computed as follows, using the model described in sections 3 and 4:

- The sensitivity to credit spreads is computed by reducing all the credit spreads on the 100 reference credits by 10 basis points, recomputing the default probability distribution, then using (7)–(10) to compute the increase in each tranche's mark-to-market value.

- The expected loss up to the CDO's maturity is computed using (7) with T_i set at the CDO's final maturity date.

---

[14]Other measures of UL are also commonly used among market practitioners. For example, a high percentile of the loss distribution, rather than a standard deviation-based measure, is often used for economic capital purposes.

- The standard deviation of loss up to payment date T_i is equal to

$$SD_i = \left(\sum_{l=0}^{N} p(l, T_i) \left[\max(\min(lA(1-R), H) - L, 0) - EL_i \right]^2 \right)^{1/2} \qquad (12)$$

and unexpected loss $UL_i = EL_i + SD_i$.

Before discussing the three risk measures, it is worth noting that risk measures are generally not additive. That is, the risk of a portfolio is generally not equal to the sum of the risk of the underlying elements of the portfolio. This is true for two of three risk measures discussed here (sensitivity and UL). However, because it is a mathematical expectation, expected loss EL is additive.

Table 4 shows the three risk measures for the tranches of the hypothetical CDO. All three risk measures tell the same basic story: The equity tranche, with only 3 percent of the notional exposure, is by far the riskiest. The senior tranche, with 90 percent of the notional exposure, is the least risky. The equity and mezzanine tranches together account for only 10 percent of the CDO's notional amount but 70–90 percent of the risk, depending on the risk measure used. The message from Table 4 is that, for CDO tranches, notional exposure is not a meaningful way to measure risk. Many argue that disclosures focusing on balance sheet exposure amounts have long been inadequate for judging a firm's true credit risk profile. One implication of the disconnect between exposure and risk of CDO tranches is that such disclosures are especially inadequate fore firms that are active in the credit derivatives market.

The risk measures in Table 4 can be used to assess the leverage of CDO tranches. Table 5 uses the three risk measures to compare the leverage of the equity, mezzanine, and senior tranches. Leverage is computed by dividing each tranche's risk, expressed as a percent of notional amount, by the risk of the entire portfolio. When computed this way, a cash investment in bonds or loans has leverage equal to one by construction.

The leverage numbers in Table 5 suggest that the hypothetical equity tranche has 13–17 times the risk of a cash investment in bonds or loans. The mezzanine tranche, although it may have an investment grade credit rating, has 5–7 times the risk of a cash investment. The senior tranche has much less risk than a cash investment.

5.2 Risk and leverage vary with credit spreads and credit enhancement

Credit spreads on the CDO's reference porfolio and a tranche's location in the CDO's capital structure affect the risk and leverage measures presented above. In this section, I allow both to vary to show how a tranche's risk and leverage change when the reference portfolio becomes riskier (credit spreads increase) or the tranche's location in the capital structure shifts.

TABLE 4. RISK MEASURES OF HYPOTHETICAL CDO TRANCHES

(Dollar amounts in millions)

Tranche	Attachment points (percent)	Notional amount	Sensitivity to 10bp shock to credit spreads		Expected loss (EL)		Unexpected loss (UL)	
			Dollar amount	Percent of notional	Dollar amount	Percent of notional	Dollar amount	Percent of notional
Equity	0–3	30	1.9	6	15.3	51	27.6	92
Mezzanine	3–10	70	2.0	3	10.5	15	31.5	45
Senior	10–100	900	1.0	0.1	3.4	0.4	22.1	2.5
Entire portfolio	0–100	1,000	4.2	0.4	29.2	2.9	71.3	7.1

Tranche	Attachment points (percent)	Notional amount ($ millions)	Leverage computed using:		
			Sensitivity	Expected loss (EL)	Unexpected loss (UL)
Equity	0–3	30	15x	17x	13x
Mezzanine	3–10	70	7x	5x	6x
Senior	10–100	900	0.3x	0.1x	0.3x
Entire portfolio	0–100	1,000	1x	1x	1x

I consider three reference portfolios with increasing credit spreads (decreasing credit quality). All 100 credits in the reference portfolio have identical credit spreads and recovery rates are fixed at 40 percent:

Portfolio	Credit spread (basis points)	Default hazard (percent per year)
A	60	1
B	120	2
C	180	3

I consider mezzanine tranches whose credit enhancement (lower attachment point) varies from 2 to 10 percent, with tranche thickness held fixed at 3 percent.

