UBS Warburg

CDO Insight

MAY 30, 2003

Inside This Issue

Market Highlights

- **Primary** The increase in primary CDO issuance that began in April has carried into May. CDOs backed by high yield loans continue to be most prevalent, followed by structured finance CDOs and synthetic CDOs based on investment grade corporate credits. Given collateral spreads and the tastes of investors, a larger share of structured finance CDOs are concentrating on AA and AAA collateral and/or incorporating prohibitions against "unusual" assets. Cash SF CDOs are also copying their synthetic brethren by dividing their AAA tranche into two classes. Sometimes, as in synthetic CDOs, there is true junior-senior credit tranching of these classes. Other times, the tranching is only to provide different maturities. If the CDO experiences credit difficulties, the two tranches become pari-passu.

- **Secondary** The market continues to enjoy activity derived from the Abbey liquidation bid lists and more entrants continue to explore and participate in the market. Trading is centered among first priority notes with fewer trades getting done in mezzanine and equity tranches. What happens after the last Abbey bid list? We think the secondary market will continue to grow, albeit after a fall from the one time effect of the Abbey bid lists. Can the primary market maintain or increase issuance without secondary pricing falling closer in line with primary offerings? Two answers. First, secondary pricing has tightened relative to primary among first priority notes. Second, for whatever reason, there seem to be CDO buyers who do not wish to venture out of the primary market.

- **This month's article** is about n^{th} (1^{st}, 2^{nd}, 3^{rd}) to default swaps and notes. We explain their mechanics and compare them to single name default swaps and senior and subordinated basket swaps. We use Venn diagrams to portray the risks of various n^{ths}. N^{th} to default swaps and notes are plays on default correlation, as well as on default probability. Thus, our article, and its two lengthy appendices, is also a basic to advance guide to default correlation. It turns out that the on-off nature of default gives default correlation unusual characteristics. These characteristics may be surprising to those more familiar with the correlation of continuous variables such as stock returns or interest rates.

High Yield CLO/CBO Spreads

Rating	5/30/2003	4/30/2003	3/31/2003
AAA	60-70	58-70	60-70
AA	110-130	105-130	110-130
A	180-250	175-250	180-230
BBB	310-400	310-400	320-390
BB	850-900	850-900	850-900
B	950+	950+	950+

Spreads to 3-month LIBOR

CDO, RMBS, ABS & CMBS Strategy Group

Laurie Goodman 212-713-4141
Thomas Zimmerman 212-713-6257
Jeffrey Ho 212-713-6258
Glenn Boyd 212-713-2718
Laurent Gauthier 212-713-4725
Douglas Lucas 212-713-3440
Wilfred Wong 212-713-7804
Victoria Ye 212-713-1458

Global CDO Group

Jeff Herlyn
Michael Rosenberg 203-719-1792

James Stehli
Lirenn Tsai 212-713-4972

Nth to Default Swaps and Notes: All About Default Correlation

An *n^{th} to default swap* is a credit default swap that references a basket of underlying credits, typically three to five names. The protection seller under the swap is exposed to the default of the reference credit that defaults "n^{th}" (first, second, third…). An n^{th} to default *note* is a credit-linked note that embeds this type of default swap in its terms. The purchaser of the note is the seller of credit protection in the embedded default swap. In this *CDO Insight*, we delve into n^{th} to default swaps and notes, define their characteristics, and compare their risks to other derivative and funded instruments.

More than other financial instruments, n^{th} to default swaps and notes are plays on *default correlation*. Simply put, default correlation measures whether credit risky assets are more likely to default together or separately. For example, default correlation answers the following question: does a 10% probability of default mean that *one* out of 10 credits is going to default, or that 10% of the time, *all* 10 credits are going to default? If the answer is "in between," where in between?

Default correlation is essential to understanding the risk of n^{th} to default swaps and notes. In fact, default correlation is essential to understanding the risk of credit portfolios generally. Along with *default probability* and *loss in the event of default*, default correlation determines the credit risk of a portfolio and the economic capital required to support that portfolio. It is important to portfolio managers, investors, CFOs, bankers, rating agencies, and regulators.

Yet, compared to knowledge about other credit factors, fewer people are familiar with the basic principals of default correlation. Available papers on default correlation are often so mathematically complex as to be obscure. In the process of describing the risks of n^{th} to default swaps and notes, this *CDO Insight* provides a basic to advanced guide to default correlation, making the maximum use of graphic illustrations, the minimum use of algebra, and no use of calculus. To spare readers who do not completely share our obsession with default correlation, we have parked the more technical material on default correlation in two appendices.

Appendix I delves into pictorial representations of default probability and default correlation. But the main focus of the Appendix I is what we term "higher orders of default correlation." For example, we explore the default correlation that exists, once a pair of credits have defaulted, between the joint default probability of that pair of credits and the probability that a *third* credit will join them in default. Readers used to the correlation of continuous variables like stock prices and interest rates may be surprised that pairwise default correlation does not fully describe default distributions.

Appendix II shows the calculation and results of historic default correlation. We show that default correlations among well-diversified portfolios vary by the ratings of the credits and also by the time period over which defaults are examined. We also describe some of the problems measuring, and even thinking about, default correlation.[1]

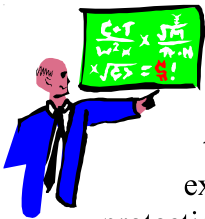

N^{th} to Default Swaps

An n^{th} to default swap is a credit default swap that references a basket of underlying credits, typically three to five names. The *protection seller*[2] under the swap is exposed to the default of the reference credit that defaults "n^{th}" (first, second, third, …). For example, if a party has sold protection on the 3^{rd} to default of five reference credits, the protection seller is responsible for default losses associated only with the default of the third reference credit to default. If only two reference credits default over the tenor of the swap, the protection seller does not have to make a default payment to the protection buyer.

