On Risk Neutral Pricing of CDOs

Roy Mashal[♦]
Columbia Business School

This Version: April 1st, 2002

Abstract

The aim of this paper is to explain the risks that are associated with the standard application of risk-neutral pricing for multi-name credit derivatives and especially to CDOs. Until the recent years CDOs were priced using rating based models, i.e. the rating companies rate the different tranches of the structure and the tranches were priced based on the spread of comparable rated bonds. This pricing method has two important shortcomings. First it is not applicable to the equity tranches, as these are not rated. Second, it ignores market data of the underlying portfolio; it resembles pricing a credit default swap based on the rating of the name. As a result, in recent years there is a strong tendency to develop new models that are based on risk-neutral pricing. Many market participants that are trying to develop such models have observed that the market prices are very different than the prices that are the outcome of their new models.

The problem with the standard application of risk-neutral pricing to multiname credit derivatives is that the risk-neutral measure for each underlying has a considerable component that reflects non-default risks. When pricing a multi-name instrument these non-default parts of the risk-neutral measure accumulate almost linearly. There is no economic reasoning for this rapid accumulation and that makes the disparity between the model prices and the market ones. We illustrate the main idea in a series of examples and give, mainly for illustration purposes, an alternative model that overcomes this problem.

J.E.L. Subject Classification: G11, G12, G13.

Key words: Credit risk, credit derivatives, portfolio models, CDO.

[♦] I would like to thank Mark Broadie, Paul Glasserman, Suresh Sundaresan, Assaf Zeevi and seminar participants at Radian Asset Assurance for helpful discussions and suggestions. Comments are very welcomed, E-mail: rm586@columbia.edu. Current version available at: http://www.columbia.edu/~rm586/.

1 Introduction

This paper is concerned with the risks involved in using the current risk neutral pricing models to price Collateralized Debt Obligations (CDOs). CDOs are of the more complicated instruments in the market; their payoffs depend mainly on the credit performance of a large portfolio of reference names as well as on a set of payoff rules, such as Over Collateralization (O/C), Interest Coverage (I/C), and unwinding option for the equity tranch. In the early years of the CDO market, its immaturity, the complexity of the instruments, and technical difficulties led to rating based pricing of CDO tranches. In this approach, the rating agencies rate the different tranches and the tranches are then priced by adding some spread over comparable rated bonds. This pricing method is obviously simple to use and rationalize. However, rating based pricing ignores market data, namely - spreads of the underlyings. An extreme comparable example of rating based pricing, would be pricing a credit default swap based on the rating of the name. As a result of the rating based pricing limitations as well as advances in technology, more sophisticated models have been developed in recent years throughout the industry. These new models are based on implementations of the risk neutral pricing methodology for pricing multiname credit derivatives. The purpose of this paper is to explore the application of risk neutral pricing to multiname credit derivatives and to highlight some very important implicit implications of the use of these models.

Throughout this paper we will decompose the observed spreads (usually referred to as risk neutral spreads - S_{RN}) in the market into two components, namely default spread (S_D) - the part of the spread that reflects the default risk, and the non-default part of the spread (S_{ND}) that is usually explained by other sources of risk such as liquidity risk, recovery risk, risk premium etc. The risk-neutral spread can therefore be decomposed by the following relationship:

$$S_{RN} = S_D + S_{ND} \qquad (1)$$

Similarly, we would decompose the premium of the instrument that is priced (e.g. first-to-default basket) – PM to the part that is due to default PM_D, and the part that is due to non-default risks $PM_{ND.}$

The important and overlooked phenomena that we would explore in this paper is the behavior of the non-default part of the premium as the dimensionality of the instrument increases under the application of risk neutral pricing methods. As we will demonstrate, the non-default spreads are monotonically increasing with the dimension (number of underlying names) of the instrument (see Figure 1). For example, the non-default part of the premium on a first-to-default basket with ten reference names would be almost double that of a first-to-default basket on five names with the same characteristics for the underlyings. The same would hold when comparing CDO with 50 underlyings to one that has 100 of them. There is no economic reason for this accumulation of the non-default spreads, rather it is an implicit outcome of the current implementations of risk neutral pricing in the real world, i.e. in a world where bond prices reflect other sources of risk other than default risk. In the following we will show several examples of the accumulation of non-default spreads, discuss in more detail the reasons for this phenomena and offer, mainly for illustration purposes, a simple model that will overcome the problem.