Figure 2 shows how the mezzanine tranche's EL (defined in section 5.1 above) varies with credit spreads and credit enhancement. For all three portfolios, as the mezzanine tranche moves higher in the CDO's capital structure (has more credit enhancement), the tranche's risk declines. As the riskiness of the reference portfolio increases moving from Portfolio A to B to C, the tranche's risk increases.

Figure 2 makes it clear that there is a tradeoff to be made in constructing CDO tranches. Holding the tranche thickness constant at 3 percent, one can achieve an EL in the vicinity of 25 percent of notional in three ways:

- the low-risk reference portfolio A and credit enhancement of around 2 1/2 percent

- the medium-risk reference portfolio B and credit enhancement of around 6 1/2 percent

- the high-risk reference portfolio C and credit enhancement of around 10 percent

This tradeoff is consistent with what is observed in the market. Tranches on low-risk reference portfolios, such as investment grade corporate bonds, typically have less credit enhancement

FIGURE 2. CREDIT ENHANCEMENT AND RISK OF SYNTHETIC CDO TRANCHES
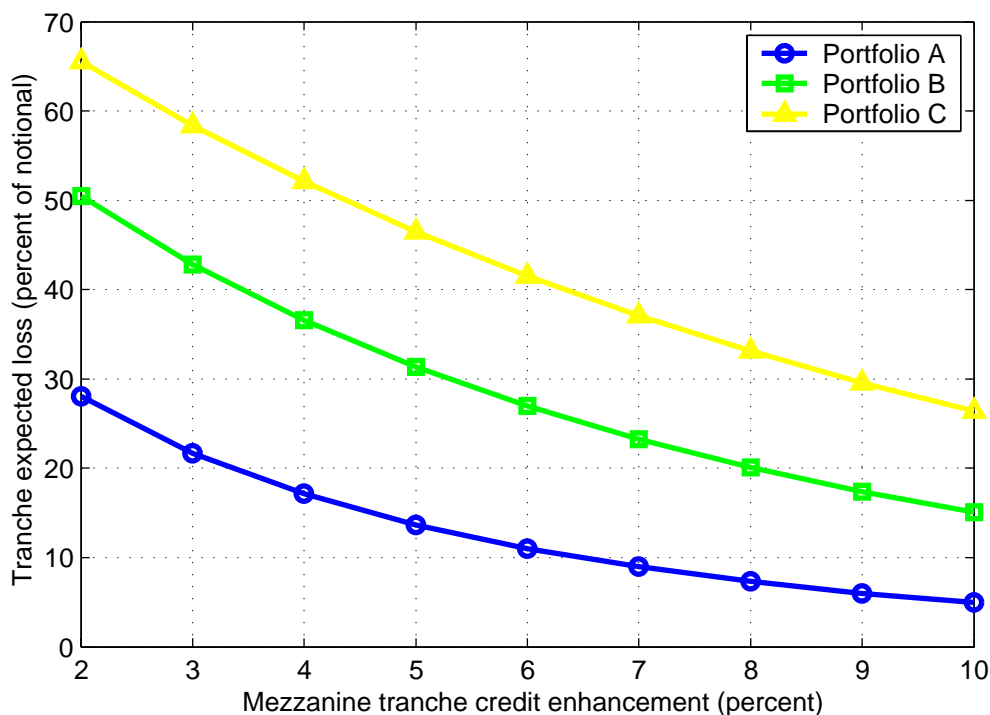
than tranches of the same seniority on high-risk reference portfolios, such as high-yield corporate bonds. Table 6 compares the average credit enhancement levels for tranches of synthetic CDOs issued in 2002–03 with high-yield and investment-grade corporate debt as collateral. The data are consistent with the tradeoff discussed above. Rating agencies have said that expected loss is a key concept for rating structured finance securities, including CDO tranches, and Table 6 supports this.

Figure 3 shows the leverage measure introduced above, where leverage is measured as expected loss (EL) on the tranche relative to the EL of a cash investment in the reference portfolio. The three lines all slope down and they all lie roughly atop one another. This implies that the riskiness of the reference portfolio has little effect on leverage, but credit enhancement has a strong effect.