Upon the default of the third credit in our example, the protection seller pays default losses associated with the reference credit to the protection buyer and the swap terminates. In physical settlement, the protection buyer of a $10 million notional swap delivers $10 million par of a *deliverable obligation* of the defaulted name to the protection seller. The protection seller pays $10 million to the protection buyer. The protection seller is then free to retain or sell the obligation as it sees fit. In cash settlement, the market price of $10 million par of a deliverable obligation of the defaulted name is determined through a specified process of dealer polling. The protection seller pays the protection buyer the difference between $10 million and the determined market value of the deliverable obligation. Note that the severity of default losses associated with the first and second defaults among the reference credits does not make a difference to the payout under a 3^{rd} to default swap.

After paying default losses on the n^{th} to default credit, the protection seller has no responsibility for subsequent defaults of reference entities. In return for this protection, the protection buyer pays a periodic fee to the protection seller. These cash flows are depicted in Exhibit 1 (below). Most n^{th}

Exhibit 1. N^{th} to Default Swap Cash Flows

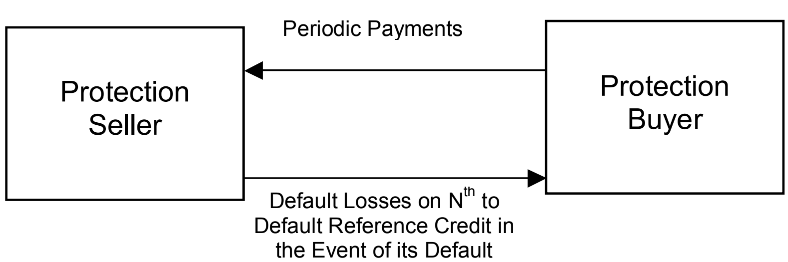

[1] Much of the material in this article and its appendices is expanded from our *Rating Cash Flow Transactions Backed by Corporate Debt*, Moody's Investors Service, September, 1989 and "Default Correlation and Credit Analysis," *Journal of Fixed Income*, March 1995. What's new is the application of default correlation to n^{th} to default swaps in the article, the quantification of higher orders of default correlation (beyond pairwise default correlation) in Appendix I, and discussion of the difficulties in relying upon historic empirical default correlations in Appendix II.

[2] The expert in credit default swap documentation will note that we use a mix of official ISDA defined terms and expressions whose meanings we feel are easier to understand. Thus "protection buyer" for "Fixed Rate Payer," "protection seller" for "Floating Rate Payer," "default" for the occurrence of a "Credit Event," and "default losses" for the difference between par and market value of the cheapest "Deliverable Obligation."

Exhibit 2. N^{th} to Default Note Cash Flows

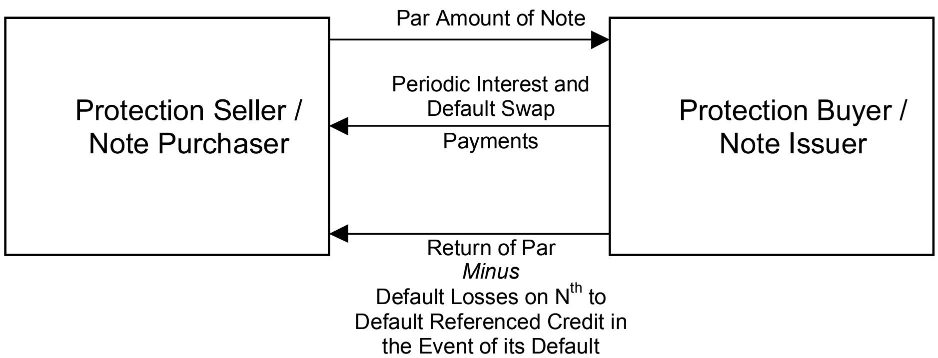

to default swaps are 1^{st} to default swaps on three to five underlying names, each rated single-A or above. Higher "n^{ths}" generally reference more names. Most portfolios are comprised of names from different industries but often the same geographical region, i.e., U.S. or Europe. The most popular n^{ths} are the first, second, and third.

It bears mentioning that like other credit default swaps, an n^{th} to default swap is purely a play on credit risk. It is not affected by interest rates movements and even currency risk is usually structured away.

N^{th} to Default Notes

An *n^{th} to default note* is a credit-linked note that embeds an n^{th} to default swap in its terms. The purchaser of the note is also the seller of credit protection on the embedded default swap. The note issuer is also the buyer of credit protection on the embedded default swap. An n^{th} to default note therefore combines two separate financial instruments in one: (1) a note and (2) an n^{th} to default swap. Most notes are based on 1^{st} to default swaps with three to five underlying reference names rated A or AA. If rated, the n^{th} to default note itself is usually rated A or BBB. The cash flows of an n^{th} to default note are shown in Exhibit 2 (above).

The interaction or intersection between the *credit swap* part the n^{th} to default note and the *note* part of the n^{th} to default note occurs when the protection seller/note purchaser owes default losses to the protection buyer/note issuer. The protection seller does not make a payment of default losses to the protection buyer. Rather, the protection buyer either delivers the deliverable obligation (physical settlement) or the protection buyer pays the protection seller the recovery value of the deliverable obligation. The note is thereby redeemed and the embedded swap terminated.

