

Implied Correlations: Smiles or Smirks?

ŞENAY AĞCA, DEEPAK AGRAWAL, AND SAIYID ISLAM

ŞENAY AĞCA

is an associate professor of finance at George Washington University in Washington, DC.
sagca@gwu.edu

DEEPAK AGRAWAL

is director of quantitative research at Diversified Credit Investments in San Francisco, CA.
dagrawal@dcinv.com

SAIYID ISLAM

is director of quantitative analytics and research at Standard & Poor's in New York, NY.
saiyid_islam@standardandpoors.com

We investigate whether the commonly observed "implied correlation smile" pattern in the market for standardized CDO tranches can arise from the various model simplifications in the industry standard one-factor Gaussian default time model that is used to compute the implied correlations. Our evidence shows that by ignoring empirical regularities like fat-tailed return distributions, heterogeneous pair-wise correlations, heterogeneous spreads, and correlation between default probabilities and recovery rates, the Gaussian model can give rise to smile patterns even when all the tranches are fairly priced. We find that among the various simplifications of the Gaussian model, the assumption of normal asset returns when they are actually fat-tailed is the most problematic one in the sense of generating the most pronounced implied correlation smile. Therefore, relaxing this assumption seems to be the most promising direction for CDO modeling. We also find that the assumption of homogenous spreads is the most benign one and may be retained if it helps with maintaining tractability. We also show that while the standard Gaussian model and the corresponding implied correlations can lead to erroneous inferences for valuation across tranches, they perform well for pricing a given tranche across time, even in periods of market dislocations.

Modern credit markets have a prominent focus on correlations among credits. A number of securities whose values depend

on correlations among a set of credits trade actively. The most popular ones include collateralized debt obligation (CDO) tranches and basket default swaps. One of the most important drivers of active trading of such securities has been the rapid standardization in the credit derivatives market. CDS indices such as the CDX North America Investment Grade (CDX.NA.IG) and iTraxx Europe are standardized, equally weighted, tradable portfolios of credit default swaps that have high liquidity. Furthermore, tranches on these standardized indices have themselves been standardized, and they also trade actively.¹

Correlation affects the values of these tranches because of its pronounced impact on the shape of the default loss distribution of the credit portfolios that underlie these securities. With active trading of these tranches, market participants have the opportunity to trade correlations, just like option traders trade volatility when they trade options. Borrowing from the analogy of "implied volatility" in the options market, it has become commonplace to talk about "implied correlation" to refer to correlation extracted from the observed tranche prices using a simple, standard model that can price the tranches as a function of the correlation among the underlying credits.

Market participants rely on these implied correlations for several types of inferences. They use them for comparisons across various tranches of a given reference portfolio.²

A common application is to interpolate or extrapolate the implied correlations of standardized tranches to infer the prices of nonstandard or bespoke tranches on the same or similar reference portfolios. This analysis is done using the implied correlations or a variant called “base” correlations. Market participants may also compare the implied correlations against empirically estimated historical correlations in the reference portfolio to infer any anomalies in the pricing of tranches. These and other applications of implied correlations raise a question about the drivers of implied correlations and, in particular, about any model dependencies in these numbers.

Just as the Black and Scholes [1973] model is the standard model used to obtain implied volatilities in the options market, one needs a standard tranche valuation model to infer implied correlations from the observed CDO tranche prices. It has become an industry practice to use a simple, one-factor Gaussian copula model introduced by Li [2000],³ to obtain implied correlations from tranche prices. This model expresses a tranche value as a function of the (risk-neutral) default probabilities on the underlying credits, the corresponding recovery rates, and correlations among credits. It further assumes that all pair-wise correlations among the credits are homogeneous. Therefore, given the individual (risk-neutral) default probabilities and recovery rates, one can back-solve the model to compute one implied correlation number from the observed price of each tranche.

When implied correlations are computed from a set of tranches on a given reference portfolio, the resulting correlations should all be theoretically identical because they all refer to the assumed homogeneous correlation among the credits in the reference portfolio. In practice, however, implied correlations computed from observed prices of standardized tranches on CDS indices show a pronounced and persistent U-shaped pattern. This pattern is often called the “correlation smile” and is analogous to the “volatility smile” seen in the equity options markets.

Why do we see a correlation smile? While mispricing across tranches cannot be completely ruled out, the pervasiveness of the smile across various tranches, across various reference portfolios, and across time suggests that the correlation patterns are driven by more fundamental factors that are common to all CDOs and across time. Implied correlations are computed using an industry standard one-factor Gaussian copula model (hereafter referred to as the standard Gaussian model) with a number of strong simplifying assumptions. For instance, it assumes

that the asset returns of the underlying names in the reference portfolio have a normal distribution. Credit spreads as well as pair-wise correlations across credits are assumed to be identical, and recoveries in the event of defaults are considered to be constant. Due to these simplifying assumptions, it is likely that this model is misspecified in several dimensions. Various model misspecifications can potentially cause implied correlations not only to be different from the true underlying correlations, but also to differ across the tranches of the same underlying reference portfolio. This behavior can show up as non-flat patterns in the implied correlation curves even when all the tranches are fairly priced.

While model misspecification is a potential channel for non-flat correlation curves, several questions remain. For instance, among the several possible misspecifications contained in the standard Gaussian model, which one(s) are the most crucial in generating a non-flat correlation curve? Conversely, which ones are the most benign and hence can be retained for the sake of simplicity or computational expediency? What are the implications for the way the correlation curves are used in the marketplace? In this article, we address these questions.

We use the following methodology. We assume a hypothetical reference portfolio of 125 credits with characteristics (such as spreads, recoveries, etc.) very similar to those of a standardized CDS index, namely, the CDX.NA.IG. We then compute the fair tranche values of a set of CDO tranches on this reference portfolio using a price-generating model. The tranches we use are the same as the standardized tranches on the CDX.NA.IG index that trade quite actively. Our price-generating model is a modification of the standard Gaussian model that relaxes its various assumptions, one at a time. Using this model, we generate the fair values of various tranches. The fair values are then inverted using the standard Gaussian model to compute the implied correlation curves. These curves represent the implied correlation patterns that would be obtained from the market prices if the real world price-generating process looked similar to the one considered by us.

We show that a non-flat correlation pattern can arise due to the various assumptions contained in the standard Gaussian model. Among them, we find that it is the assumption of Gaussian return distributions that is the most critical for distorting the shape of the implied correlation curves. We show that this assumption can cause a pronounced smile pattern that appears very similar to

what is commonly seen empirically using market observed tranche prices. On the other hand, we find that the assumption of identical spreads across names is the most benign as the effect of relaxing it causes only marginal distortions to a flat implied correlation pattern. When we relax the assumption of a constant recovery rate and allow for recoveries to be negatively correlated with the systematic default risk in the portfolio, consistent with empirical observations, we find that the resulting implied correlation patterns are distorted in a direction opposite of the commonly seen smile patterns. In this case, stochastic recoveries cause a “frown pattern” rather than a “smile pattern” in the implied correlation curve. This suggests that stochastic recoveries, negatively correlated with the default risk, are not the primary driver of the observed correlation smile patterns.

Given these results, can the implied correlations obtained from the standard Gaussian model still be useful for pricing purposes? We examine the performance of implied correlations computed from the standard Gaussian model for pricing a given tranche on a future date. We find that the standard Gaussian model prices various tranches successfully when the implied correlation of the same tranche from a previous date is used as a correlation input to the model. Furthermore, we find that it works well for pricing tranches across time, even in periods of credit market stress.

Our overall evidence suggests that the industry standard Gaussian copula default time model and implied correlations obtained from this model can be useful for pricing tranches across time, but should be used with caution for making inferences regarding mispricing between similar tranches of different CDOs or across different tranches of the same CDO. In particular, the assumption that returns are normally distributed rather than fat tailed can cause correlation smile patterns similar to those seen in market prices. Hence, if one has to focus efforts in a particular direction to improve the standard model, relaxing this assumption and allowing for fat tails in the return distributions would be the most promising.

IMPLIED VOLATILITY AND IMPLIED CORRELATION

Because the idea of implied correlations developed as a direct analog of the implied


volatility idea in the option pricing literature, we begin with a brief discussion of implied volatility and volatility smiles in equity options markets. We then examine implied correlation and correlation smile patterns commonly seen in the liquid market of standardized tranches on CDS indices.

Implied Volatility and Volatility Smile

Implied volatility is a commonly used construct in the equity options market. It refers to the volatility parameter in the Black–Scholes option pricing model that represents the volatility of returns on the underlying stock. It is computed by equating the observed option price to the Black–Scholes model value and then solving for the volatility parameter. The assumption is that all other inputs to the model are known. Implicitly, this calculation presumes that observed market prices are adequately described by the Black–Scholes model. If this assumption were indeed true, then implied volatilities computed from a set of options on the same underlying stock but with different strike prices would all be identical because they all represent the volatility of returns on the same underlying stock. However, in practice, the implied volatilities from such a set of options are not all equal. They exhibit a skew, very often a U-shaped pattern, commonly known as the volatility smile or skew. As an illustration, Exhibit 1 shows the skew observed in implied volatilities computed from a set of call options on the S&P 500 Index on a randomly selected day.

EXHIBIT 1 Implied Volatility Smile on S&P 500 Index Options

This exhibit shows implied volatilities obtained using call prices as of April 7, 2004, for options on the S&P 500 Index expiring on June 17, 2004. S&P 500 closed at 1140.53 on April 7, 2004.



The volatility smile curve suggests that the implied volatility is higher for out-of-the-money (OTM) options than for at-the-money (ATM) options. Prior research shows that the reason for this smile is mainly the discrepancy between the assumptions of the Black-Scholes model and the real world.⁴ The construct of implied volatility is based on the Black-Scholes model, but many of the assumptions of this model are not empirically valid. In particular, 1) return distributions are known to be fat-tailed while the Black-Scholes model uses a geometric Brownian motion to model asset returns, thus assuming a normal distribution, and 2) return volatility is known to be stochastic but the Black-Scholes model assumes it to be constant. The impact of either of these two violations is that large changes in the price of the underlying stock may be observed empirically with a higher frequency than what is assumed in the Black-Scholes model. OTM options, therefore, have a higher chance of becoming in the money than what is implied by the model. So, the market price of an OTM option is more than the model price. In computing implied volatility, one can match the market price with the Black-Scholes model price only by increasing the volatility parameter in the model. As a result, the implied volatility of OTM options is higher than that of ATM options. There are a number of studies that examine deviations from Black-Scholes model assumptions to fit the observed implied volatilities. For example, Das and Sundaram [1999] consider the impact of jump in stock returns as well as stochastic volatility to explain the term structure of implied volatilities.