Putting Figures 2 and 3 together, we can see the interaction of risk and leverage in the construction of CDO tranches. For a low-risk reference portfolio like Portfolio A to yield a mezzanine tranche with EL around 25 percent (high enough to get a low investment-grade credit rating), the credit enhancement must be small, around 2–3 percent. This implies leverage of 8 or 9 times. To construct a mezzanine tranche with similar EL but a high-risk reference portfolio like Portfolio C, the credit enhancement would need to be large, around 10 percent, implying leverage of only around 3 times. In sum, low-risk assets need to be more leveraged to generate a mezzanine tranche level of risk (and return).

TABLE 6. CREDIT ENHANCEMENT OF SYNTHETIC CDO TRANCHES BY COLLATERAL TYPE

	(percent)	
	Reference portfolio collateral type	
Tranche	High-yield	Investment-grade
Super-senior	20	12
Senior (AAA)	15	6
Junior mezzanine (A/BBB)	8	4

Data cover balance-sheet synthetic CDOs reported by www.creditflux.com over 2002–03.

FIGURE 3. CREDIT ENHANCEMENT AND LEVERAGE OF SYNTHETIC CDO TRANCHES
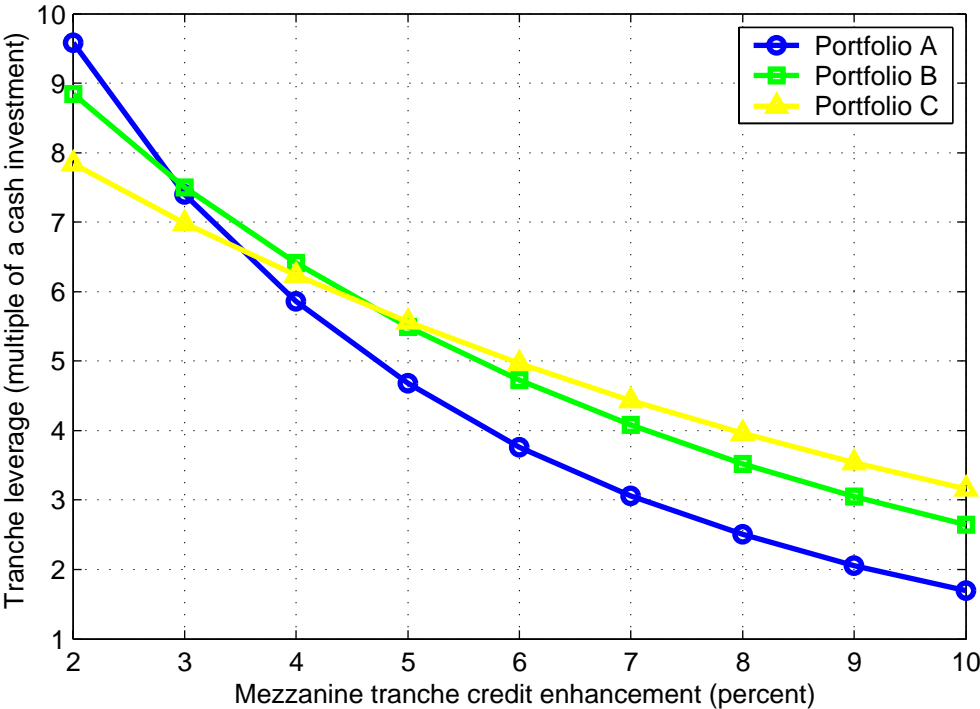
Judging from the types of CDO deals observed in the market, investors appear to prefer mezzanine tranches on higher-risk reference portfolios with less leverage and more credit enhancement. Mezzanine tranches on low-risk, low-spread reference portfolios with more leverage and less credit enhancement are rarely issued.

5.3 Risk transfer is limited when the equity tranche is retained

In CDOs that transfer exposures off a bank's balance sheet, the originating bank often retains the first-loss equity tranche. There are two sound reasons for this practice. First, to limit informational asymmetry: because the bank chooses which credits are put into the reference portfolio, and often retains monitoring or servicing responsibility for credits it originated, it can influence the performance of the reference portfolio. CDO investors prefer the originating bank to have its own money at stake, in a first-loss position. This puts the CDO originator's profitability, as well as its reputation, at stake. Second, the CDO originator may find the equity tranche to be a good match for its desired risk-return and diversification targets.

I earlier described how the equity tranche bears a large share of the CDO's risk but only a small share of the notional exposure amount. When a bank transfers credits from its balance sheet into a CDO but retains the equity tranche, it has transferred nearly all of the notional exposure but much less than all of the credit risk. Banks often did such balance sheet CDOs where equity tranche risk was retained, especially in the early days of the credit derivatives market (late 1990s), but some market participants may not have understood that the amount of risk transfer occurring in these transactions was smaller than the notional amounts involved may have suggested.