2 Some Preliminaries

The new pricing models for multiname credit derivatives are based on risk neutral pricing. The pricing process usually follows the following steps:

1. Infer risk neutral single name default probabilities from outstanding single name instruments.[1]
2. Apply some risk neutral dependence structure among the names.
3. Price the multiname instrument.

Different models usually differ in step 2; the dependence structure is sometimes assumed to be induced by correlated hazard rates for the different names (Duffie [1998], Duffie

[1] See for example Hull and White (2000).

and Singleton [1999b], Duffie and Gurleanu [2001]), infectious defaults (Jarrow and Yu [2000] and Davis and Lo [1999a, 199b]), Normal copula (Hull and White [2001]), or a t-copula (Mashal and Naldi [2002]). However different, all the models rely on the risk neutral pricing technology, and ultimately share the same pitfall that is of our current concern.

We will try to illustrate the point of this paper in a series of examples, as we believe it would be the clearest method. In the examples that follow we will fix the default and non-default parts of observed spreads and explore the breakdown of the premium for multiname credit derivatives. The ideas will be illustrated using Nth-to default baskets and CDOs. Before proceeding to the examples, we will review some recent empirical papers that try to determine the part of the observed market spreads that is due to the risk of default (more details in the Appendix).

Elton et. al. (2001) compare observed spreads to historical default and recovery rates and find that default risk by itself explains a very small part of the observed spreads in the market. Huang and Huang (2001) take a different approach, calibrating four different structural default models to historical data, but not to spreads. They find, as well, that the risk of default explains only a small part of observed spreads. Overall, the part of the spread that is explained by the risk of default usually accounts for as much as 30%-70% of observed spreads, even after taking into account the tax asymmetry between treasuries and corporate bonds. The objective measure of default is only 30%-70% of the risk neutral one. The difference between the two measures of default is usually explained by risk premium, liquidity risk, recovery rate risk and other sources of risk, all of which we refer to as non-default risks.

3 Illustrations and Examples

In the examples that follow we make a similar breakdown of the overall risk neutral premium into two parts: a) The part that is due to default by pricing using the objective measure of default and b) The part that is due to the non-default risks, implied by the difference between the overall premium and the premium that is due to default (as in Eq. 1):

$$PM_{ND} = PM - PM_D$$

In order to demonstrate our argument we will fix in the following examples the part of the objective measure of default risk (for the underlyings) to be 50% of the risk neutral (observed) one and the part of the non-default risks, to account for the other 50%. Hence, if we observe a spread of 1% (over Libor) it reflects that the spread that is due to default risk is 0.5% and the spread that is due to non-default risks is also 0.5%.

Example 1: Credit default swap

Suppose that we want to price a credit default swap on a bond that pays a spread of 1% over LIBOR, where this spread is again composed of 0.5% for default and 0.5% for non-default risks. Under some simplifying assumptions[2] the credit default swap can be replicated with a money market account and the underlying bond, therefore its premium should equal the spread paid by the underlying bond. The result of the above argument is that the premium for the CDS would be PM=1% that can be decomposed again to PM_D=0.5% for default and PM_{ND}=0.5% for the non-default risks. This can be thought of as a 'one dimensional version' of a first-to-default basket.

Our next examples demonstrate first-to-default baskets valuation. For simplicity and to avoid model dependent prices we will assume the following:[3]

1. One-year horizon.
2. Zero short rate.
3. Independence among the names in the basket.
4. Zero recovery rate (R=0).

In such a simplified case, the premium of the basket should account for the probability of having at least one default while taking into account the recovery rate. For example, two

---

[2] See Tolk (2001).

[3] All of those assumptions are for ease of demonstration only; qualitatively the results will follow for any values of these parameters (with correlation<1 and R<1).

firms with probability of default through the horizon P(A defaults)=P(B defaults)=1% would have:

$$P(Default) = P(A \bigcup B) = 1 - P(\overline{A} \bigcap \overline{B}) = 1 - P(\overline{A}) \cdot P(\overline{B}) = 1 - 0.99^2 = 0.0199$$

therefore the premium, PM, should be:

$$PM = E[payoff] = 0 \cdot [1 - P(Default)] + (1 - R) \cdot P(Default) = 1 \cdot 0.0199 = 1.99\% \qquad (2)$$

And in the general case where there are n names in the basket, and, p=Pr(name j defaults within the horizon) for all j.