Implied Correlation and Correlation Smile

Implied correlation is a construct that is analogous to implied volatility. Just as implied volatility is based on the Black-Scholes model, implied correlation also has to be based on a standard and widely accepted model that expresses the price of a CDO tranche in terms of a correlation parameter, in addition to other known inputs. In the CDO market, a single-factor Gaussian copula default time model with homogeneous correlation assumption is commonly used as the standard model for this purpose.

The implied correlation parameter is extracted from a given tranche price, assuming that all other inputs to the model are known.

If implied correlation is indeed a measure of correlation among underlying credits, then one could compute it from different CDO tranches with the same underlying reference portfolio and get identical numbers. However, in practice, this is not the case. The implied correlations computed from observed prices of tranches on a given reference portfolio typically exhibit a U-shaped pattern, which is commonly referred to as the "correlation smile," analogous to the "volatility smile." Exhibit 2 shows a typical correlation smile pattern obtained from the market prices of standardized tranches on the CDX.NA.IG index. Moreover, the implied correlation smile is quite pervasive, and the pattern seen in Exhibit 2 is exceedingly common, with implied correlations being lowest for the [3%, 7%] mezzanine tranche and highest for the [15%, 30%] senior tranche. Such observations suggest that the correlation smile is a fairly persistent phenomenon.


Interpretation of Correlation Smile

If the correlation curve is translated back to tranche prices, it is easy to infer the correlation smile as suggesting relative value opportunities across tranches. For instance,

EXHIBIT 2

Implied Correlation Smile for the CDX.NA.IG Index

The exhibit shows implied correlations obtained from market prices of standardized tranches on the CDX.NA.IG index as of September 22, 2005.



the correlation smile suggests a higher than average implied correlation for senior-most tranches. This can be interpreted as senior tranche spreads being too high, i.e., senior tranches being cheap. In other words, when correlations among credits increase, the likelihood of many defaults occurring together increases, thereby increasing the likelihood of an extreme loss on the portfolio. Since senior tranches suffer losses only when the underlying portfolio has an extreme loss, the likelihood of senior tranches suffering a loss increases with a rise in correlations. This leads to a widening in their spreads. Thus, if the tranche pricing model is correctly specified, one can translate a higher than average implied correlation as suggesting too large a spread for the senior tranche.

Spreads on equity tranches, on the other hand, decrease with a rise in correlation across assets. When correlations increase, the likelihoods of both very large losses and very small losses increase. Equity, being the first-loss tranche, gains from a higher likelihood of very small (including zero) losses. Since the loss on the equity tranche is capped, it does not suffer in the same proportion from a higher likelihood of very large losses. As a result, the equity tranche overall becomes safer as correlations increase and its spread declines with rising correlations.⁵ A typical correlation smile may suggest that the equity tranche has a higher than average implied correlation and is thus overpriced or rich.

Such interpretations assume that the standard Gaussian model is an accurate description of the real behavior of underlying credits in the reference portfolios. As we outline in the next section, the standard Gaussian model has strong simplifying assumptions built into it. It is possible that the observed smile patterns arise, in substantial part, if not completely, from the fact that many of the assumptions in the standard Gaussian model do not reflect empirical reality. In the next section, we explore specific assumptions of this model and examine whether the commonly observed correlation smile is an artifact of model misspecifications.

STANDARD GAUSSIAN MODEL AND IMPLIED CORRELATIONS

We first review the industry standard Gaussian model for CDO tranche pricing. Then, we look at the impact of certain assumptions of this model on implied correlations. Finally, we examine the performance of this model for pricing a given CDO tranche through time.

The Standard Gaussian Model

The standard Gaussian model is a one-factor Gaussian copula default time model introduced by Li [2000]. It is the industry standard model that is used to price CDO tranches and to compute implied correlations due to its computational efficiency and easy intuitive appeal. The model first computes the loss distribution of the collateral portfolio by taking into account the default probability term structures of the underlying assets and their correlations and assumptions about recovery in the instance of default. The various tranches are then priced according to this loss distribution.

The standard Gaussian model makes the following assumptions:⁶

1. The marginal distributions of asset returns of the names in the reference portfolio are Gaussian.
2. The dependence structure of asset returns is also given by a Gaussian copula.
3. The correlations among asset returns are driven by a single common factor.
4. All pair-wise correlations among asset returns are identical (homogeneous correlation assumption). This assumption ensures that there is only one correlation parameter in the model, and thus one can solve for a single implied correlation number from a given tranche price.
5. All the credits in the reference portfolio have identical spreads (homogeneous spread assumption) equal to the average spread of the portfolio.
6. Recovery rates on underlying credits are homogeneous across credits and are constant through time.

Assume that there are n names in the reference portfolio. This model defines a latent variable X_i , $i = 1, 2, \dots, n$, for each name, driven by a one-factor model, as follows:

$$X_i = a_i Z + \sqrt{1 - a_i^2} \varepsilon_i \quad (1)$$

Here Z is the common risk factor, ε_i are the identically distributed idiosyncratic shocks independent of Z and of each other, and Z and ε_i have standard normal distributions. The above structure implies that the correlation between X_i and X_j is $a_i a_j$, and that X_i , being a linear combination of standard normal variables, also has a standard

normal distribution. It can be shown that X_i can be interpreted as asset returns of individual firms that have a one-factor structure. Under the assumptions of the standard Gaussian model, all assets have the same pair-wise correlation, i.e., $a_i = a_j = a$, for all i, j ($a \geq 0$).

The above dependence structure is used to compute a distribution of default times as follows. Let t_i be the default time of the i th asset and $P_i(t)$ be the cumulative risk-neutral probability of asset i defaulting before time t , i.e., the probability that $t_i < t$. The $P_i(t)$ values for all assets are determined from the known or given CDS spreads where the risk-neutral default intensity for the i th asset, λ_i , is the CDS spread⁷ per unit of loss given default (LGD).⁸ Accordingly, the risk-neutral probability of default of asset i by time t is, $P_i(t) = 1 - \exp(-\lambda_i t)$. Note that as the spreads and LGDs are assumed to be identical across all the firms in the standard Gaussian model, all assets have the same default intensity and risk-neutral default probability term structure, i.e., $\lambda_i = \lambda$ and $P_i(t) = P(t)$ for all i . To compute the default time, the normal variates X_i are transformed to uniform variates U_i such that $U_i = \Phi(X_i)$, where Φ is the standard cumulative normal distribution function. If U_i is greater than $P_i(t)$, then asset i does not default by time t . If $P_i(t_1) < U_i < P_i(t_2)$, then asset i defaults between time t_1 and t_2 . Thus, low values of X_i , especially high negative values, lead to more defaults. The standard Gaussian model thus allows us to compute a risk-neutral distribution of default times and a risk-neutral distribution of portfolio losses using CDS spreads, LGD, and a correlation parameter. These distributions are then used to value the various CDO tranches using standard risk-neutral valuation methodology. Equivalently, if a tranche value is given, one can back-solve for a correlation parameter.

Model Misspecifications and Implied Correlation Smile

In this section, we first describe our methodology to explore the impact of various assumptions in the standard Gaussian model on implied correlation patterns. Then we present our findings related to misspecifications arising from the standard model assumptions by considering them one at a time.

Methodology. We start with assuming a reference portfolio that resembles an actively traded CDS index and focus on pricing CDO tranches referencing this index. The tranche structure that we choose is similar to the

standardized tranche structure that is commonly seen in the marketplace.

We compute the fair values of these tranches using a price-generating model. The price-generating model is assumed to be the “true model” that drives the fair tranche values. In other words, it is the model that takes the reference portfolio and tranche characteristics as given and generates the fair values of tranches. Next, we back-solve for the implied correlation parameter from the fair values of tranches, but this time using the standard Gaussian model, as is the common industry practice.

When the price-generating model is the same as the model used to extract correlations (i.e., both are standard Gaussian models), we should not observe any smile pattern in the computed implied correlations, that is, we should recover the same correlation parameter value for all tranches. We allow the price-generating model to deviate from the standard Gaussian model to reflect various empirical regularities and yet continue to use the standard Gaussian model to compute implied correlations, consistent with market practice. We then examine the resulting implied correlation curves. The following deviations from the standard Gaussian model are considered, each motivated by empirical evidence: 1) a fat-tailed return distribution rather than a Gaussian distribution, 2) heterogeneous spreads on underlying credits rather than homogenous spreads, 3) heterogeneous pair-wise correlations among underlying credits rather than homogenous correlations, and 4) recovery rates correlated with default probabilities rather than constant and homogenous recovery rates. We study each deviation separately to isolate its marginal impact on the implied correlations curve.

We assume a reference portfolio of 125 credit default swaps, each with a tenure of five years and equal weight, similar to the CDX.NA.IG index. The tranches on this portfolio are assumed to have breakpoints at 3%, 7%, 10%, 15%, and 30%, again similar to the standardized tranches on the CDX.NA.IG index. In our baseline standard Gaussian model, we assume that all credits in the reference portfolio have identical spread of 49 bps, which is the spread on the CDX.NA.IG Series 5 index on September 22, 2005, just after it became an on-the-run series.⁹ We also assume the loss given default to be constant at 50%.

We employ a Monte Carlo simulation framework for computing the tranche fair values. Each price-generating model that we use has a one-factor structure, thus correlated asset returns for individual names in the reference

portfolio are simulated using a one-factor model. If the systematic factor as well as idiosyncratic shocks are normally distributed, then the distribution of individual asset returns X_i is itself normal and can be computed analytically. If any of the component distributions is nonnormal, we compute the distribution of individual asset returns X_i numerically. The timings of individual asset defaults are determined in a manner similar to as described in our discussion of the standard Gaussian model—i.e., if U_i is greater than $P_i(t)$, then asset i does not default by time t ; however, if $P_i(t_1) < U_i < P_i(t_2)$, then asset i defaults between times t_1 and t_2 . In cases where the systematic factor as well as idiosyncratic shocks are normally distributed, we use the standard normal distribution function to transform asset returns to a uniform (0, 1) space. In cases where nonnormal distributions are involved, the transformation is done using the numerically computed distribution functions. Numerical estimations of nonnormal distributions are done using 10 million data points.

We follow market convention for quoting the fair values of tranches. The fair values of all tranches except the equity tranche are expressed in terms of the fixed annual spread earned on the notional balance of the tranche. Fair spread on the equity tranche is decomposed into two components—a fixed annual spread of 5% that is earned on the notional balance of the equity tranche, called the running spread, and the remainder that is paid in advance as an upfront fee. In Exhibit 3, we report the fair values of standard CDO tranches obtained using the standard Gaussian model.