5.4 CDO structuring

So far the modeling and discussion of CDOs has been stylized. In reality, the structure of a CDO transaction can be idiosyncratic and can dramatically affect the returns that CDO investors earn. Three important elements of a CDO's structure, none which are present in the stylized model used in this paper, are:

- a payment waterfall: contractual terms that determine the priority of cashflows out of the CDO;

- "trigger" provisions that give additional credit enhancement to senior tranches by diverting cashflows away from more junior tranches to more senior tranches if the reference portfolio par value or interest proceeds decline below a certain level;

- covenants restricting the CDO asset manager's choice of credits; for example, a lower bound on the average credit rating or average spread of the reference portfolio, or an upper bound on the percent of the reference portfolio invested in assets with a low credit rating.

17

These features are more common in cash CDOs.

In the early days of the credit derivatives market, some CDO investors paid too little attention to the details of the CDO structures they were investing in. These investors did not look past the CDO tranche's credit rating to fully understand the structural features of the CDO. Subsequently, these investors were surprised by the poor performance of some of their CDO investments.

To the extent that these structuring considerations are important for a given CDO, the stylized model used in this paper would not be adequate for pricing or risk measurement. The model in this paper would be more appropriate for transactions with simple structures, generally including synthetic CDOs and specifically, single-tranche synthetic CDOs, discussed in section 6.1 below.

6   CDO tranches are exposed to correlation risk

The pricing of CDO tranches reflects investors' expectation of the correlation of defaults in the reference portfolio throughout the life of the CDO. Since this correlation is unobservable, CDO tranches are exposed to correlation risk. Figure 4 shows how the mark-to-market value of the three hypothetical CDO tranches varies with the correlation among the credits in the reference portfolio. Recall that the baseline assumption is asset correlation across credits equal to 0.3. By taking the baseline par spreads from Table 2 and revaluing each tranche as the asset correlation across credits is varied between 0 and 0.9, Figure 4 shows how correlation affects the value of the three CDO tranches.

The effect of correlation on CDO tranches is intuitive. A higher correlation of defaults implies a greater likelihood that losses will wipe out the equity and mezzanine tranches and inflict losses on the senior tranche. Thus, the value of the senior tranche falls as correlation rises. Conversely, higher correlation also makes the extreme case of very few defaults more likely. Thus, the value of the equity tranche rises as correlation rises. Equity tranche investors gain more in a scenario with very few defaults than they lose from a scenario with many defaults (they are only exposed to the first few defaults). Mezzanine tranches are subject to both effects, which can broadly cancel each other out make mezzanine tranches less sensitive to correlation, as in Figure 4.

Correlation effects are particularly strong for some of the most recent innovations in credit markets, including single-tranche CDOs and first-to-default basket swaps, which are discussed below.

Figure 4. Effect of correlation on CDO tranches

Figure 4. Effect of correlation on CDO tranches
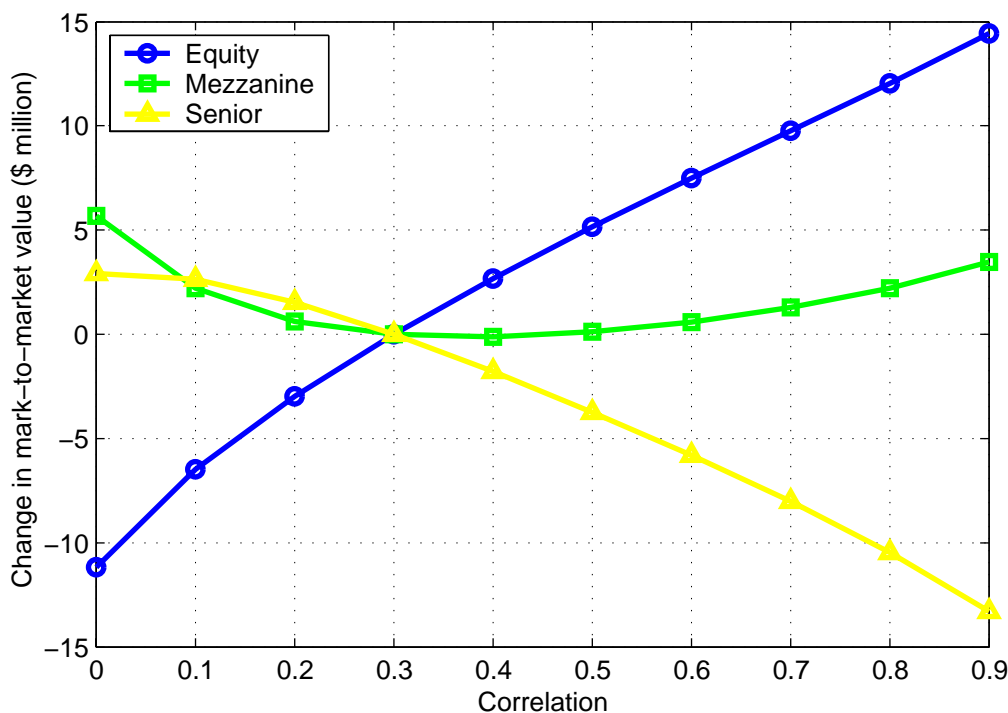
6.1 Single-tranche CDOs