It is sometimes convenient for the protection seller that the n^{th} to default swap be embedded in a note. There might be internal or external regulatory reasons for doing so, for example. For example, the investor might be prohibited from entering into a derivative contract, but an nth to default note might not be considered a derivative. An n^{th} to default *note* isolates the protection buyer from the credit risk of the protection seller. By buying the note, the protection seller has effectively collateralized its potential obligation to the protection buyer. Almost the same result could be obtained, however, by having the protection seller of an n^{th} to default swap collateralize the mark-to-market of the swap to the protection buyer. Within limits, this would give the protection seller the ability to choose what collateral it owns and posts to the protection buyer. The protection seller can also replace one acceptable collateral instrument with another. But in an n^{th} to default note, the protection seller must buy the protection buyer's debt instrument.[3]

Also, in note form, the protection seller is now at risk to the protection buyer for what we have termed the protection seller's collateral, a.k.a. the principal amount of the note. An n^{th} to default protection seller/note purchaser must consider not only the credit risks of the reference credits in the swap basket, but also the credit risk of the note issuer.

Comparison of N^{th} to Default Swaps to Other Credit Swaps

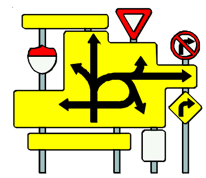

To help understand the risks of n^{th} to default swaps, we will compare them to single name default swaps and basket default swaps. We will look at the following six credit default swaps:

- a $10 million *1st* to default swap on five underlying credits;
- a $10 million *5th* to default swap on five underlying credits;
- a portfolio of *five separate $2 million single name credit default swaps*;
- a portfolio of *five separate $10 million single name credit default swaps*;
- a *$10 million subordinate basket credit default swap* responsible for the first $10 million of default losses on a portfolio of five underlying credits of $10 million each;
- a $*10 million senior basket credit default swap* responsible for default losses above $40 million on a portfolio of five underlying credits of $10 million each.

These credit swaps are shown in Exhibit 3 (next page), in order of their risk, which we define as the amount of potential default losses the protection seller might have to absorb. We assume that all credit swaps reference the same five names, and explain their relative risks.

The riskiest credit derivative in Exhibit 3 is the $50 million swap portfolio comprised of five separate $10 million single name credit default swaps. This portfolio of separate swaps is more risky than the subordinated basket swap. Protection sellers in both situations are equally exposed to the first $10 million of default losses from the five names. But the protection seller's aggregate exposure under the subordinate basket swap is capped at $10 million. The protection seller of the five single name swaps is exposed to an additional $40 million of potential default losses.

The subordinate basket swap is riskier than the 1st to default swap. Protection sellers under both these swaps are equally exposed to default losses from the first of the credits to default. But the subordinate basket swap protection seller is also exposed to the 2nd, 3rd, 4th, and 5th defaults in the portfolio up to $10 million aggregate of default losses.

[3] An alternative structure would have the protection seller buy a note for X amount of par, but be responsible for some Y amount of losses where Y > X. In this manner, the protection buyer is collateralized for the first X/Y% of default losses that might occur. X would be set relative to Y to allow for the expected loss in the event of default.

Exhibit 3. Comparison of Credit Swaps

Riskiness	Credit Default Swap	Losses
1	Portfolio of 5 separate $10 million single name credit default swaps	Default loss on each credit up to $10 million; capped at $50 million in aggregate
2	Subordinate basket swap. First $10 million of default losses on portfolio of 5 names @ $10 million each name	Default loss on each credit up to $10 million; capped at $10 million in aggregate
3	$10 million 1st to default swap on 5 names	Default loss on first credit to default; capped at $10 million
4	Portfolio of 5 separate $2 million single name credit default swaps	Default loss on each credit up to $2 million; capped at $10 million in aggregate
5	$10 million 5th to default swap on 5 names	Default loss on fifth credit to default; capped at $10 million
6	Senior basket swap. Last $10 million of default losses on portfolio of 5 names @ $10 million each name	Default loss after $40 million of losses have already occurred; capped at $10 million

The 1st to default swap is more risky than the $10 million swap portfolio comprised of five separate $2 million single name credit default swaps, but the situation is not as clear cut. First, assume that default losses on any credit will be of equal percentage of par, say 50%. Upon the first default among the five names, the protection seller of the 1st to default swap pays $5 million and the protection seller of the swap portfolio comprised of five separate $2 million single name credit default swaps pays $1 million. The only way for the protection seller on the portfolio of swaps to pay $5 million is if *all* five underlying credits default.

The only way that the protection seller on the portfolio of swaps would have to make larger total default loss payments is if there are specific patterns of defaults and percentage default losses among underlying credits. For example, suppose that default losses from the first defaulting credit were only 10% of par. The protection seller under the 1st to default swap would have to pay $1 million while the protection seller on the portfolio of swaps would have to pay $200,000. If one more credit defaulted with a default loss percentage above 40%, or if all four remaining credits defaulted with an average default loss percentage above 10%, then the protection seller of the portfolio of swaps would pay more than the protection seller of the 1st to default swap.

Note that if the loss percentage on the first credit to default was above 25%, it would take at least three defaults for the protection seller of the portfolio to possible pay more than the protection seller of the 1st to default swap. While certain scenarios of this sort are possible, they are not as likely as situations where the 1st to default swap protection seller has to pay more default losses.

Similarly, the $10 million swap portfolio comprised of five separate $2 million single name credit default swaps is more risky than the 5th to default swap. The only way that the protection seller of the 5th to default swap would have to make the larger default loss payment is if (1) all five credits default *and* (2) the average percentage loss on all five defaults was smaller than the default loss percentage of the fifth credit to default.