In the following subsections, we consider various deviations from the standard Gaussian model, one at a time, and investigate the resulting implied correlation patterns.

Fat-tailed return distributions. It has been well documented that the observed returns on financial assets have fat-tailed empirical distributions (e.g., see Fama [1965]). The standard Gaussian model, however, assumes that asset returns are normally distributed. To examine the impact of this assumption on implied correlations when the true asset return distribution has fat tails, we assume a price-generating model that incorporates a t -distribution for asset returns. We consider t -distributions with degrees of freedom parameter ν set to 4, 7, and 12 to achieve varying degrees of fat tails. As shown in Exhibit 4, a t -distribution is a convenient way to generate a fat-tailed return distribution. It nests the normal distribution as a special case when the degrees of freedom parameter is large (typically $\nu > 30$). We work with three different cases, 1) both the systematic and the idiosyncratic parts of asset returns are t -distributed (referred to as the double t -distribution model); 2) the systematic part is t -distributed while the idiosyncratic part is normal; and 3) the systematic part is normally distributed and the idiosyncratic part is t -distributed. The remaining assumptions in our price-generating model are the same as in the standard Gaussian model. The distribution of the total asset returns is determined numerically as described in the previous subsection. As before, the correlation between X_i and X_j is $a_i a_j$ and under the homogeneous correlation assumption, $a_i = a$, for all i .

EXHIBIT 3

Fair Tranche Values Using the Standard Gaussian Model


This exhibit shows fair values using the standard Gaussian model. For this analysis, we use standardized tranches on a reference portfolio of 125 credit default swaps. All credits are assumed to have an identical spread of 49 bps and a recovery value of 50%. Pair-wise correlations are assumed to be homogenous. Their assumed values are shown in the first column. Fair spreads are given in basis points, except the equity (0%–3%) tranche, which is quoted in upfront points assuming a 5% running spread.

Correlation	0%–3% Tranche (Upfront)	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	53.36%	77.37	0.00	0.00	0.00
0.05	47.32%	173.34	6.07	0.18	0.00
0.10	42.27%	234.88	27.31	3.18	0.05
0.15	37.80%	272.82	54.24	10.62	0.42
0.20	33.63%	295.31	78.90	21.93	1.66
0.25	30.23%	317.67	104.45	36.02	3.93
0.30	26.71%	322.59	122.30	48.88	7.57

EXHIBIT 4

Probability Densities of a Standard Normal Distribution and a t -Distributions

The exhibit compares the tails of a standard normal distribution and a t -distribution with four degrees of freedom.



We estimate fair prices for our standard CDO tranches using the above price-generating model. For each asset return correlation parameter, a new empirical distribution is generated. The fair tranche values for the double t -distribution case are reported in Exhibit 5. Both the t -distributions are assumed to have the same degrees-of-freedom parameter. Panel A and Panel B show the results corresponding to four and seven degrees of freedom, respectively.¹⁰ In Exhibit 5, Panel C and Panel D show the implied correlations computed from the fair values in Panel A and Panel B, respectively, using the standard Gaussian model.

Exhibit 6, Panel A, shows these implied correlations for $\nu = 4$ case. There is a clear U-shaped pattern in these implied correlations. Thus the evidence suggests that using the standard Gaussian model, when in reality asset returns may be fat-tailed, can lead to a correlation smile pattern similar to the one derived from market tranche prices. Furthermore, the smile patterns are more pronounced when the underlying correlation across asset returns are higher. In Exhibit 6, Panel B, the underlying asset correlation is fixed at 30% and different curves correspond to different degrees-of-freedom parameters in the double t -distribution. As we increase the degrees-of-freedom parameter from 4 to 7, and then to 12, the implied correlation curve flattens. This result is expected since the t -distribution

approaches the standard normal distribution as degrees of freedom increase. The fact that we see a significant correlation smile even for mildly fat-tailed returns corresponding to $\nu = 12$, suggests that the tail behavior of returns has a substantial impact on the correlation smile. Thus, the fat-tailed nature of returns is a key empirical regularity that an alternative to the standard Gaussian model should strive to capture.

The above analysis brings out another important finding. Apart from a non-flat shape of the implied correlation curve, the levels of implied correlations can be very different from the actual correlations. A model misspecification, such as assuming normally distributed returns when they are fat tailed, drives a substantial wedge between true correlations and the implied correlations. In Exhibit 6, Panel A, each curve represents a particular level of true correlation and the corresponding implied correlations can be read off the vertical axis. For example, when the true correlation is 30%, the implied correlation ranges from 7% to 44%, depending on the tranche.

This illustration brings out the pitfalls of interpreting implied correlation as true correlation, particularly for the purpose of relative value assessments. As an example, suppose the historically observed average level of correlation for investment-grade names is 20% and an investor expects these levels to remain the same in the future. Under this assumption, when we examine the implied correlation on the equity tranche in Exhibit 5, Panel C, it appears that the equity tranche is priced at an implied correlation level of 15% instead of 20%. This may suggest that the equity tranche spread is high relative to historical levels (or, equivalently, the equity tranche value is lower relative to its historical levels). If so, one may infer that equity is cheap, i.e., selling protection (or going long on the equity tranche) would be a profitable strategy since equity would become less risky and its spread would decrease as correlations move up to their historical levels. An investor using this relative value strategy would be making an erroneous correlation trade, because in our example actual correlation is 20% and equity is, in fact, fairly priced by the double t -distribution price-generating model.

Next, to explore the relative impact of fat tails in the distribution of the systematic factor versus the idiosyncratic factor, we examine cases where only one of them is t -distributed while the other is normally distributed.¹¹ Exhibit 7 presents our findings for these cases

EXHIBIT 5

Fair Tranche Values and Implied Correlations: A Double t -Distribution Model

This exhibit shows tranche fair values and implied correlations using fat-tailed asset return distributions, a feature that is not incorporated in the standard Gaussian model. For this analysis, we use standardized tranches on a reference portfolio of 125 credit default swaps. All credits are assumed to have a spread of 49 bps and a recovery value of 50%. Pair-wise correlations are assumed to be homogenous. Their assumed values are shown in the first column. Panel A and Panel B show tranche fair-values generated with a double t -distribution model with degrees of freedom parameter, ν , set equal to four and seven, respectively. Panel C and Panel D show the implied correlations computed using the standard Gaussian model from the fair tranche values shown in Panel A and Panel B, respectively. Fair spreads are given in basis points, except for the equity (0%–3%) tranche, which is quoted in upfront points assuming a 5% running spread.

Panel A: Fair values of tranches using a double t -distribution ($\nu = 4$)

Correlation	0%–3% Tranche (Upfront)	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	53.24%	77.67	0.00	0.00	0.00
0.05	49.12%	126.92	8.85	3.10	0.89
0.10	45.41%	158.52	24.01	9.77	3.11
0.15	41.45%	176.60	40.76	18.95	6.70
0.20	37.66%	187.72	53.46	26.62	9.83
0.25	34.66%	199.81	68.44	36.39	14.97
0.30	31.21%	201.93	78.34	45.14	20.13

Panel B: Fair values of tranches using a double t -distribution ($\nu = 7$)

Correlation	0%–3% Tranche (Upfront)	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	53.39%	77.56	0.00	0.00	0.00
0.05	48.89%	147.20	7.37	1.10	0.11
0.10	44.11%	184.31	23.77	6.23	1.01
0.15	40.47%	215.38	44.80	15.03	2.96
0.20	36.87%	233.75	62.55	24.64	5.76
0.25	33.42%	249.54	81.34	36.31	9.97
0.30	29.49%	248.08	92.71	44.78	13.76

Panel C: Implied correlations computed from the fair values in Panel A above

Correlation	0%–3% Tranche	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	0.000	0.000	0.000	0.000	0.000
0.05	0.034	0.024	0.060	0.099	0.171
0.10	0.067	0.041	0.094	0.145	0.235
0.15	0.108	0.053	0.127	0.186	0.292
0.20	0.152	0.062	0.149	0.217	0.325
0.25	0.188	0.070	0.178	0.252	0.383
0.30	0.234	0.072	0.199	0.288	0.422


Panel D: Implied correlations computed from the fair values in Panel B above

Correlation	0%–3% Tranche	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	0.000	0.000	0.000	0.000	0.000
0.05	0.035	0.035	0.055	0.077	0.111
0.10	0.080	0.059	0.094	0.125	0.176
0.15	0.117	0.082	0.135	0.170	0.232
0.20	0.161	0.099	0.168	0.209	0.280
0.25	0.204	0.119	0.204	0.251	0.326
0.30	0.260	0.117	0.230	0.286	0.367

EXHIBIT 6

Implied Correlations Computed from Fair Tranche Prices Where Asset Returns Follow a Double t -Distribution

Panel A shows implied correlations for a double t -distribution model with four degrees of freedom. Panel B shows the effect of increasing the degrees of freedom of t -distributions on the implied correlation smile in the models with a double t -distribution, where the actual underlying asset correlation is 30%. All credits are assumed to have an identical spread of 49 bps.



using t -distributions with $\nu = 4$.¹² In Exhibit 8, the implied correlation results from this case are compared against the model with a double t -distribution, assuming a 30% homogeneous correlation across the underlying asset returns. As can be observed in the exhibit, the tail behavior of the systematic factor has a far greater impact on the shape of the implied correlation curve than the tail behavior of the idiosyncratic returns.

Exhibit 7, Panel A, shows the fair tranche values with t -distributed systematic factor and normally distributed idiosyncratic factor. The corresponding values with double t -distribution can be seen in Exhibit 5, Panel A. We observe that, in comparison to the results with double t -distribution reported in Exhibit 5, as correlation increases,

the riskiness of both the equity tranche and senior tranches changes considerably. Let's focus on the behavior of the values of the equity tranche (0%–3% tranche) and a senior tranche, such as the 10%–15% tranche, as the underlying correlation increases from 0% to 30%. In Exhibit 5, Panel A, using the double t -distributed price-generating model, as the underlying asset correlation increases from 0% to 30%, the equity tranche upfront price declines from 53 to 31 bps while the 10%–15% tranche spread increases from 0 to 45 bps. In Exhibit 7, Panel A, the equity upfront price declines substantially from 54 to 16 bps and the 10%–15% tranche spread increases drastically from 0 to 83 bps. As a contrast, we can look at Exhibit 7, Panel B, which reports the tranche fair values when the systematic

EXHIBIT 7

Fair Tranche Values and Implied Correlations: t and Gaussian Distribution Mixture Models

This exhibit shows tranche fair values and implied correlations using a model where asset returns are generated using a mixture of t and normal distributions. For this analysis, we use standardized tranches on a reference portfolio of 125 credit default swaps. All credits are assumed to have a spread of 49 bps and a recovery value of 50%. Pair-wise correlations are assumed to be homogenous. Their assumed values are shown in the first column. Panel A shows tranche fair values for the case where the underlying asset returns have a t -distributed (with four degrees of freedom) systematic component and a normally distributed idiosyncratic component. Panel B shows tranche fair values where asset returns have a normally distributed systematic component and t -distributed (with four degrees of freedom) idiosyncratic components. Panels C and Panel D show the implied correlations computed using the standard Gaussian model from the fair tranche values shown in Panel A and Panel B, respectively. Fair spreads are given in basis points, except the equity (0%–3%) tranche, which is quoted in upfront points assuming a 5% running spread.