In a traditional CDO, the underwriter purchases and warehouses a portfolio of bonds, loans, or credit default swaps and, as close to simultaneously as possible, places the equity, mezzanine, and senior tranches of the CDO's capital structure with investors. Except when the underwriter retains a tranche on its books or earns an ongoing fee as CDO asset manager, its risk exposure is limited to the time from when the credit risk is warehoused until the CDO tranches are sold. After all the tranches of a traditional CDO are sold, the underwriter no longer has exposure to the credits in the CDO's portfolio.

One difficulty for underwriters in a traditional CDO is that all tranches, across the entire capital structure, must be sold in order to do the deal. If no senior investors can be found, the deal cannot be done no matter how eager the equity and mezzanine investors are. Because it has become difficult for dealers to find investors across the entire capital structure at the same time, dealers have recently begun to offer single-tranche CDOs to investors. Thanks to the added flexibility of single-tranche CDOs, they now outnumber full capital structure CDOs among new deals, according to market observers.

In a single-tranche CDO, only one tranche of the capital structure is sold to an investor. Instead of purchasing and warehousing the entire reference portfolio, the dealer hedges its

			(Dollar amounts in millions)	
Tranche	Attachment points (percent)	Notional amount	Protection sold on each reference credit	Total protection sold on 100 reference credits
Equity	0–3	30	4.83	483
Mezzanine	3–10	70	4.71	471
Senior	10–100	900	2.17	217

exposure on just the single tranche that is sold.[15] The dealer hedges by selling credit protection on the names in the reference portfolio in an amount that offsets the dealer's exposure to spread movements or defaults in those names. These amounts change over time as the level of credit spreads changes, thus single-tranche CDOs require dynamic hedging.

The three tranches of the hypothetical CDO presented earlier can be used to illustrate the calculation of single-tranche CDO hedge amounts. Table 7 shows the hedge positions that would be needed if a dealer sold one of these tranches singly.[16] The hedge instruments are 5-year single-name CDS on the 100 underlying credits in the reference portfolio. Because the hypothetical CDO assumes that all credits in the reference portfolio have the same single-name CDS spread, the hedge position is the same for each of the 100 reference credits. For example, the second row of the table shows that a dealer who has sold the mezzanine tranche to an investor, effectively buying credit protection from the investor, would sell $4.71 million of protection on each of the 100 names in the reference portfolio, or $471 million in total, to hedge the spread risk of the mezzanine tranche.

The dynamic hedging of single CDO tranches presents dealers with four new risk management challenges that are not present in traditional CDOs. First, the dealer must calculate the hedge positions (or "deltas") of each credit default swap in the reference portfolio. These deltas are model-dependent, thus single-tranche CDOs leave dealers with model risk that would not exist with a traditional CDO. Second, as deltas change over time and the dealer dynamically adjusts its hedges, it is exposed to liquidity risk. The credit default swap market may not have enough liquidity for the dealer to adjust its hedge as desired without incurring high trading costs. Third, a dealer prefers to be hedged against both small moves in credit spreads ("spread risk") and unexpected defaults ("jump-to-default risk"). Hedging against both risks adds complexity that the simple example in Table 7 ignores. Fourth, the value of a single tranche CDO depends on the correlation assumption that is used to price and

---

[15]Or, the dealer can purchase the entire portfolio and hold the unsold tranches in its portfolio.