Finally, the 5th to default swap is more risky than the senior basket swap. Given that the protection seller under the senior basket swap has $40 million of subordination below it, it is not exposed to the 5th credit to default unless previous default losses exceed $30 million. Only if previous default losses have been $40 million (i.e., if all four credits have defaulted and their default losses were $10 million each) would it be as exposed to the 5th credit to default in a way equal to the protection seller under the 5th to default swap.

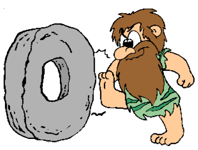

Picturing Nth to Default Risk

In this section, we are going to explore nth to default risk another way, using Venn diagrams. These are those intersecting circles that show the overlap or lack of overlap between two or more events or conditions. The area of circles labeled A, B, and C in Exhibit 4 (below) represent the probability that underlying credits A, B, or C are going to default over the term of an nth to default swap. In Exhibit 4, the circles have some overlap with one another. These overlaps represent the probability that more than one credit is going to default over the term of the nth to default swap. The "2s" in the exhibit indicate the area where two of the circles overlap and therefore represent the probability that two of the three credits will default. There are three such overlaps, between circles (and credits) AB, BC, and AC. There is also an area where all three circles overlap, representing the probability that all three credits will default over the term of the nth to default swap.

Exhibit 4 says something about the relative risks of 1st, 2nd, and 3rd to default swaps. The area within all three circles represents the probability that one or more credits will default. This is the probability that the protection seller of a 1st to default swap on these three underlying credits will have to pay default losses. The sections of Exhibit 4 labeled "2" and "3" represent the probability that two or three credits will default. This is the probability that the protection seller of a 2nd to default swap on these three underlying credits will have to pay default losses. Finally, the section of Exhibit 4 labeled "3" represents the probability that three credits will default. This is the probability that the protection seller of a 3rd to default swap on these three underlying credits will have to pay default losses. (The undefined area outside the three circles represents the probability that *no* credits will default.)

Exhibit 4. Number of Defaulting Credits

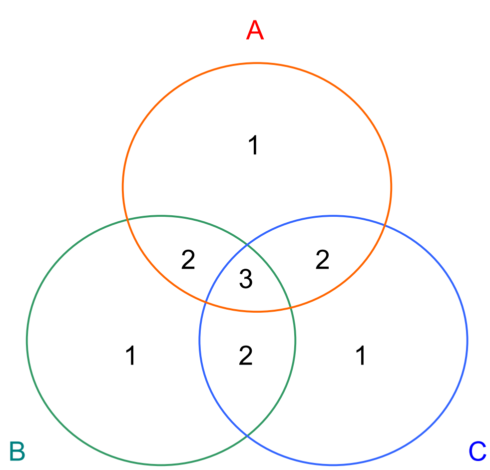

From Exhibit 4, we can tell a lot about the probability of a default loss pay out under an nth to default swap. However, this particular drawing is not the only way that credits A, B, and C might behave with respect to defaulting. It might be, for example, that Exhibit 5 (next page) is a more accurate representation of how defaults occur among these credits.

What we see in Exhibit 5 is the situation where credits A, B, and C *never* default at the same time. This is represented in the exhibit by the lack of overlap among circles A, B, and C. In this

Exhibit 5. Defaulting Credits Default Separately

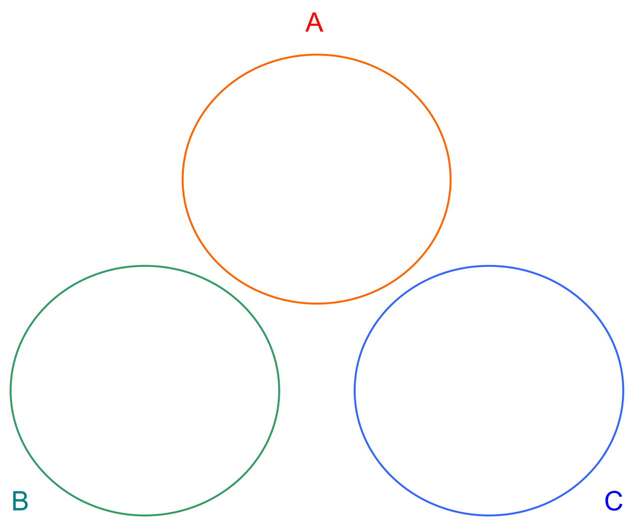

scenario, there is no risk to the protection seller of a 2nd or 3rd to default swap because 2nd and 3rd defaults among the portfolio will never occur. However, comparison of Exhibits 4 and 5 shows that there is *more* risk of one credit defaulting. The overlaps in Exhibit 4 are spread out in Exhibit 5 with the result that there is greater area covered by the circles representing greater probability that *one* credit will default.

Exhibit 6 (below) shows the opposite situation as Exhibit 5. Instead of defaults being spread out, and never occurring together, in Exhibit 6 defaults are bunched up and never occur separately. (Note that we draw the circles in Exhibit 6 a little offset so you can see that there are three of them. In theory, they rest exactly on top of each other.) Compared to both Exhibits 4 and 5, there is a lot of probability that two or three of the credits will default. In this scenario, there is a relatively more risk to the protection seller of a 2nd or 3rd to default swap. However, there is less risk of one or more credits defaulting and therefore less risk to the protection seller of a 1st to default swap.