Panel A: Fair values of tranches using a t -distribution ($\nu = 4$) for systematic factor and standard normal distribution for idiosyncratic factor

Correlation	0%–3% Tranche (Upfront)	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	53.54%	77.81	0.02	0.00	0.00
0.05	42.57%	191.37	31.41	10.95	2.57
0.10	35.48%	238.91	67.39	29.13	8.50
0.15	29.06%	252.58	91.31	44.81	14.56
0.20	24.25%	265.27	111.58	60.53	22.28
0.25	19.07%	261.28	123.24	70.78	27.78
0.30	15.50%	262.70	135.53	82.86	36.19

Panel B: Fair values of tranches using a standard normal distribution for systematic factor and t -distribution ($\nu = 4$) for idiosyncratic factor

Correlation	0%–3% Tranche (Upfront)	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	53.38%	77.92	0.01	0.00	0.00
0.05	51.25%	109.83	0.41	0.01	0.00
0.10	49.19%	142.30	3.16	0.10	0.00
0.15	47.06%	170.55	9.94	0.89	0.01
0.20	45.07%	194.51	20.02	2.83	0.08
0.25	42.68%	212.74	32.31	6.31	0.37
0.30	40.15%	228.79	45.38	11.39	1.02

Panel C: Implied correlations computed from the fair values in Panel A above

Correlation	0%–3% Tranche	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	0.000	0.000	0.000	0.000	0.000
0.05	0.097	0.064	0.108	0.152	0.220
0.10	0.180	0.104	0.176	0.227	0.309
0.15	0.266	0.123	0.224	0.286	0.379
0.20	0.337	0.138	0.274	0.346	0.444
0.25	0.418	0.133	0.302	0.383	0.484
0.30	0.479	0.135	0.344	0.434	0.541

Panel D: Implied correlations computed from the fair values in Panel B above


Correlation	0%–3% Tranche	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	0.000	0.000	0.000	0.000	0.000
0.05	0.016	0.015	0.020	0.031	0.050*
0.10	0.033	0.032	0.040	0.045	0.050*
0.15	0.052	0.048	0.063	0.074	0.086
0.20	0.070	0.067	0.087	0.096	0.106
0.25	0.095	0.080	0.110	0.126	0.139
0.30	0.121	0.095	0.136	0.154	0.176

Note: *Unique values cannot be backed out since any correlation value between 0%–5% prices the 15%–30% tranche to 0.00 bps. We report the implied correlation as 5% because the deviation between implied correlation and actual correlation is minimal in this case.

EXHIBIT 8

Implied Correlations Computed from Fair Tranche Prices Where Asset Returns Are Derived from a Mixture of t and Normal Distributions

This exhibit shows a comparison between implied correlations from models where asset returns are generated from a mixture of t (with four degrees of freedom) and normal distributions to a double t -distribution model with four degrees of freedom. The actual underlying correlation is assumed to be 30%. All credits are assumed to have an identical spread of 49 bps.



factor is normally distributed and the idiosyncratic factor has a t -distribution with $\nu = 4$. In this case, as the underlying correlation increases, the change in the risk (and hence the value) of both the equity tranche and senior tranches is the least pronounced. When correlation increases from 0% to 30%, the equity upfront price declines from 53 to 40 bps while the 10%–15% senior tranche spread increases from 0 to just 11 bps.

We explain these rather interesting findings as follows: In a CDO, the process of structuring has the effect of distributing the systematic and idiosyncratic risks in the underlying portfolio of assets unevenly across the structured securities or tranches. While the junior tranches, especially the equity tranche, are affected by any idiosyncratic default in the reference pool, senior tranches experience loss stresses primarily in the event of systematic joint defaults. In other words, at the senior tranche levels, the idiosyncratic risk gets diversified away, and it is primarily the systematic risk that drives the performance of these securities. In the case where individual asset returns have a combination of a t -distributed systematic component and a normally distributed idiosyncratic component, such as those reported in Panel A of


Exhibit 7, the fatter tails of the t -distribution result in a greater weight on the systematic part of the asset return. Thus defaults are driven more due to systematic events than idiosyncratic events. This results in the equity tranche becoming relatively safer and the senior tranches relatively riskier. The reverse argument holds for the case where the systematic latent factor has a normal distribution and the idiosyncratic factor has a t -distribution. In this case, defaults in the underlying portfolio are driven mainly by idiosyncratic shocks, thereby making the 0%–3% tranche riskier and the senior tranches relatively safer.

Further insights into the behavior of the tranche prices can be obtained by examining the distributions of the number of defaults, as generated by different models. Exhibit 9 shows the (risk-neutral) number of default distributions over a five-year horizon, where the underlying asset return correlation is assumed to be 30%. As can be observed in Exhibit 9, equity is likely to be the least risky when the systematic factor is t -distributed and the idiosyncratic factor is normally distributed (due to a high likelihood of having no or very few defaults) and to be the most risky when the systematic factor is normally distributed and the idiosyncratic factor is t -distributed.

EXHIBIT 9

Risk-Neutral Default Distribution

This exhibit shows the risk-neutral default distributions where systematic and idiosyncratic latent factors of asset returns follow either a standard normal distribution or a t -distribution with four degrees of freedom. The underlying asset return correlation in each case is 30%. All credits are assumed to have an identical spread of 49 bps.



Conversely, senior tranches are the most risky when the systematic factor is t -distributed and the idiosyncratic factor is normally distributed (due to the extended fat tail of the default distribution) and the least risky when the systematic factor is normally distributed and the idiosyncratic factor is t -distributed.

The results presented in Exhibit 7 can be further explained with the following example. In this study, a 49 bps underlying CDS spread with 50% recovery translates to about a 5% (risk-neutral) default probability over a five-year horizon.¹³ For the 30% correlation case, this 5% default probability translates to an asset return threshold of 1) -1.8223 for the model where the systematic latent factor has a t -distribution and the idiosyncratic factor has a standard normal distribution, 2) -2.0004 for the model where the systematic latent factor has a standard normal distribution and the idiosyncratic factor has a t -distribution, and 3) -2.1861 for the double t -distribution model. The likelihoods of a systematic shock exceeding these default thresholds (i.e., being more negative) are 7.12%, 2.27%, and 4.71%, respectively.¹⁴ This is consistent with the notion that, in a relative sense, systematic defaults drive the results when the systematic factor has a t -distribution and the

idiosyncratic factor has a standard normal distribution, whereas idiosyncratic defaults drive the results when the systematic factor is normally distributed and the idiosyncratic factor has a t -distribution.¹⁵ As a result, the implied correlation smile is more pronounced when only the systematic factor has a t -distribution and less pronounced when only the idiosyncratic factor has t -distribution, with the double t -distribution results falling in between.

Overall, our results suggest that the assumption of a normally distributed latent factor in the standard Gaussian model when the actual systematic factor has fat tails contributes substantially to a correlation smile. In contrast, the assumption of normally distributed idiosyncratic returns when actual idiosyncratic returns are fat tailed has a far smaller influence in generating a correlation smile pattern. Later in this section, we will examine the relative importance of fat-tailed return distributions for implied correlation smile patterns compared with other misspecifications of the Gaussian copula model that are considered in this study.

Heterogeneous correlations. A critical assumption of the standard Gaussian model is that of homogeneous correlations, i.e., all pair-wise correlations across asset returns

in the reference portfolio are considered to be identical. This assumption is critical if one has to back out a single correlation number from an observed tranche price. In reality, however, one can expect a significant heterogeneity in pair-wise correlations for the names in various CDS indices. For example, Exhibit 10 shows the distribution of KMV estimates of pair-wise asset return correlations among the names in CDX.NA.IG index on September 22, 2005.¹⁶ These pair-wise correlations have a wide range from 0.12 to 0.63, thus confirming that the assumption of homogenous correlations is a strong one.

How crucial is the assumption of a homogenous dependence structure when true underlying correlations are heterogeneous? To examine this question, we carry out the following experiment. We randomly assign factor loadings a_i in the model represented in Equation (1) to each of the 125 names in our reference portfolio. These factor loadings are drawn from a uniform distribution over the interval $[0, 1]$. Assigning a different (random) factor loading or sensitivity to the systematic factor Z results in a well-behaved heterogeneous dependence structure. All the remaining assumptions of the standard Gaussian model are retained unchanged. Using this price-generating model, we compute the fair values for our standardized tranches and then invert them to compute the implied correlations using the standard Gaussian model.¹⁷


We repeat this analysis four times, each time generating a different (though random) heterogeneous correlation structure as a robustness check. Note that in each of these scenarios, the generated average correlations are close to 0.25 since the mean individual asset return systematic factor loading, a_i , is 0.5 (being drawn from a uniform distribution between zero and one). While acknowledging that real world correlations might be indeed very different from our randomly generated dependence structures, these analyses nonetheless allow us to study the impact of different heterogeneous correlation structures on implied correlations.

The results for our five scenarios are presented in Exhibit 11. Panel A gives fair values of tranches obtained in five different random heterogeneous correlations scenarios. In Panel B, for comparison purposes, we show fair values of the same tranches as generated by the standard Gaussian model with a homogenous correlation value equal to the average of all pair-wise correlations from the corresponding randomly generated heterogeneous correlations. Finally, we back out implied correlations from

EXHIBIT 10

Pair-Wise Asset Return Correlation Estimates

This exhibit shows the distribution of KMV estimates of pair-wise asset correlations among the names in CDX.NA.IG index on September 22, 2005.



the fair values given in Panel A. These results are reported in Panel C and plotted in Exhibit 12, Panel A. The results show that these implied correlations also display a U-shaped pattern similar to that observed empirically. Thus an implied correlation smile can also arise due to the homogeneous correlation assumption of the standard Gaussian model.

The evidence presented so far is based on a set of randomly generated correlation structures. We next conduct a similar analysis using empirically estimated pair-wise asset return correlations for our reference portfolio of 125 names as provided by KMV.¹⁸ The reference portfolio comprises the same names as the CDX.NA.IG Series 5.