$$PM = (1 - R) \cdot P(Default) = (1 - R) \cdot [1 - (1 - p)^n] \qquad (3)$$

Example 2: First-to-default on a two name basket

As seen in Eq. 2, the premium of a first-to-default basket on two names would be PM=1.99%. Since we assume specific default and non-default spreads we can also calculate the breakdown of the premium for this basket:[4] pricing with the objective measure of default (using Eq. 3 and p=0.5%) gives:

$$PM_D = 1\text{-}(1\text{-}0.005)^2 = 0.9975\%$$

and we can deduce that the implied non-default premium is:

$$PM_{ND} = PM - PM_D = 1.99\%\text{-}0.9975\% = 0.9925\%.$$

[4] It can be shown that replication of first-to-default baskets with single name instruments, e.g. credit default swaps, is not possible and therefore the argument of pricing by replication does not hold.

Example 3: First-to-default on a three name basket

The same calculations for a basket with three names would result in:

$$PM = [1-(1-0.01)^3] = 2.97\%$$

$$PM_D = [1-(1-0.005)^3] = 1.49\%$$

And therefore, the non-default would be:

$$PM_{ND} = PM - PM_D = 2.97\% - 1.49\% = 1.48\%$$

These examples show that increasing the number of names in the basket from 1 to 2 to 3 results in an increase in the non-default part of the premium from 0.5% to 0.99% to 1.48%. Again, there is no economic justification for this (almost linear) increase in the non-default part of the overall premium of the instrument.

Example 4: First-to-default dimensionality study

The previous examples show in detail the behavior of the non-default part of the spread for baskets with few names. Table 1 continues this line of investigation for baskets of size up to 50 names. Although first-to-default baskets usually have up to ten reference names we provide this table because the same logic would apply to CDOs that usually have 50-100 assets. Table 1 follows the previous calculations for baskets with varying number of reference names, each has a risk-neutral default probability of 1% and an objective measure of default of 0.5%. The table extends our previous observation to higher dimension baskets.

Number of reference names	PM calculated using S_{RN}	PM_D calculated using S_D	Implied non-default premium PM_{ND}
1	0.01	0.005	0.005
2	0.0199	0.009975	0.009925
3	0.029701	0.014925	0.014776
4	0.039404	0.01985	0.019553
5	0.04901	0.024751	0.024259
10	0.095618	0.04889	0.046728
15	0.139942	0.072431	0.067511
20	0.182093	0.09539	0.086704
30	0.2603	0.139616	0.120684
40	0.331028	0.18168	0.149348
50	0.394994	0.221687	0.173306

Table 1: Breakdown of the risk neutral premium.
The table shows the risk neutral premium, the part of it that is due to default risk, and the part of it that is due to the non-default risks.

The premiums for the non-default risks are monotonically increasing in the number of names as can be seen in Figure 1.