Exhibit 6. Defaulting Credits Default Together

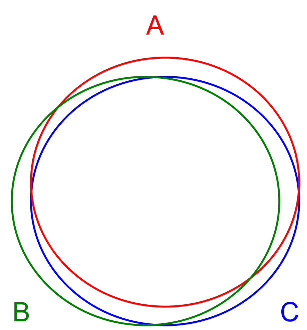

The difference between the default probability pictures in Exhibits 4, 5, and 6 is *default correlation*.

Default Correlation Defined

Default correlation is the phenomenon that the likelihood of one obligor defaulting on its debt is affected by whether or not *another* obligor has defaulted on *its* debts. A simple example of this is if one firm is the creditor of another: if Credit A defaults on its obligations to Credit B, we think it is more likely that Credit B will be unable to pay its own obligations. This is an example of *positive* default correlation. The default of one credit makes it *more* likely the other credit will default.

There could also be *negative* default correlation. Suppose that Credit A and Credit B are competitors. If Credit A defaults and goes out of business, it might be the case that Credit B will get Credit A's customers and be able to get price concessions from Credit A's suppliers. If this is true, the default of one credit makes it *less* likely the other credit will default. This would be an example of *negative* correlation.

But default correlation is not normally discussed with respect to the particular business relationship between one credit and another. And the existence of default correlation does not imply that one credit's default directly *causes* the change in another credit's default probability. It is a maxim of statistics that *correlation does not imply causation*. Nor do we think negative default correlation is very common. Primarily, we think *positive* default correlation generally exists among credits because the fortunes of individual companies are linked together via the health of the general economy or the health of broad subsets of the general economy.

Drivers of Default Correlation

The pattern of yearly default rates for U.S. corporations since 1920, shown in Exhibit 7 (below), is notable for the high concentrations of defaults around 1933, 1970, 1991, and 2001. A good number of firms in almost all industries defaulted on their credit obligations in these depressions and recessions. The boom years of the 1950s and 1960s, however, produced very few defaults. To varying degrees, all businesses tend to be affected by the health of the general economy, regardless of their specific characteristics. The phenomena of companies tending to default together or not default together is indicative of positive default correlation.

But defaults can also be caused by industry-specific events that only affect firms in those particular industries. Despite a favorable overall economy, low oil prices caused 22 companies in the oil industry to default on rated debt between 1982 and 1986. Bad investments or perhaps bad regulation caused 19 thrifts to default in 1989 and 1990. Recently, we experienced the defaults of numerous dot coms due to the correction of "irrational exuberance." Again, the phenomena of companies in a particular industry tending to default together or not default together is indicative of positive default correlation.

Exhibit 7. U.S. Corporate Default Rates Since 1920

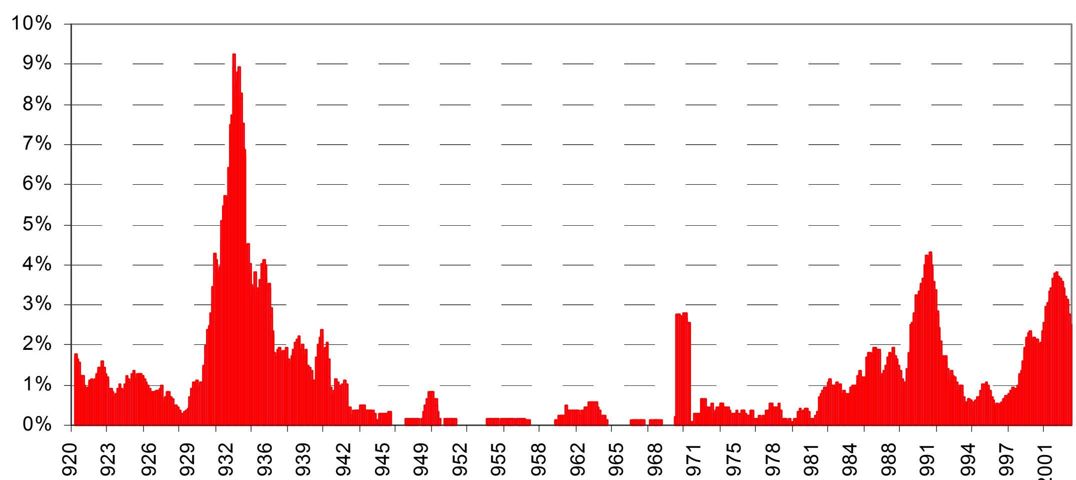

Source: Moody's Investors Service

There are other default-risk relationships among businesses that do not become obvious until they occur. The effect of low oil prices rippled through the Texas economy in the 1980s affecting just about every industry and credit in the state. A spike in the price of silver once negatively affected both film manufacturers and silverware makers. The failure of the South American anchovy harvest one year drove up the price of grains and put cattle farmers under pressure. These default-producing characteristics hide until, because of the defaults they cause, their presence becomes obvious.

Finally, there are truly company-specific default factors such as the health of a company's founder or the chance a warehouse will be destroyed by fire. These factors do not transfer default contagion to other credits.

Defaults are therefore the result of an unknown and unspecified multi-factor model of default that seems akin to a multi-factor equity pricing model. Default correlation occurs when, for example, economy-wide or industry-wide default-causing variables assume particular values and cause widespread havoc. Uncorrelated defaults occur when company-specific default-causing variables cause trouble.

Why We Care About Default Correlation

Default correlation is very important in understanding and predicting the behavior of credit *portfolios,* including n^{th} to default swaps. It directly affects the risk-return profile of investors in credit risky assets and is therefore important to the creditors and regulators of these investors. Default correlation even has implications for industrial companies that expose themselves to the credit risk of their suppliers and customers through the normal course of business. We will prove these assertions via an example.