The results are reported in Exhibit 13. The first row gives the fair values of tranches obtained from the heterogeneous correlation model. The second row shows the fair values under the standard Gaussian model with homogenous correlation equal to the average of pair-wise KMV correlations for our reference portfolio (0.347). The third row reports the implied correlations that are backed out from tranche fair values given in the first row. The implied correlations obtained are also plotted in Panel B of Exhibit 12.

As shown in Exhibit 13 and Panel B of Exhibit 12, the shape of the implied correlation smile using KMV estimated correlations is similar to the ones obtained from a random heterogeneous correlation model given in Exhibit 12, Panel A. In both cases, it is primarily the mezzanine tranche that has a large effect on the correlation smile. As can be observed in Exhibit 11, Panel C, for an average input correlation of 30%, equity and senior tranches produce implied correlations that are close to the average input correlation, whereas the implied

EXHIBIT 11

Fair Tranche Values and Implied Correlations: A Gaussian Copula Model with Heterogeneous Input Correlations

This exhibit shows tranche fair values and the corresponding implied correlations considering the heterogeneous nature of pair-wise asset correlations, a feature that is not incorporated in the standard Gaussian model. For this analysis, we use standardized tranches on a reference portfolio of 125 credit default swaps. All credits are assumed to have a spread of 49 bps and a recovery value of 50%. Heterogeneous correlations are created by assigning a random value between zero and one to the correlation between each credit and the systematic factor. The rest of the assumptions of the standard Gaussian model are retained. Panel A shows tranche fair values generated by considering heterogeneous pair-wise correlations. The average values of pair-wise correlations are shown in the first column. For comparison purposes, Panel B shows fair values of the same tranches using the standard Gaussian model with homogenous correlation equal to the average of pair-wise heterogeneous correlations. Panel C shows the implied correlations computed from the fair tranche values shown in Panel A using a standard Gaussian model with homogenous correlation equal to the average pair-wise heterogeneous correlations. Fair tranche spreads are given in basis points, except for the equity (0%–3%) tranche, which is quoted in upfront points assuming a 5% running spread.

Panel A: Fair values of tranches for a heterogeneous correlation structure in a Gaussian copula model

Average Correlation	0%–3% Tranche (Upfront)	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.223	31.2%	241	108	57	8
0.239	30.1%	240	112	62	9
0.246	29.4%	241	116	65	11
0.269	28.3%	258	123	68	13
0.300	26.1%	258	126	73	16

Panel B: Fair values of tranches for the standard Gaussian copula model

Constant Correlation	0%–3% Tranche (Upfront)	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.223	31.6%	297	84	26	3
0.239	30.7%	306	95	31	3
0.246	30.5%	318	102	36	4
0.269	28.8%	317	110	41	5
0.300	27.3%	336	126	52	7

Panel C: Implied correlations computed from the fair values in Panel A above

Input Average Correlation	0%–3% Tranche	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.223	0.236	0.107	0.262	0.327	0.302
0.239	0.251	0.106	0.267	0.349	0.318
0.246	0.260	0.107	0.280	0.359	0.347
0.269	0.277	0.130	0.301	0.374	0.363
0.300	0.310	0.130	0.310	0.395	0.389



correlation from the mezzanine tranche deviates considerably from the input number. A similar pattern is observed in Exhibit 13 using KMV correlation estimates. With an average input correlation of 34.7%, it is mainly the mezzanine tranche that leads to the correlation smile. The deviation of implied correlation from the actual correlation estimate is minimal for other tranches. Thus, while both analyses confirm that the homogenous correlation assumption of the standard Gaussian model

is another likely explanation for the observance of a correlation smile in tranche prices, there is clear indication that the effect is driven more by the mezzanine tranche than the other tranches. This is primarily due to the location of the mezzanine tranche in the capital structure. This tranche becomes relatively correlation insensitive once correlation increases beyond a certain threshold. We will discuss the correlation insensitivity of the mezzanine tranche in detail later in this section.

EXHIBIT 12

Implied Correlations Computed from Fair Tranche Prices with Heterogeneous Pair-Wise Correlations

This exhibit shows implied correlations backed out using the standard Gaussian model where homogenous correlation is equal to the average of heterogeneous pair-wise correlations. In Panel A, heterogeneous pair-wise correlations are assigned randomly. In Panel B, heterogeneous pair-wise correlations are from the global correlation model estimated by KMV.



Heterogeneous spreads (default probabilities). The standard Gaussian model assumes that all the underlying credits in the reference portfolio have the same spread, often set equal to the average spread on the portfolio. Exhibit 14 shows the distribution of spreads for the CDX.NA.IG index constituents, again as of September 22, 2005, our representative analysis date. The spreads are rather heterogeneous and range from a few basis points to more than 300. Here, we examine the impact of the spread homogeneity assumption on the implied correlation curve.

We first estimate the fair values of tranches using the standard Gaussian model except that we allow each name in our reference portfolio to have a distinct CDS spread. Observed CDS spreads for the 125 CDX.NA.IG index components on September 22, 2005, are used for this purpose. The average of these spread values is 49 bps, a value that is used in the homogenous spread model when solving for implied correlations. All other assumptions of the standard Gaussian model are retained to isolate the impact of heterogeneity in spreads.

EXHIBIT 13

Fair Tranche Values and Implied Correlations: A Heterogeneous Correlation Structure with Empirically Estimated Correlations

This exhibit shows tranche fair values and the corresponding implied correlations considering the heterogeneous nature of pair-wise asset correlations, a feature that is not incorporated in the standard Gaussian model. For this analysis, we use standardized tranches on a reference portfolio of 125 credit default swaps. All credits are assumed to have an identical spread of 49 bps and a recovery value of 50%. The heterogeneous correlations are obtained from the global correlation model of KMV. The average of these correlations is 0.347. Row A shows tranche fair values generated by assuming heterogeneous pair-wise correlations. The remaining assumptions of the standard Gaussian model are retained. For comparison purposes, Row B shows fair values of the same tranches using the standard Gaussian model with homogenous correlation equal to average correlation of 0.347. Row C shows the implied correlations computed from the fair tranche values shown in Row A using a standard Gaussian copula model with homogenous correlation of 0.347. Tranche fair spreads are given in basis points, except for the equity (0%–3%) tranche, which is quoted in upfront points assuming a 5% running spread.

	0%–3% Tranche	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
A. Fair values with heterogeneous correlations	23.4% upfront	315	137	63	13
B. Fair values with constant correlations	23.7% upfront	330	138	61	11
C. Implied Correlations computed from fair values in Row A	0.347	0.257	0.347	0.354	0.356

EXHIBIT 14

Distribution of CDS Spreads for the CDX.NA.IG Index

This exhibit shows the distribution of five-year CDS spreads for the constituents of the CDX.NA.IG index on September 22, 2005. The spreads have substantial heterogeneity.




Exhibit 15, Panel A, shows tranche fair values when spreads are heterogeneous. The corresponding fair values with a homogenous spread of 49 bps were shown in Exhibit 3. We translate the fair values in Exhibit 15, Panel A, to implied correlations using the standard Gaussian model with homogenous spread of 49 bps. The results are shown in Exhibit 15, Panel B, as well as in Exhibit 16. As can be observed, the homogenous spread assumption in the standard Gaussian model also induces a correlation smile, although once again the smile appears to be driven mainly by the mezzanine tranche and becomes pronounced only as the “true,” or

input, correlations increase beyond small values. This finding is again consistent with our earlier argument about the general correlation insensitivity of the mezzanine tranche at medium to high levels of input correlation. Overall, the evidence presented here suggests that assuming homogenous spreads when they have a substantial heterogeneity can also lead to a correlation smile. However, the correlation smiles generated by the heterogeneous spread assumption are far milder than those generated by fat-tailed return distributions and heterogeneous correlation assumptions.

Correlated recovery rate and probability of default.

Another major assumption in the standard Gaussian model is that recovery rates upon default are constant. However, there is growing evidence in the literature that recovery rates are stochastic and correlated with the state of the economy.¹⁹ For example, Altman, Resti, and Sironi [2004] find a strong negative correlation between default rates and recovery rates in the U.S. corporate bond market. Hull, Predescu, and White [2005] develop a model for CDO pricing that incorporates the relationship between recovery rates and default rates.

In this section, we investigate implied correlation patterns when recoveries are stochastic and correlated with the systematic factor—our proxy for the state of the economy. High default rates correspond to low realizations of our systematic factor (bad state of the economy). Therefore, the empirically observed negative correlation

between default rates and recovery rates translates into a positive correlation between our systematic factor and recovery rates.

We assume the following recovery model, wherein a latent variable r_i driving the recovery rate R_i on instrument i is decomposed into systematic and idiosyncratic components,

$$\begin{aligned} r_i &= \rho_{RR} Z + \sqrt{1 - \rho_{RR}^2} Y_i \\ R_i &= \Phi(r_i) \end{aligned} \quad (2)$$

where Φ is the standard normal cumulative distribution function; Z is the common systematic factor with a standard normal distribution that drives both the default and recovery processes; Y_i is the idiosyncratic recovery factor, also with a standard normal distribution; and ρ_{RR} is the

EXHIBIT 15

Fair Tranche Values and Implied Correlations: A Heterogeneous Spread Model

This exhibit shows fair values and implied correlations considering the heterogeneous nature of spreads on credits in the reference portfolio of a CDO, a feature that is not incorporated in the standard Gaussian model. For this analysis, we use standardized tranches on a reference portfolio of 125 credit default swaps. All credits are assumed to have a recovery value of 50%. We use CDS spreads on 125 CDX.NA.IG index components on September 22, 2005, for heterogeneous spreads. Panel A shows tranche fair values generated by considering heterogeneous spreads in a Gaussian model. The pair-wise correlations are assumed to be homogeneous and their values are shown in the first column. Panel B shows the implied correlations computed from the fair tranche values shown in Panel A using a standard Gaussian model with homogenous spread equal to the average spread on credits, which is 49 bps. Fair tranche spreads are given in basis points, except the equity (0%–3%) tranche, which is quoted in upfront points assuming a 5% running spread.