[16]Because the tranche payoffs are nonlinear, the hedge amounts cannot simply be added together to determine the hedge position if a dealer has sold more than one tranche. In a previous version of this working paper, I erroneously suggested they could. I thank Janet Tavakoli for pointing out my error.

hedge it. Since default correlation is unobservable, a dealer is exposed to the risk that its correlation assumption is wrong ("correlation risk").

Given the less-than-perfect liquidity of the credit default swap market and the many new risks that a seller of single tranche CDOs has to manage, market participants report that risk management of these products is currently more of an art than a science.

6.2 First-to-default baskets and correlation trading

The growth of new credit portfolio products, such as single tranche CDOs, has given rise to what is known as "correlation trading." By delta hedging a single-tranche CDO, a dealer can eliminate most, if not quite all, of the sensitivity to individual credit spreads in the reference portfolio. What remains unhedged is the exposure to correlation. When market participants disagree about the correct correlation to apply to a CDO tranche, they will put different values on the tranche. Those who estimate a lower value will sell to those who estimate a higher value. Different perceptions of hedging costs, different model structures or different model parameters can also lead market participants to disagree about the value of single CDO tranches.

First-to-default baskets are another example of a correlation product. An investor in a first-to-default basket receives a periodic spread payment and agrees to bear a loss when any one of a specified group of credits defaults. Like equity tranches of CDOs, first-to-default baskets have a first loss exposure to the reference portfolio, offer leveraged exposures to credit risk, and their value depends on correlation. Unlike CDO tranches, first-to-default baskets typically involve small reference portfolios (five names is typical). First-to-default baskets can be valued using the same model presented earlier for valuing CDO tranches.

Table 8 shows a hypothetical example of a first-to-default basket. The five credits in the basket are shown in the table, along with their single-name CDS spreads at the five-year maturity. The basket contains $10 million notional amount of each credit. At the first default of any credit in the basket, the protection buyer delivers $10 million notional amount of the defaulted credit's debt to the protection seller in exchange for $10 million cash. After one default the basket terminates.

Table 9 shows why first-to-default baskets are a "correlation product." The table shows the par spread at which the first-to-default basket would be priced, for different levels of correlation across credits. An investor who thinks the correct correlation is 0.7 could purchase the first-to-default basket from a dealer who thinks the correct correlation is 0.3 and receive 301 basis points of spread for a risk that the investor thinks should only be worth 217 basis points. The table also shows how the mark-to-market value of the first-to-default basket changes with correlation, assuming an initial correlation of 0.3. As the example shows, investors with different views on correlation will see value in trading correlation products.

TABLE 8. REFERENCE PORTFOLIO OF A HYPOTHETICAL FIRST-TO-DEFAULT BASKET

Credit	5-year CDS spread
BAE Systems	80
Deutsche Telekom	67
General Motors	146
Nokia	25
Norfolk Southern	34

Note: The basket contains $10 million notional amount of each credit.

TABLE 9. CORRELATION RISK IN THE HYPOTHETICAL FIRST-TO-DEFAULT BASKET

Correlation	Par spread (basis points)	Change in basket's mark-to-market value from correlation=0.3 ($ thousands)
0	352	−195
0.1	337	−136
0.2	320	−71
0.3	301	0
0.4	282	76
0.5	261	158
0.6	240	245
0.7	217	338
0.8	193	438
0.9	168	548

Correlation trading is the most sophisticated segment of the credit derivatives market. Reportedly, large dealers and sophisticated hedge funds are most active in this market segment.

7   CDO tranches are sensitive to the business cycle

Because CDO tranches are sensitive to correlation, and correlation of defaults is typically driven by the business cycle, the correlation risk of CDO tranches can also be characterized, and measured, as "business cycle risk." Using the model described in section 3 above, and interpreting the common factor as business cycle risk that is common to all credits, I can compute the exposure of each of the tranches of the hypothetical CDO to business cycle conditions. Specifically, I can compute the expected loss (EL) on the CDO tranche as defined above, *conditional* on a certain value of the common factor.

Table 10 shows such a calculation. Three different business cycle conditions are considered: boom, trend growth, and recession. These correspond to setting M, the common factor driving defaults, at its 10th, 50th, and 90th percentiles, respectively. The table shows both the dollar amount of each tranche's conditional EL in the boom, trend growth and recession scenarios, and the conditional EL as a percent of the tranche's notional amount.

The equity tranche, in a first-loss position, expects to bear defaults of about half its notional amount in a trend growth macroeconomic scenario and expects to lose its entire notional amount in a recession.