Suppose we wish to understand the risk of a bond portfolio and we know that each of the 10 bonds in the portfolio has a 10% probability of default over the next five years. What does this tell us about the behavior of the portfolio as a whole? Not much, it turns out, unless we also understand the default correlation among credits in the portfolio.

It could be, for example, that all the bonds in the portfolio always default together. Or to put it another way, if one of the 10 bonds default, they all default. If so, this would be an example of "perfect" *positive* default correlation. Combined with the fact that each bond has a 10% probability of default, we can make a conclusion about how this portfolio will perform. There is a 10% probability that *all* the bonds in the portfolio will default. And there is a 90% probability that *none* of the bonds will default. Perfect positive default correlation, the fact that all the bonds will either default together or not default at all, combines with the 10% probability of default to produce this extreme distribution, as shown in Exhibit 8 (next page).

At the other extreme, it could be the case that bonds in the portfolio *always* default separately. Or to put it another way, if one of the 10 bonds default, no other bonds default. This would be an example of "perfect" *negative* default correlation. Combined with the fact that each bond has a 10% probability of default, we can make a conclusion about how this portfolio will perform: there is a 100% probability that *one* and only one bond in the portfolio will default. Perfect negative default correlation, the fact that when one bond defaults no other bonds default, combines with the 10% probability of default to produce this extreme distribution, as shown in Exhibit 9 (next page).

Exhibit 8. Extreme Positive Default Correlation

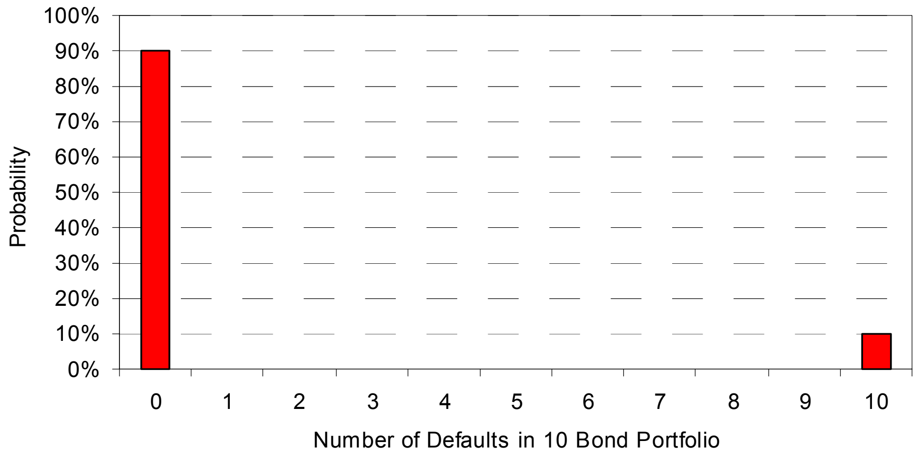

Exhibit 9. Extreme Negative Default Correlation

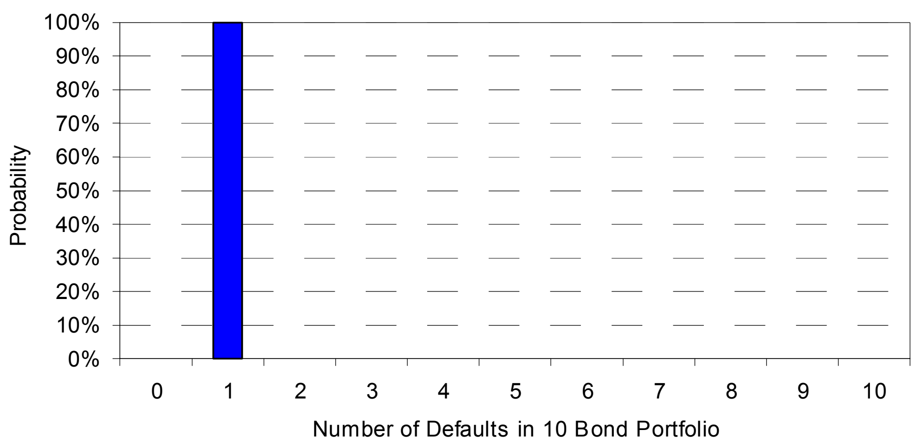

The difference in the distributions depicted in Exhibit 8 and 9 has profound implications for investors in these portfolios. Remember that in both cases, the default probability of bonds in the portfolio is 10%. The expected number of defaults in both portfolios is one. But one knows with *certainty* the result of the portfolio depicted in Exhibit 9: one and only one bond is going to default. This certainty would be of comfort to a lender to this investor. The lender knows with certainty that nine of the bonds are going to perform and that par and interest from those nine performing bonds will be available to repay the investor's indebtedness.

The investor in the portfolio depicted in Exhibit 8 has the greatest uncertainty. Ninety percent of the time the portfolio will have no defaults and 10% of the time every bond in the portfolio will default. A lender to an investor with this portfolio has a 10% risk that *no* bonds in the portfolio will perform.

A complete analysis of the risk of these two example portfolios would depend on the distribution of default recoveries. But it is obvious that the portfolio depicted in Exhibit 8 is much more risky than the portfolio depicted in Exhibit 9, even though the default probabilities of bonds in the portfolios are the same. The difference in risk profiles, which is due only to default correlation, has profound implications to investors, lenders, rating agencies, and regulators. Debt backed by the portfolio depicted in Exhibit 8 should bear a higher premium for credit risk and be rated lower. It this is a regulated entity, it should be required to have more capital.