Panel A: Fair values of tranches for a heterogeneous spread model

Correlation	0%–3% Tranche (Upfront)	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	53.1%	71	0	0	0
0.05	47.3%	155	4	0	0
0.10	42.9%	215	21	2	0
0.15	38.8%	257	44	8	0
0.20	35.2%	284	68	17	1
0.25	31.8%	301	89	28	3
0.30	28.6%	313	108	40	6


Panel B: Implied correlations computed from the fair values in Panel A above

Input Correlation	0%–3% Tranche	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	0.002	0	0	0	0
0.05	0.049	0.039	0.044	0.04	0.05
0.10	0.092	0.08	0.088	0.089	0.092
0.15	0.137	0.128	0.132	0.134	0.134
0.20	0.18	0.172	0.178	0.181	0.181
0.25	0.228	0.219	0.221	0.224	0.231
0.30	0.274	0.244	0.263	0.265	0.278

EXHIBIT 16

Implied Correlations Computed from Fair Tranche Prices with Heterogeneous Spreads

This exhibit shows implied correlations backed out from the fair tranche values that are generated assuming heterogeneous spreads for the CDO reference portfolio. The spreads used are the actual spreads on the names in the CDX.NA.IG portfolio as of September 22, 2005. The average of these spreads is 49 bps. While the fair tranche values are generated assuming heterogeneous spreads, the implied correlations shown are computed by assuming a homogeneous spread of 49 bps. Different curves correspond to different correlation assumptions.



correlation between the systematic factor and the latent recovery variable r_i . This specification ensures that the average recovery is 50%, which is consistent with our fixed recovery of 50% assumption in the previous sections.

We illustrate the impact of stochastic recoveries by setting ρ_{RR} to be 0.29, following the empirical evidence of Altman, Brooks, Resti, and Sironi [2005] on the correlation between weighted average recovery of all corporate bond defaults between 1982 and 2001 and the annual GDP growth in the U.S.²⁰ The fair values of the tranches, generated using the above specification of stochastic recoveries, are given in Exhibit 17, Panel A. We extract implied correlations from these fair tranche values using the standard Gaussian model with a fixed recovery rate of 50%. These are shown in Exhibit 17, Panel B, and Exhibit 18.

These results show some interesting patterns in implied correlations. First of all, the correlation smile gets inverted when recovery is stochastic and positively correlated with the systematic factor. Moreover, the standard Gaussian model fails to converge to an implied correlation solution for the mezzanine tranche for cases where the underlying asset correlations are in excess of 10%. In these cases, mezzanine tranche spreads are substantially

high in the "true" model (with the systematic factor and recovery rates correlated) and the standard Gaussian model with a constant recovery rate of 50% is not able to match these prices for any level of constant pair-wise asset correlation beyond 10%.

Why are the results of this section so different from those of the previous sections in terms of the shape of the implied correlation curve? The fair value of a tranche is a function of the risk-neutral loss distribution of underlying collateral assets. When recoveries are fixed, as is the case in previous sections, the loss distribution for the collateral has the same shape as the distribution of number of defaults because the recovery (or LGD) simply scales the default distribution to a loss distribution. Moreover, in this case, the mean of the loss distribution is simply the average number of defaults multiplied by the LGD. When the recovery rate is positively correlated with the systematic factor, however, good states of the world that correspond to a small number of defaults also have high recovery values. Therefore, levels of loss are low in good states of the world relative to a constant recovery model. Conversely, bad states of the world that correspond to a large number of defaults have low recovery values. In these cases, the losses are higher relative to a constant

EXHIBIT 17

Fair Tranche Values and Implied Correlations: A Model with Recovery Rates Correlated with Systematic Factor

This exhibit shows the fair values of standardized tranches on a CDO and implied correlations when recovery rates on credits in the reference portfolio are stochastic and correlated with the systematic factor, a feature that is not incorporated in the standard Gaussian copula model. The systematic factor drives the asset returns. The correlation between the systematic factor and the recovery rate, ρ_{RR} , is considered as 0.29. For this analysis, we use standardized tranches on a reference portfolio of 125 credit default swaps. Panel A shows the fair values of CDO tranches. Panel B presents the implied correlations obtained from fair values given in Panel A using the standard Gaussian model with a recovery rate of 50%. Fair spreads are given in basis points, except the equity (0%–3%) tranche which is quoted in upfront points assuming a 5% running spread. When an implied correlation could not be backed out using the standard Gaussian model, it is represented as not available, N/A.

Panel A: Fair values of tranches for a model with recovery rates correlated with systematic factor

Correlation	0%–3% Tranche (Upfront)	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	50.94%	116.09	0.37	0.00	0.00
0.05	45.43%	270.62	34.73	4.39	0.08
0.10	41.04%	330.67	76.16	19.13	1.30
0.15	37.23%	367.14	112.23	37.15	4.38
0.20	33.59%	381.06	140.48	57.26	9.70
0.25	30.25%	391.20	162.48	74.58	16.30
0.30	26.87%	391.00	178.34	89.27	22.82

Panel B: Implied correlations computed from the fair values in Panel A above

Input Correlation	0%–3% Tranche	3%–7% Tranche	7%–10% Tranche	10%–15% Tranche	15%–30% Tranche
0.00	0.018	0.019	0.019	0.000	0.000
0.05	0.067	0.146	0.115	0.110	0.106
0.10	0.112	0.283	0.195	0.187	0.189
0.15	0.156	N/A	0.277	0.257	0.258
0.20	0.201	N/A	0.361	0.329	0.325
0.25	0.250	N/A	0.458	0.402	0.390
0.30	0.297	N/A	0.596	0.466	0.448

recovery model. The relation between expected recovery rates and default severity observed when ρ_{RR} is 29% is given in Exhibit 19. As shown in the exhibit, the positive correlation between recovery rate and the systematic factor leads to an increase in the expected loss.²¹ Therefore, all tranches except equity become riskier and their spreads increase.²² The standard Gaussian model with 50% fixed recovery can match the corresponding rise in the prices only by increasing the correlation. Since the mezzanine tranche is the least correlation-sensitive among others, a substantial rise in correlation is required to match the corresponding prices. This leads to the inverted implied correlation curves observed in this section.

Discussion. The overall evidence presented in the previous subsections suggests that deviations from the assumptions of the standard Gaussian model can lead to implied correlation smile patterns that are very similar to those implied from market prices of tranches. Given that the standard Gaussian model is widely used due to its

tractability, it is important to explore which of the assumptions are more severe in terms of their impact on the extent of the non-flat implied correlation patterns.²³ In this respect, for each of the assumptions of the standard Gaussian model that is relaxed in this study, we examine the mean absolute deviation of implied correlations from the actual ones as well as from the average implied correlation. These statistics are calculated across the tranches. The deviation of implied correlations from the actual input correlation indicates whether model-implied correlations are good estimates of actual underlying correlations. These are reported in Exhibit 20, Panel A. In Exhibit 20, Panel B, we report the mean absolute deviation of implied correlations from the average implied correlation that is calculated across the tranches for each correlation level considered. These results are useful to understand which assumption of the Gaussian model leads to non-flat patterns in implied correlations.

EXHIBIT 18

Implied Correlations with Stochastic Recovery

This exhibit shows implied correlation curves for the case where recovery is stochastic. The correlation between recovery rate and systematic factor, ρ_{RR} , is 0.29. Different curves correspond to different underlying asset return correlation assumptions




EXHIBIT 19

Recovery Rates for Different Default Percentages

This exhibit shows the expected recovery rate corresponding to varying levels of default severity 1) when there is positive correlation of 29% between recovery rates and systematic risk factor and 2) when recovery rate is constant and equal to 50%.




EXHIBIT 20

A Comparison of Various Misspecifications of the Standard Gaussian Model

This exhibit shows mean absolute deviations of implied correlations when certain assumptions of the standard Gaussian model are relaxed. Panel A reports mean absolute deviations from actual input correlation, i.e., a small mean absolute deviation indicates that the model implied correlations are good estimates of true underlying correlation. Panel B reports mean absolute deviations from average implied correlations, i.e., a small mean absolute deviation indicates a flat implied correlation pattern. All deviations are calculated using the implied correlations reported in Exhibits 7, 11, 13, 15, and 17. For the stochastic recovery model, if there are missing values in calculating these statistics, the calculations are done by omitting the tranche whose implied correlation value could not be estimated. For the t -distribution models, v represents the degrees of freedom. Statistics for the heterogeneous correlation model are shown separately due to the difference in the input correlations for this case. Bold values represent the maximum deviation observed among the considered cases, whereas italics shows the minimum deviation observed among the considered cases.

Panel A: Mean Absolute Deviation from Actual Correlation

Input Correlation	0%	5%	10%	15%	20%	25%	30%
Fat-tailed returns							
Double t ($v = 4$)	0.000	0.046	0.055	0.068	0.078	0.088	0.108
Double t ($v = 7$)	0.000	0.026	0.034	0.045	0.052	0.063	0.078
Systematic t - ($v = 4$), idiosyncratic normal	0.000	0.078	0.100	0.117	0.129	0.141	0.151
Systematic normal, idiosyncratic t ($v = 4$)—	0.000	0.034	0.071	0.085	0.113	0.140	0.163
Heterogeneous spreads	0.000	<i>0.006</i>	<i>0.012</i>	<i>0.017</i>	<i>0.022</i>	<i>0.025</i>	<i>0.035</i>
Recovery rate correlated with systematic factor	0.011	0.064	0.072	0.090	0.114	0.123	
Average Input Correlation	22.3%	23.9%	24.6%	26.9%	30.0%		34.7%
Heterogeneous correlations	0.070	0.072	0.080	0.076	0.075		0.021

Panel B: Mean Absolute Deviation from Average Implied Correlation

Input Correlation	0%	5%	10%	15%	20%	25%	30%
Fat-tailed returns							
Double t ($v = 4$)	0.000	0.046	0.059	0.069	0.072	0.083	0.090
Double t ($v = 7$)	0.000	0.025	0.035	0.043	0.049	0.054	0.063
Systematic t - ($v = 4$), idiosyncratic normal	0.000	0.046	0.055	0.066	0.081	0.101	0.118
Systematic normal, idiosyncratic t ($v = 4$)—	0.000	0.011	0.006	0.012	0.013	0.018	0.023
Heterogeneous spreads	0.001	<i>0.004</i>	<i>0.003</i>	<i>0.002</i>	<i>0.003</i>	<i>0.004</i>	<i>0.009</i>
Recovery rate correlated with systematic factor	0.009	0.018	0.037	0.041	0.052	0.063	0.079
Average Input Correlation	22.3%	23.9%	24.6%	26.9%	30.0%		34.7%
Heterogeneous correlations	0.060	0.064	0.07	0.068	0.071		0.030

The results in Exhibit 20, Panels A and B, suggest that the most severe assumption of the standard Gaussian model is that of normal distribution for both the systematic and idiosyncratic latent factors in modeling asset returns. When any of the latent factors are non-normal, the implied correlations obtained using the standard Gaussian model leads to a pronounced implied correlation

smile. Deviation of implied correlations from the true ones is particularly severe when the systematic latent factor has a t -distribution (with a small degrees-of-freedom parameter) and the idiosyncratic component has a normal distribution. The smile pattern is quite significant also when both the systematic and idiosyncratic factors have t -distributions with a small degrees-of-freedom parameter.