The mezzanine tranche, in a second-loss position, suffers no losses in a boom and minimal loss in a trend growth scenario, but suffers most of the portfolio's EL in a recession. In this sense, mezzanine tranches are leveraged bets on business cycle risk. Recall the hypothetical CDO's mezzanine tranche. Its par spread is 315 basis points, compared with 60 basis points on the reference portfolio. In exchange for this higher return, the mezzanine investor is exposed to a loss of 64 percent of principal in a recession scenario, compared with 7.6 percent on the reference portfolio.

The senior tranche expects to suffer very little loss, even in a recession scenario. Figure 5 shows the expected loss on the three tranches across a full range of macroeconomic shocks (1st to 99th percentile). Beyond the 96th percentile common factor shock, corresponding to a less than 4-in-100 or less-than-once-per-25-years shock, the senior tranche begins to see its principal significantly eroded by additional losses. While the senior tranche is not exposed to "recession risk," it could be said to be exposed to "depression risk."

Comparing the conditional EL for trend growth in Table 10 with the unconditional EL in Table 4, we see that the two are similar for the equity tranche. However, for the mezzanine and senior tranches, the unconditional EL is much larger than the EL conditional on trend

TABLE 10. EXPECTED LOSS ON HYPOTHETICAL CDO TRANCHES CONDITIONAL ON THE COMMON MACROECONOMIC RISK FACTOR

(Dollar amounts in millions)

| | | | Macroeconomic environment | | | | | |
| | | | Boom | | Trend growth | | Recession | |
Tranche	Attachment points (percent)	Notional amount	Expected loss	Percent of notional	Expected loss	Percent of notional	Expected loss	Percent of notional
Equity	0–3	30	1.4	5	14	47	30.0	100
Mezzanine	3–10	70	0	0	0.3	0.4	44.8	64
Senior	10–100	900	0	0	0	0	1.3	0.1
Entire portfolio	0–100	1,000	1.4	0.14	14.3	1.4	76.1	7.6

Note: "Boom" corresponds to setting the common macroeconomic risk factor at its 10th percentile, "trend growth" corresponds to the 50th percentile, and "recession" corresponds to the 90th percentile.

FIGURE 5. EXPECTED LOSS ON HYPOTHETICAL CDO TRANCHES IS CORRELATED WITH MACROECONOMIC CONDITIONS
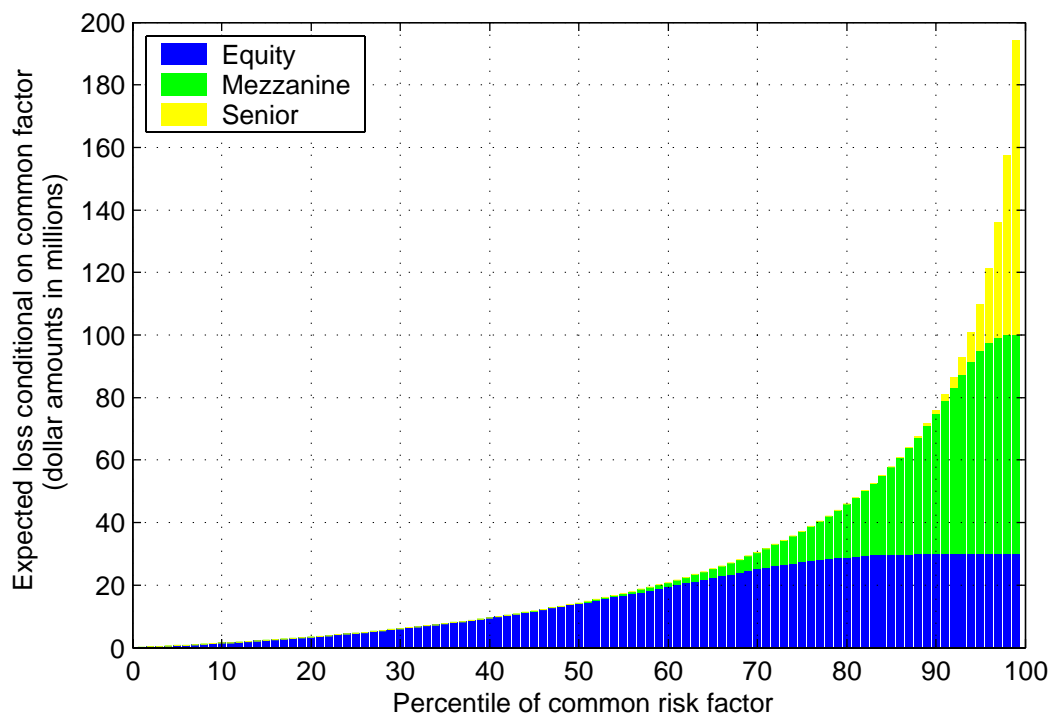
growth. This reflects the skewness of the loss distributions for the mezzanine and senior tranches and the lack of skewness in the equity tranche's loss distribution.[17]