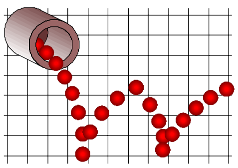

Default Probability and Default Correlation in 1st to Default Swaps

In Exhibit 10 (below), we relate default correlation to 1st to default swaps. The exhibit shows the probability that *at least one* of five credits will default. This is the risk to the protection seller of a 1st to default swap on five reference names. The exhibit shows this probability assuming different levels of default correlation, from –0.05 to 1.00.[4] The line labeled "5% Default Probability" also incorporates the assumption that each credit in the reference portfolio has a 5% probability of default.

Focusing on that line, note that over different default correlation assumptions, it moves from 25% probability of at least one reference credit defaulting to 5% probability of at least one reference credit defaulting. At the most extreme negative default correlation, which is –0.05 for this default probability, none of the five underlying credits default at the same time. The probability of one underlying credit defaulting is therefore 5 * 5% or 25%. At 1.00 default correlation, when one of the credits default, they all default. The probability of all five credits defaulting is 5%.

Exhibit 10. Correlation and 1st to Default Risk

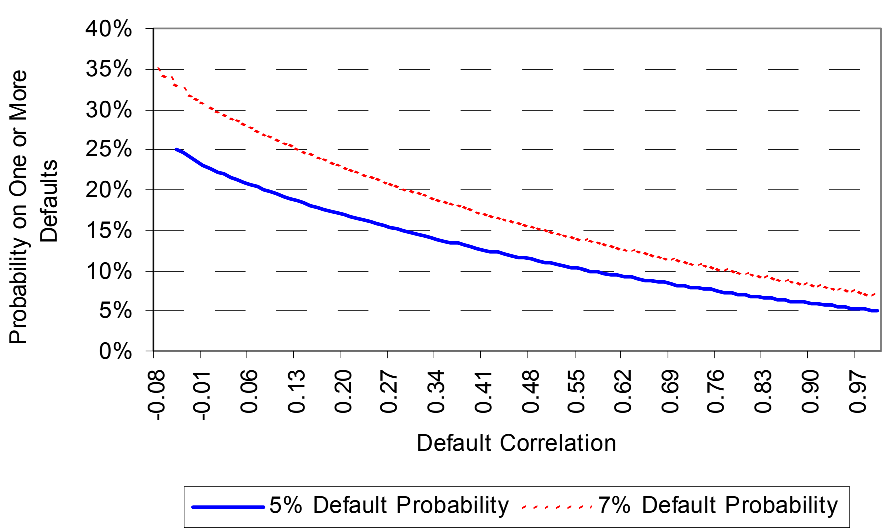

[4] We assume that default correlations between pairs of credits, the default correlation between a pair of credits and a third credit, the default correlation between three credits and a forth credit, and the default correlation between four credits and a fifth credit are all equal. In Appendix I, we discuss the insufficiently of pairwise correlations in describing the probability distributions of binomial events like default and the effect of "higher order" default correlations.

Higher default correlation *always* reduces the risk to a 1^{st} to default protection seller. Recall from Exhibit 6 that when defaults happen together, there is more of a chance that *no* defaults will occur. (In the exhibit, the probability of no defaults occurring is represented by the area outside circles A, B, and C.) As default correlation increases, we are moving from a picture of defaults like the one shown in Exhibit 5 to a picture of default like the one shown in Exhibit 6. A probability distribution like the one pictured in Exhibits 6 is more likely to produce cases where there are *multiple* defaults and more likely to produce cases where there are *no* defaults. The increase in the probability of no underlying reference defaults reduces 1^{st} to default risk.

Also in Exhibit 10, is a line labeled "7% Default Probability," which incorporates the assumption that every reference credit has a 7% probability of default. Just like the 5% line, it moves, left to right, from only one credit defaulting at a time to all credits defaulting together. In this case, that means a 35% probability of a swap payout at –0.08 default correlation and a 7% probability of a payout at 1.00 correlation. From being 10% above the 5% default line in extreme negative correlations, the 7% default line declines at a faster rate than the 5% default line until it is 2% higher than the 5% line.

We can also compare, in Exhibit 10, the relative effects of underlying reference credit default probability and default correlation on the probability of a payout under a 1^{st} to default swap. But first we need to decide where to look.

In Appendix II, we display a wealth of data on historical default correlations. The historical evidence is that default correlation among *well diversified* investment grade credits over five years is 0.00 or 0.01. In Appendix II, we also show how, if anything, our measure of default correlation is biased to report results that are too high. Other researchers looking at intra-industry default correlation have estimated default correlation for names in the same industry at as much as 0.10 or 0.20. So a very generous range of likely default correlations for a 1^{st} to default swap is from 0.00 to 0.30. On the other hand, it seems pretty easy to us for someone to get default probability wrong by 2% over five years.

Looking at the line for 5% underlying reference credit default probability over the range of default correlation from 0.00 to 0.30, the risk that at least one reference credit will default ranges from 22.6% to 15.8%. This is a difference of 6.9%. Changing the underlying reference credit default probability from 5% to 7% increases the risk that at least one reference credit will default an average of 6.9%. It seems to us much more likely that a protection seller will be off 2% on the true probability of default for the underlying reference credits than off by 0.30 in their estimation of default correlation. Since both mistakes would have the same result, this shows that 1^{st} to default swaps are much more sensitive to estimates of default probability than estimates of default correlation.