As discussed before, in these cases, defaults are driven more by the realization of systematic credit events, which is equivalent to assuming a higher correlation parameter in the standard Gaussian framework. Hence, the equity tranche becomes safer whereas senior tranches become riskier, resulting in a pronounced implied correlation curve. Thus the normality assumption of the systematic latent factor in the standard Gaussian model is the one that seems most problematic.

Among other assumptions of the standard Gaussian model, the least problematic one is that of heterogeneous spreads. As can be observed in Exhibit 20, Panel A, this assumption has the least mean absolute deviation of implied correlation from true model correlation in almost all cases considered. Also, the implied correlation pattern is fairly flat as observed in Exhibit 20, Panel B. Therefore, the homogenous spread assumption of the standard model can take the last priority if one has to be selective in relaxing the assumptions of the standard Gaussian model.

The assumptions of homogenous correlations and constant recovery fall somewhere in between the homogenous spreads and normally distributed asset returns assumptions in terms of both the severity of the implied correlation smile as well as the deviation of the implied correlation from actual underlying correlations. Therefore, one has to weigh the benefits of relaxing these assumptions against the increased cost in the form of model complexity.

Overall, our results suggest that the observed implied correlation smile patterns are, to a large extent, driven by the simplifying assumptions of the standard Gaussian model. The assumption of normally distributed asset returns instead of fat-tailed ones is the most crucial drawback of the model. Therefore, an alternative model that concentrates on relaxing this assumption seems to be the most fruitful direction for addressing the implied correlation smile patterns observed in the market.

Another important finding observed throughout this study is that in most cases, the deviation of the implied correlation from "true" correlation is at its maximum for the mezzanine tranche (3%–7% tranche). We explain this finding as follows. For the investment-grade index with an average spread of 49 bps and recovery of 50%, the risk-neutral expected number of defaults in the collateral portfolio over a horizon of five years is about 5%. Since the mezzanine tranche covers portfolio losses in the range of 3%–7%, the expected loss on the portfolio lies in the mezzanine tranche interval. Changes in the correlation struc-

ture have no impact on the first moment of the default distribution, which is the expected number of defaults. These changes only affect the second moment or the spread of the default distribution. Therefore, the mezzanine tranche, due to its location in the capital structure, becomes relatively correlation-insensitive for a range of correlation values. With an increase in correlation, senior tranches become riskier due to the added weight in the negative tail of the underlying portfolio loss distribution, whereas the equity tranche becomes less risky due to the added weight in the positive tail of the loss distribution. This behavior is also referred to as senior tranches being short on correlation and the equity tranche being long on correlation. Therefore, as one moves from equity to senior tranches, there must exist a point in the CDO capital structure where the effect of correlation is reduced, and typically the mezzanine tranche displays this characteristic.


Exhibit 21 plots the fair values of equity, senior, and mezzanine tranches for different input correlation levels using the standard Gaussian model assumptions. As can be observed, the risk of the equity tranche decreases while that of the senior tranche increases with increasing correlation, whereas the mezzanine tranche is insensitive to correlation when correlation is higher than 25% and lower than 50% in our example. The insensitivity of the mezzanine tranche to correlation values implies that fairly similar tranche prices, when inverted, can produce substantially different implied correlations for this tranche, thus sometimes causing large deviations between implied correlation and true correlation.

As discussed earlier, in such cases as the heterogeneous correlations and heterogeneous spreads, the smile pattern is driven mainly by the mezzanine tranche. When the price-generating model is different from the standard Gaussian model (e.g., because of the heterogeneous correlations), the 3%–7% mezzanine tranche value may shift only slightly, but in terms of implied correlation, this translates into a much larger deviation from the input correlation. Furthermore, rather than obtaining a unique solution for the implied correlation, we are sometimes likely to get two alternate solutions. As can be inferred from Exhibit 21, one of the solutions suggests an unrealistically large value of the implied correlation (and a correlation frown instead of a smile);²⁴ hence it is a common practice to choose the lower correlation value as the solution. Consistent with this practice, we also report the lower correlation value in our results. The smile produced by the mezzanine tranche

EXHIBIT 21

Fair Values of CDO Tranches for Different Correlation Levels

This exhibit shows the fair values of 0%–3% equity tranche, 3%–7% mezzanine tranche, and 7%–10%, 10%–15% and 15%–30% senior tranches for different asset correlation levels using the standard Gaussian model. Spreads are plotted on the primary (left) y-axis for all tranches except the equity tranche. For the equity tranche, upfront points are shown on the secondary y-axis. The plot highlights the insensitivity of the mezzanine tranche to a wide range of correlation values.



becomes more pronounced as the actual levels of “true,” or input, correlations increase due to the general correlation insensitivity of the mezzanine tranche at medium to high levels of input correlation.²⁵

Implied correlations for pricing tranches through time.

In the previous sections, we show that implied correlations obtained using the standard Gaussian model are rather inaccurate measures of true correlations among underlying credits. They can deviate substantially from true correlations because many of the simplifying assumptions in the standard Gaussian model do not describe empirical regularities realistically. We also emphasize the perils of executing correlation trades across tranches on the basis of implied correlations.


Despite the poor performance of the standard Gaussian model for relative value assessments across different tranches using implied correlations, the model is still quite appealing for its simplicity and efficiency in

implementation. In this respect, a natural question is whether the model is useful for pricing a given CDO tranche through time. To address this question, we conduct an empirical analysis using actual market data on the CDX.NA.IG index and its standardized tranches, i.e., we price (0%–3%), (3%–7%), (7%–10%), and (10%–15%) tranches through time. We do not include the (15%–30%) tranche due to a large number of missing market quotes for this tranche in our dataset. Our analysis is based on the following methodology. We use observed CDX tranche prices (from GFNet database), recovery estimates (from GFNet database),²⁶ and CDS spreads on names in the index (from Mark-It database) to back out an implied correlation for each tranche for a given date t . Next, we compute the model price of each tranche for the following period date, i.e., $t + \Delta t$, using the standard Gaussian model. In our case Δt corresponds to a week or seven calendar days. The inputs to the model are the CDS spreads and recovery

EXHIBIT 22

Actual and Model Prices of CDX.NA.IG Index Tranches

Panels A, B, C, and D show the actual and model prices of equity, (3%–7%) mezzanine tranche, (7%–10%) tranche, and (10%–15%) tranche, respectively. Model prices are computed using the standard Gaussian copula model where the lagged implied correlation of the tranche is used as the input correlation.




estimates as of date $t + \Delta t$ as well as the implied correlation for this tranche estimated as of date t . We then compare our model tranche prices with actual tranche prices on date $t + \Delta t$ to gauge the relative pricing performance of the standard Gaussian model. We repeat this process weekly until the end of the analysis period.

Our analysis covers the period from March 3, 2005, to August 17, 2005. We choose this particular period since it includes a period of stress in the credit derivatives

market that was triggered by rating downgrades of GM and Ford, both components of the CDX.NA.IG index. Both credits were simultaneously downgraded by S&P on May 5, 2005. Thus, our sample period allows us to test the performance of the standard Gaussian model in normal times²⁷ as well as in periods of market stress, i.e., around May 5, 2005.

EXHIBIT 22 (continued)



The time series of model and actual tranche prices are presented in Exhibit 22 for various standardized tranches. The results suggest that implied correlations computed from the standard Gaussian model perform well in pricing a given tranche through time. This is valid across different tranches and through both stable and turbulent periods in the markets. The model prices for the mezzanine tranche occasionally have a higher deviation from actual prices. This occurs at points where implied correlations could not be backed out for this tranche and zero correlation is used instead.²⁸ In particular, the standard Gaussian model is able to capture the considerable widening of tranche premiums around the downgrades of

GM and Ford. Thus this model appears valuable for pricing a given tranche through time.

Overall, the presented evidence suggests that although the implied correlations based on a standard Gaussian model are not good indicators of relative values *across tranches*, these correlations are still quite useful for pricing a given tranche *through time* and allow the user to benefit from the simplicity of the standard Gaussian model.

CONCLUSION

Given the importance of correlation trading and the pervasiveness of a smile pattern in implied correlations, it is important to understand why such a pattern arises.

It is standard practice in the marketplace to back out implied correlations using a simple Gaussian model with homogenous pair-wise correlations, homogenous spreads, and constant recoveries for underlying assets.

While the simplicity of this model is appealing, the evidence presented in this study shows that its assumptions are in violation of empirical regularities. Ignoring these violations can give rise to correlation smile patterns even when tranches are fairly priced. Specifically, we show that when the true price-generating model is different from the standard Gaussian model (but the latter is still used to back out implied correlations), non-flat patterns can arise due to any of the following assumptions: 1) Gaussian distributions of underlying asset returns, 2) homogenous pair-wise correlations across assets, 3) homogenous spreads across the underlying credits, and 4) default probabilities being independent of recovery rates. Among the various assumptions, the one of normally distributed returns is the most problematic. In reality, returns are likely to have distributions that are somewhat fat tailed. When this empirical regularity is ignored in favor of Gaussian distributions, the correlation smile is pronounced and is very similar to what is observed in practice, based on market prices of tranches. Thus, if one has to focus on one particular direction to make the standard Gaussian model more realistic, relaxing the distributional assumptions to account for the fat tails seems to be the most promising direction. We have also shown that making the distribution of the common factor fat-tailed has a much bigger payoff than relaxing the normality of the idiosyncratic components. On the other hand, we also find that assumptions like the homogeneity of spreads, although almost always violated in real world portfolios, turn out to not have a material impact in creating a correlation smile for common investment-grade portfolios like the CDX.NA.IG. Thus, one may be able to retain such assumptions in order to keep the model more tractable.

Furthermore, we show that, although this model is not reliable for valuations across different tranches, it works reasonably well for pricing a given tranche across time—through periods of market stability as well as dislocations. Our results thus highlight the strengths and weaknesses of the standard Gaussian model, by far the most commonly used paradigm for valuing CDO tranches. These findings can be used by the modeling community to decide the enhancements that will make the model more realistic, while retaining its appeal in terms of simplicity and tractability.