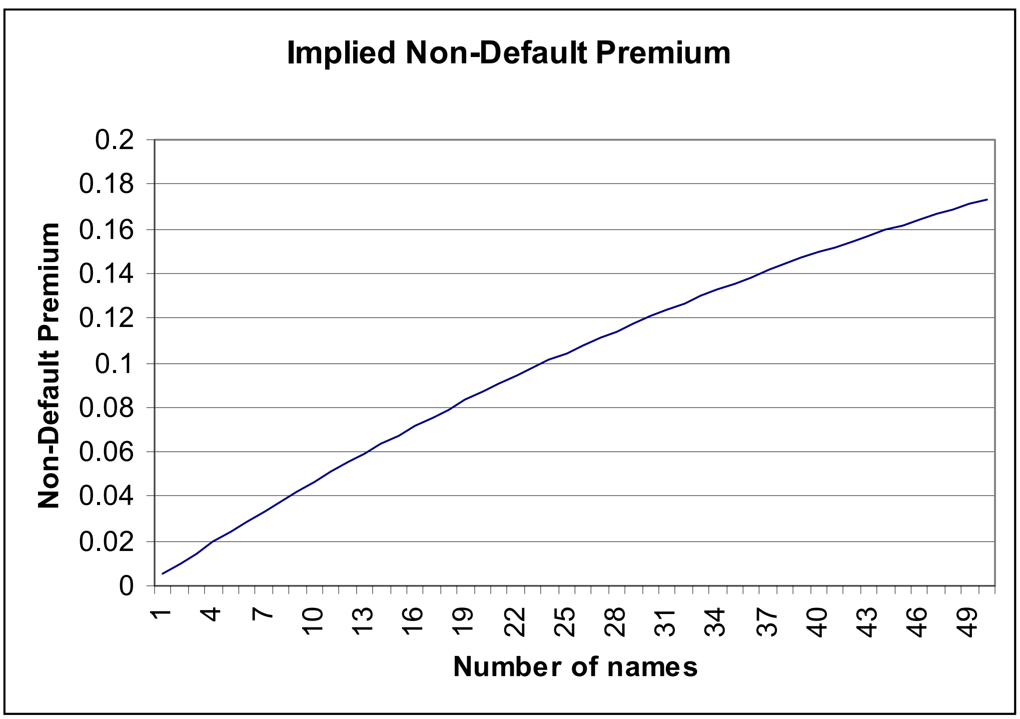

Figure 1: Non-default premium as a function of the number of names in the basket.
The figure shows the implied non-default premium as a function of the number of names in the basket. These premiums reflect the liquidity and other risk premiums that are implicitly priced by the use of risk neutral pricing.

After detailing the impact of the application of risk neutral pricing to first-to-default baskets, we next show the impact on the pricing of CDOs in the next example.

Example 5: CDO Valuation

Consider a simplified CDO, underlying are 50 names that yield Libor+350 bps each. Libor is constant at 5%; recovery rate is assumed 50% and paid at maturity. Three-year horizon. The CDO has three tranches:

1. Senior tranch promises Libor+50 bps.
2. Mezzanine tranch promises Libor+200 bps.
3. Equity tranch gets the reminder.

Suppose also that the CDO experiences constant annual default probabilities (CADP) for all the underlyings throughout the horizon. Table 2 gives the yields of the three tranches as a function of the CADP, assuming they sell for par.[5]

CADP	Equity	Mezzanine	Senior
0%	31.00%	7.00%	5.50%
1%	26.42%	6.98%	5.50%
2%	21.70%	6.88%	5.50%
3%	16.92%	6.68%	5.50%
4%	12.06%	6.34%	5.50%
5%	7.24%	5.90%	5.50%
6%	2.46%	5.34%	5.50%
7%	-2.32%	4.66%	5.50%
8%	-7.02%	3.90%	5.50%
9%	-11.70%	3.04%	5.50%
10%	-16.32%	2.10%	5.48%
11%	-20.84%	1.08%	5.48%
12%	-25.32%	0.00%	5.48%
13%	-29.74%	-1.12%	5.46%
14%	-34.10%	-2.34%	5.46%
15%	-38.40%	-3.58%	5.44%

Table 2: CDO Yields as a Function of CADP.
The table displays the yields of the different tranches of the simplified CDO assuming they sell at par.

[5] The yields were computed using simulation under a normal dependence structure with uniform 20% correlation.

For example, assume that we observe spreads that correspond to 6% default probabilities. Using conventional risk neutral pricing we would get 2.46% yield for the equity. However, the default-only yield corresponds to a default probability of 3%, i.e. 16.92% yield for the same equity tranch. This latter yield totally neglects the non-default risk premium. But the difference in the yields is staggering, and would grow monotonically with the number of underlying names in the CDO, similarly to Figure 1.

4 An alternative – the additive model

We next describe a simple alternative model that overcomes the problem that has been illustrated in the previous sections. The main aim in presenting this model is to add to the understanding of the problem and possible ways to overcome it. Note that this model is not applicable with the data currently at hand.

In this model investors should determine the premium that they desire to receive on top of the premium that is due to default. Instead of implicitly accumulating the non-default premium using the risk neutral pricing, investors specifically determine the premium that they require for all the non-default risks (PM_{ND}) and add it to the objective default premium (PM_D). This method consists of three steps:

1) Infer objective measure of default for the names, S_D for the underlyings.
2) Price the instrument based on the objective measures of default S_D, and calculate the premium for the instrument, PM_D.
3) Determine some non-default premium for the specific instrument PM_{ND}.
4) Add up and get the premium of the instrument: $PM = PM_D + PM_{ND}$.

Using this method one does not impose arbitrary non-default premium for the instrument, but rather imposes intended premium for the specific instrument. For example, suppose one wants to price a first-to-default basket with 10 names. She observes outstanding debt of the names and infers that the risk neutral probability of default (S_{RN}) is 1%. She further knows (for example using KMV's EDF) that the non-default spreads account for 50% of the observed spreads. That leads to the breakdown of the spreads to 50 bps due to the objective probability of default (S_D) and 50 bps due to non-default risks (S_{ND}). The investor determines that the non-default premium of the basket should be PM_{ND}=100 bps.