8 Conclusions

Synthetic CDOs are popular vehicles for transferring the credit risk of a portfolio of assets. Using a pricing model for CDO tranches that does not require Monte Carlo simulation, I analyze the risk characteristics of the tranches of synthetic CDOs. In a hypothetical CDO, the equity and mezzanine tranches contain 10 percent of the notional amount of the CDO's reference portfolio but 70–90 percent of the credit risk. This implies that credit risk disclosures relying on notional amounts are especially inadequate for firms that invest in CDOs.

A basic result is that equity and mezzanine tranches are leveraged exposures to the underlying credit risk of the CDO's reference portfolio. I explore several implications of this result. First, even though mezzanine tranches are typically rated low-investment-grade, the leverage they possess implies their risk (and expected return) can be many times that of a low-investment-

---

[17]Amato and Remolona (2003) discuss the skewness of loss distributions in more detail.

grade corporate bond. Second, a mezzanine tranche's risk and leverage depend on the riskiness of the CDO's reference portfolio and the tranche's credit enhancement. Third, because the equity tranche contains a large fraction of the CDO's total risk, risk transfer is limited when the CDO originator retains the equity tranche.

CDO tranches and other innovative credit products, such as single-tranche CDOs and first-to-default basket swaps, are sensitive to the correlation of defaults among the credits in the reference portfolio. Because correlation is unobservable, differences of opinion among market participants as to the correct default correlation can create trading opportunities as well as "correlation risk" to be managed. Finally, the paper shows how the dependence of CDO tranches on default correlation can also be characterized and measured as an exposure to the business cycle, or as "business cycle risk." A mezzanine tranche, in particular, is highly sensitive to business cycle risk.

The paper highlights key issues for investors in synthetic CDO tranches as well as for dealers who structure synthetic CDOs for clients. Investors in mezzanine CDO tranches are taking on leveraged exposures to the underlying credit risk of the reference portfolio. A tranche's credit rating does not convey all aspects of the tranche's risk. If investors disclose the notional amounts of their portfolio, broken down by credit rating, the leveraged nature of a mezzanine tranche's risk exposure would not be obvious.

The paper also touches on some of the risks to dealers who structure and make markets in synthetic CDO tranches for clients. A complete set of synthetic CDO tranches may not be fully hedged by selling single-name credit default swap protection on the CDO's reference portfolio if the CDO tranches are structured as swaps whose payouts do not depend on the flow of income from the reference portfolio. Dealers in CDO tranches, including those who structure single-tranche CDOs, are exposed to model risk and, because of the dynamic hedging required, liquidity risk that are not present in traditional cash CDOs.

References

Amato, Jeffery and Eli Remolona, 2003. The credit spread puzzle. *BIS Quarterly Review* (December), 51–63.

Andersen, Leif, Jakob Sidenius, and Susanta Basu, 2003. All your hedges in one basket. *Risk* (November), 67–72.

Boscher, Hans and Ian Ward, 2002. Long or short in CDOs. *Risk* (June), 125–129.

Duffie, Darrell and Nicolae Garleanu, 2001. Risk and Valuation of Collateralized Debt Obligations. *Financial Analysts Journal* (January/February), 41–59.

Goodman, Laurie S., 2002. Synthetic CDOs: An Introduction. *Journal of Derivatives* (Spring), 60–72.

Hull, John and Alan White, 2003. Valuation of a CDO and an n^{th} to default CDS without Monte Carlo simulation. Working paper, University of Toronto.

Li, David, 2000. On default correlation: a copula function approach. *Journal of Fixed Income* (March), 115–118.

Schmidt, Wolfgang and Ian Ward, 2002. Pricing default baskets. *Risk* (January), 115–118.

Vasicek, Oldrich, 1987. Probability of loss on loan portfolio. Working paper, Moody's KMV (can be downloaded at `http://www.moodyskmv.com/research/whitepaper/Probability_of_Loss_on_Loan_Portfolio.pdf`).