ENDNOTES

This study was initiated while Deepak Agrawal and Saiyid Islam were at KMV. We acknowledge KMV's support in providing some of the data for this study. We would also like to thank Don Chance, Anurag Gupta, Bill Morokoff, Yim Lee, the editor, and conference participants at the 2005 Hedge Funds World Conference, 2005 Advanced Correlation Modeling Conference, 2007 Financial Management Association Meeting, 2007 Washington Area Finance Conference, and seminar participants at the FDIC for helpful discussions. We are especially grateful to the editor, Stephen Figlewski, for constructive comments that have vastly improved the article. Any remaining errors are our responsibility. The views expressed in this article are the authors' own and do not necessarily represent the views of Standard & Poor's, KMV or Diversified Credit Investments.

¹Amato and Gyntelberg [2005] provide an overview of how standardization has been one of the main drivers of active trading in CDS indices and index tranche markets.

²For example, an article in the December 2004 issue of *Credit* magazine suggests that one can infer relative mispricing across tranches by looking at the implied correlation chart. It says "investors are not being properly compensated for the risk of buying equity and junior mezzanine tranches.... The implied correlation smile shows the current strength of demand for equity and mezzanine tranches." The article suggests "...advising clients to take advantage of the richness of the 3%–6% tranche in particular, either to express a bearish view on the market or as a cheap hedge against other tranches...."

³See Laurent and Gregory [2005] and Burtschell, Gregory, and Laurent [2005] for a comparative analysis of alternative CDO pricing models. See Hull and White [2004] for details of CDO pricing.

⁴Chriss [1996] provides a detailed discussion of implied volatility.

⁵Equity tranches are usually quoted on an upfront fee basis. A fall in spread is equivalent to a fall in the upfront fee.

⁶For example, see Hull and White [2004], who state "the standard market model has become a one-factor Gaussian copula model with constant pairwise correlations, constant CDS spreads, and constant default intensities for all companies in the reference portfolio."

⁷In this framework, CDS spreads are assumed to price only default risk and other risk factors, such as liquidity, demand-supply conditions, etc., are assumed to be negligible.

⁸Suppose that for a time period from 0 to T , default occurs between t and dt . The present value of expected spread earnings for a constant default intensity λ and discount rate r is $\int_0^T se^{-(\lambda+r)t} dt$. The present value of expected loss is $\int_0^T LGD * \lambda * e^{-(\lambda+r)t} dt$. Equating present value of expected loss

with the present value of expected spread earnings results in $\lambda = s/LGD$.

⁹We obtain index spread data from Mark-It Group. CDX.NA.IG Series 5 is initiated in September 20, 2005. To avoid any beginning of the week effect, we use the index spread as of September 22, 2005, which is a Wednesday.

¹⁰Results with 12 degrees of freedom are not presented in Exhibit 5 to conserve space.

¹¹We thank the editor, Stephen Figlewski, for this suggestion.

¹²Hull and White [2004] find a good fit between model prices and market quotes for the iTraxx EUR index tranches using a degree of freedom equal to four.

¹³In the standard model, default intensity, λ , is equal to the credit spread divided by loss given default. Therefore, for a spread of 49 bps and a recovery rate of 50%, default intensity is 0.0098. As a result, default probability over a horizon of five years is 4.8% ($=1 - \exp(-\lambda t) = 1 - \exp(-0.0098 * 5)$).

¹⁴Cumulative distribution function (c.d.f) of t density with four degrees of freedom is evaluated at -1.8223, c.d.f of standard normal density is evaluated at -2.0004, and c.d.f of t density with four degrees of freedom is evaluated at 2.1861. These default thresholds are estimated using empirical distribution discussed in the "Methodology" subsection.

¹⁵Note that at a portfolio level the impact of fat tails due to t -distributed idiosyncratic factors is expected to be less due to the law of large numbers. However, since the process of CDO structuring distributes the systematic and idiosyncratic risks unevenly, fat tails of individual factors become important in a CDO context.

¹⁶KMV correlation estimates are available by subscription.

¹⁷Mashal, Naldi, and Tejwani [2004] and Hager and Schobel [2005] also discuss the impact of homogenous correlation assumption on implied correlations.

¹⁸KMV uses a factor model to measure asset return correlations between firms. More specifically, KMV estimates an R^2 value for each firm, which is the proportion of a firm's risk that can be explained by systematic factors. Thus R^2 captures the systematic component of a firm's risk and in a single-factor world, the product $R_i R_j$ is a measure of the correlation between firm i and j . Using KMV estimates of R^2 for the 125 names in our reference portfolio, we construct a heterogeneous correlation structure that is more reflective of real-world correlations and uses these correlations to estimate the tranche fair values and corresponding implied correlations.

¹⁹See Acharya, Bharath, and Srinivasan [2007]; Altman, Brooks, Resti, and Sironi [2005]; Carey and Gordy [2004]; and Frye [2000] for empirical evidence. Bakshi, Madan, and Zhang [2006] and Das and Hanouna [2006] suggest methods to extract implied recoveries from market data considering the stochastic nature of the recovery levels. Altman, Resti and Sironi [2004]

provide a detailed review on the theoretical and empirical developments related to recovery rates.

²⁰See Table 3 in Altman, Brooks, Resti and Sironi [2005, p. 2215].

²¹The increase in the mean of loss distribution is mainly driven with higher recovery rates in low default environments and lower recovery rates in high default environments due to the positive correlation between recovery rates and the systematic factor. Note that default distributions do not change whether recovery is stochastically linked to the systematic factor or it is constant. This is because default distributions are functions of the collateral asset default probabilities and their pair-wise asset return correlations only.

²²Exhibit 19 also shows that for low levels of default, i.e., those that impact the equity tranche, the average recovery is higher than 50%, thereby resulting in equity becoming less risky when recovery rates are correlated with the systematic factor.

²³We thank the editor, Stephen Figlewski, for this suggestion.

²⁴Correlations are believed to be at the lower end of the spectrum because the underlying index constituents typically represent a diverse set of industries or sectors.

²⁵In Exhibit 21, we can see that the correlation insensitivity of the 3%–7% mezzanine tranche manifests for input correlations greater than about 25% and continues to a correlation of about 50%. From our simulations, we have seen that these thresholds can vary depending on the characteristics of the reference portfolio and the structure of the tranches. Initially, when input correlation is small, the portfolio loss distribution at any given horizon is unimodal. However, as the input correlation increases, at some point the loss distribution becomes bimodal. As an extreme case, when the correlation is 100%, the whole reference portfolio behaves like one exposure in the case of a homogeneous portfolio and its loss distribution is bi-modal, i.e., loss is either zero or its maximum value. The value of the mezzanine tranche is a function of the probability mass covering a loss greater than 3%. This probability mass initially increases as correlation increases, then becomes almost constant, and finally starts decreasing again as the shape of the distribution changes. This behavior is reflected in the spread of the mezzanine tranche and explains why the thresholds exist between which the spread is almost flat.

²⁶GFINet is an inter-dealer/broker CDS trading platform that also provides CDS data on a subscription basis.

²⁷We consider normal times to be periods of low credit-spread volatility and absence of any market disrupting event, such as the downgrade of GM and Ford that took place on May 5, 2005.

²⁸In these cases, a zero correlation produced model prices closest to market prices.

REFERENCES

- Acharya, V., S.T. Bharath, and A. Srinivasan. "Does Industry-Wide Distress Affect Defaulted Firms? Evidence from Creditor Recoveries." *Journal of Financial Economics*, Vol. 85, No. 3 (2007), pp. 787–821.
- Altman, E., A. Resti, and A. Sironi. "Default Recovery Rates in Credit Risk Modeling: A Review of the Literature and Empirical Evidence." *Economic Notes*, Banca Monte dei Paschi di Siena SpA, 33 (2004), pp. 183–208.
- Altman, E., B. Brooks, A. Resti, and A. Sironi. "The Link between Default and Recovery Rates: Theory, Empirical Evidence and Implications." *Journal of Business*, 78 (2005), pp. 2203–2227.
- Amato, J., and J. Gyntelberg. "CDS Index Tranches and the Pricing of Credit Risk Correlations." *BIS Quarterly Review*, 2005, pp. 73–87.
- Bakshi, G., D. Madan, and F. Zhang. "Understanding the Role of Recovery in Default Risk Models: Empirical Comparisons and Implied Recovery Rates." Working paper, University of Maryland, 2006.
- Black, F., and M.S. Scholes. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, 81 (1973), pp. 637–654.
- Burtschell, X., J. Gregory, and J.P. Laurent. "A Comparative Analysis of CDO Pricing Models." Working paper, BNP-Paribas, 2005.
- Carey, M., and M. Gordy. "Systematic Risk in Recoveries Debt." Mimeo, Federal Reserve Board, Washington, DC, 2004.
- Chriss, N.A. *Black-Scholes and Beyond: Option Pricing Models*. McGraw-Hill, 1996, pp. 327–410.
- Das, S.R., and P. Hanouna. "Implied Recovery." Working Paper, Santa Clara University, 2006.
- Das, S.R., and R.K. Sundaram. "Of Smiles and Smirks: A Term Structure Perspective." *Journal of Financial and Quantitative Analysis*, 34 (1999), pp. 211–239.
- Fama, E. "The Behavior of Stock Market Prices." *Journal of Business*, 38 (1965), pp. 34–105.
- Frye, J. "Collateral Damage Detected." Federal Reserve Bank of Chicago, Emerging Issues Series, October 1–14, 2000.
- Gregory, J., and J.P. Laurent. "In the Core of Correlation." *RISK* (October 2004), pp. 87–91.
- Hager, S., and R. Schobel. "A Note on the Correlation Smile." Working Paper, Tübingen University, 2005.
- Hull, J., M. Predescu, and A. White. "The Valuation of Correlation-Dependent Credit Derivatives Using a Structural Model." Working Paper, University of Toronto, 2005.
- Hull, J., and A. White. "Valuation of a CDO and nth to Default CDS without Monte Carlo Simulation." *The Journal of Derivatives*, 12 (Winter 2004), pp. 8–23.
- Laurent, J.P., and J. Gregory. "Basket Default Swaps, CDOs and Factor Copulas." *Journal of Risk*, 7 (2005), pp. 103–122.
- Li, D. "On Default Correlation: A Copula Function Approach." *The Journal of Fixed Income*, 9 (2000), pp. 43–54.
- Mashal, R., M. Naldi, and G. Tejwani. "The Implications of Implied Correlation." Working paper, Quantitative Credit Research, Lehman Brothers, 2004.
- Merton, R.C. "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates." *Journal of Finance*, 2 (1974), pp. 449–471.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675.

©Euromoney Institutional Investor PLC. This material must be used for the customer's internal business use only and a maximum of ten (10) hard copy print-outs may be made. No further copying or transmission of this material is allowed without the express permission of Euromoney Institutional Investor PLC. Source: Journal of Derivatives and
<http://www.iijournals.com/JOD/Default.asp>