The conventional pricing would lead to a premium for the basket of 956 bps (Table 1, using S_{RN}) whereas the additive method would lead to a premium of 489 bps (Table 1, using S_D) for the default part, and 100 bps for the non-default part, or a total of 589. A difference of 367 bps between the premiums calculated using the risk neutral pricing and the additive method. What is the cause of this difference? Instead of arbitrarily imposing risk and liquidity premiums for the basket the investor uses specific desired premiums.

5 Discussion

In the above we have argued that the use of risk neutral pricing models for pricing certain contingent instruments should be reconsidered. As we have demonstrated, the use of these models for the pricing of multiname credit derivatives leads to an implicit accumulation of non-default premiums with no economic justification. The cause of this accumulation is that observed spreads of default-risky instruments in the market represent not only default risk but other risks as well. As recent research show, the part of the spread that is not due to default risk is very large, and as a consequence when one uses risk neutral pricing models the accumulation of the non-default premiums is very rapid in the number of names in the basket.

The main focus of the paper is the risk that is associated with the use of the current pricing models. For demonstration purposes we also provide an alternative pricing method that overcomes the problem involved with the use of current models. Although we illustrate the problem with multiname credit derivatives, similar deficiencies, to different extents, arise in pricing models of other instruments as well. Practitioners have to be very aware to the consequences of the use of current models. It is essential that both academics and practitioners realize that the current models are far from accurate and that there is a great need to develop new models.

Appendix

In this appendix we will review two recent empirical papers, Elton et. al. (2001) and Huang and Huang (2001). These two papers are concerned with the decomposition of spreads to parts that are due to default risk and to parts that are due to non-default risks.

Elton et. al. (2001) investigate historical spread and default data and explain the spread by three factors: default risk, tax asymmetry, and risk premium. Their results find that default risk explains a very small portion of observed spreads. We next summarize their evidence

Elton et. al. are using one year historical default rates from Moody's and S&P and make a simplified assumption that the default process is stationary and Markovian, out of which they generate default probabilities for two to ten years for the different ratings AAA to CCC. Next they use historical recovery rates for theses ratings, and the historical spreads over treasuries (industrial sector) for those ratings and maturities of two to ten years. If the default risk was the only relevant risk and investors were risk neutral one could compute the generic spreads that are due to default risk. Table 3 shows the historical spreads and the "default spread" for AA, A, and BBB industrial ratings.

Year	AA Spreads			A Spreads			BBB Spreads		
	Historical	Default	Explained	Historical	Default	Explained	Historical	Default	Explained
2	41	0	1.0%	62	5	8.5%	117	15	12.4%
3	42	1	1.9%	68	6	9.3%	121	18	15.0%
4	46	1	2.6%	72	7	10.3%	121	22	17.9%
5	49	2	3.4%	74	8	11.4%	121	25	20.9%
6	53	2	4.4%	75	10	12.6%	120	29	23.9%
7	55	3	5.1%	76	11	13.9%	119	32	26.7%
8	57	3	5.9%	77	12	15.1%	119	35	29.5%
9	59	4	7.0%	78	13	16.4%	118	38	32.1%
10	60	5	8.0%	79	14	17.8%	118	41	34.7%

Table 3: Historical Spreads and Spreads due to Default.
The table shows historical observed spreads and spreads that are due to historical default rates for different maturities and ratings

Table 3 shows that default spread is only a very small part of the overall spread of corporates over treasuries, for example the default risk for BBB bonds with 5 years to maturity explains only 21% of the observed historical spread for these bonds.

A common way to control for the tax asymmetry is to use the spread over Libor or over other corporates instead of the spread over Treasuries[6]. Table 4 shows the same data as in Table 3 but instead of using the historical spreads over treasuries it shows the historical spreads over AA rated corporate bonds.

Year	A Spreads			BBB Spreads		
	Historical	Default	Explained	Historical	Default	Explained
2	21	5	25.6%	75	15	19.3%
3	26	6	24.1%	79	18	23.0%
4	26	7	28.5%	76	22	28.7%
5	25	8	34.3%	71	25	35.4%
6	23	10	41.9%	67	29	42.5%
7	21	11	50.0%	64	32	49.8%
8	20	12	58.5%	62	35	57.1%
9	19	13	67.4%	60	38	63.9%
10	18	14	76.9%	58	41	70.9%

Table 4: Historical Spreads and Spreads due to Default Relative to AA Spreads. This table shows the same data as table 4 for A and BBB ratings relative to AA

For our purposes the main point is that credit risk accounts for a small portion (around 30-70%) of the total premium that corporate bonds offer over treasuries and over AA corporate bonds.

The second paper is Huang and Huang (2001), whose title is "How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk? A new Calibration Approach". As illustrated in the title, the aim of the paper is to quantify the part of spreads that is due to default risk. The usual method used with structural models calibration is to calibrate them to given spreads; the authors instead calibrate different models to historical default rates, recovery rates and other important parameters but <u>not</u> to spreads. Using such setting they get a very robust estimation procedure for the part of the spreads that is due to credit risk alone. Table 5 summarizes their findings.

[6] This does not fully account for the tax asymmetry because lower graded firms usually issue higher coupon bonds, and the tax burden of those is greater than that on higher grade, lower coupon bonds. However, it also reduces part of the default risk.

Rating	Historical Spread	Longstaff Schwartz		Leland-Toft		Anderson-Sundaresan-Tychon	
	Bps	Credit Spread	Explained	Credit Spread	Explained	Credit Spread	Explained
Aaa	63	10.5	16.7%	20.6	32.7%	16.7	26.5%
Aa	91	13.9	15.3%	19.2	21.1%	15.9	17.5%
A	123	20.7	16.8%	20.4	16.6%	18.6	15.1%
Bbb	194	47.7	24.6%	36	18.6%	39.5	20.4%
Bb	299	177.2	59.3%	113.2	37.9%		
B	408	375	91.9%	306	75.0%		

Table 5: Huang and Huang (2001) default explained part of spread (10 years, face recovery). This table summarizes the part of the spread that is explained by default risk using three different credit models.

Again we can see that the default risk explains a small part of observed spreads.

Table 6 computes the ratio between the default spread and the spread over a generic AA.

Rating	Historical Spread Over AA	Longstaff Schwartz		Leland-Toft		Anderson-Sundaresan-Tychon	
	Bps	Credit Spread	Explained	Credit Spread	Explained	Credit Spread	Explained
A	32	20.7	64.7%	20.4	63.8%	18.6	58.1%
Bbb	103	47.7	46.3%	36	35.0%	39.5	38.3%
Bb	208	177.2	85.2%	113.2	54.4%		
B	317	375	118.3%	306	96.5%		

Table 6: The Part of the Relative Spread Over AA Rating That is Explained By Default Risk. This table provides the same data as table 6 with historical spread relative to AA bonds.

Table 6 shows that even if we control in a rough way for the tax asymmetry portion of the spread we are still left with a very large part that is not explained by default risk.

References

[1] Davis, M. and V. Lo. 1999a. Infectious defaults. Working paper, Imperial College, London.

[2] Davis, M. and V. Lo. 1999b. Modelling default correlation in bond portfolios. In Carol Alexander, editor, *ICBI Report on Credit Risk*.

[3] Duffie, D. 1998. First-to-default valuation. Working paper, Graduate School of Business, Stanford University.

[4] Duffie, D. and N. Garleanu. 2001. Risk and valuation of collateralized debt obligations. Working paper, Graduate School of Business, Stanford University

[5] Duffie, D. and K. J. Singleton. 1999a. Modeling term structures of defaultable bonds. *Review of Financial Studies* **12** 687-720.

[6] Duffie, D. and K. J. Singleton. 1999b. Simulating correlated defaults. Working paper, Graduate School of Business, Stanford University.

[7] Elton, E., M. Gruber, D. Agrawal, and C. Mann. 2001. Explaining the rate spread on corporate bonds. *The Journal of Finance* **56(1)** 247-277.

[8] Huang, J. and M. Huang. 2001. How much of the corporate-treasury yield spread is due to credit risk? results from a new calibration approach. Working paper, Graduate School of Business, Stanford University.

[9] Hull, J. and A. White. 2000. Valuing credit default swaps I: no counterparty default risk. *Journal of Derivatives* **8(1)** 29-40.

[10] Hull, J. and A. White. 2001. Valuing credit default swaps II: modeling default correlation. *Journal of Derivatives* **8(3)** 12-22.

[11] Jarrow, R. A. and F. Yu. 2000. Counterparty risk and the pricing of defaultable securities. Working paper, Johnson GSM, Cornell University.

[12] Mashal R., and M. Naldi. 2002. Pricing multiname credit derivatives: heavy tailed hybrid approach. Working paper, Graduate School of Business, Columbia University.

[13] Tolk J. 2001. Understanding the risks in credit default swaps. Special report, Structured Finance, Moody's Investors Service.