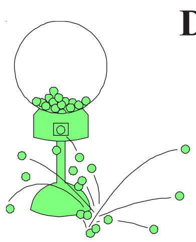

Default Probability and Default Correlation in 2^{nd} to Default Swaps

In Exhibit 11 (next page), we address the risk of 2^{nd} to default swaps, or the risk that at least *two* defaults will occur in our five-credit five-year 2^{nd} to default swap example. The situation with respect to default correlation is ambiguous. At first, default correlation increases the risk of a 2^{nd} to default swap and then it decreases that risk. What happens is that as default correlation first increases the probability that credits default together,

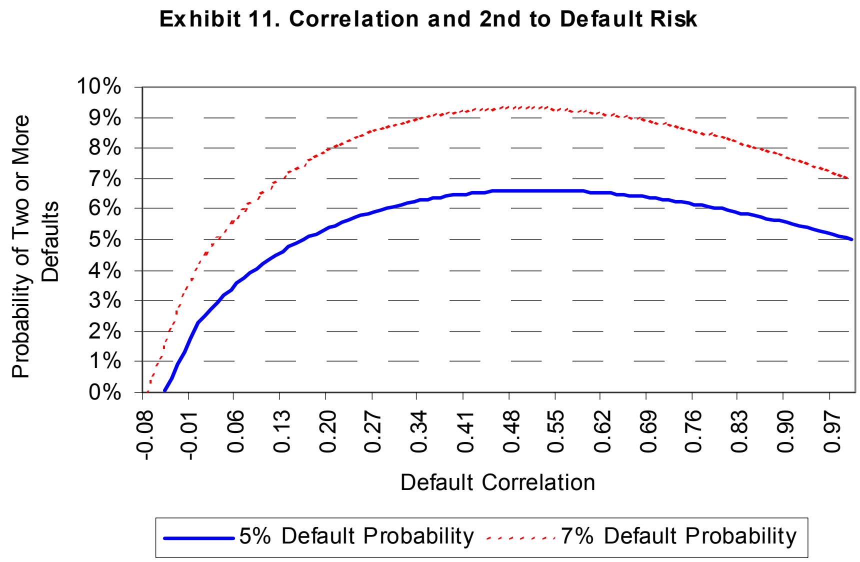

there is a greater chance that two, three, four, and five credits will default together. The chance of *two or more* defaults increases. But as default correlation increases still more, the chance of *exactly* two defaults peaks and then declines while the probability of three, four, and five credits defaulting continues to increase. It happens that the decrease in the probability of exactly two defaults is greater than the increase in the probability of *two or more* defaults. Thus, 2^{nd} to default risk decreases. We show this effect in Exhibit 12 (below).

In the relevant range of default correlation, 0.00 to 0.30, the probability of a 2^{nd} to default payoff, assuming 5% underlying reference credit default probability, ranges from 0.0% to 6.1%. Over this same range of default correlations, an increase in underlying reference default probability from 5% to 7% increases the probability of a 2^{nd} to default payout an average of 2.4%. So we see that default correlation is much more important in the pricing of 2^{nd} to default swaps than it is in the

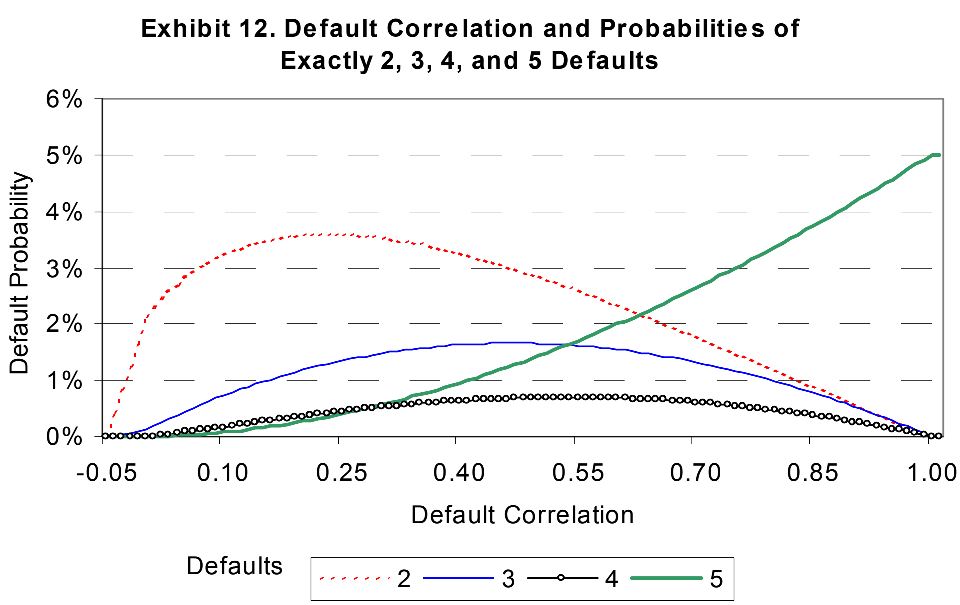

pricing of 1st to default swaps. This result generally holds for higher "nths" and different underlying reference credit default probabilities.

Conclusion

In this *CDO Insight*, we explained the mechanics of nth to default swaps and notes and then compared them to portfolios of single name default swaps and to senior and subordinated credit swaps on a portfolio of names. We ranked different credit swaps in order of their risk to the protection seller. We next "pictured" nth to default risk via the use of Venn diagrams representing credits defaulting together or separately. This lead right into a discussion of default correlation. We explained what negative and positive default correlation are, why default correlation is important in analyzing credit risky portfolios like nth to default swaps, and even the causes of default correlation.

Finally, we turned back to nth to default swaps and applied our understanding of default probability and default correlation. We addressed the probability of *at least one* credit defaulting out of a portfolio of five (1st to default risk) and the probability of at least *two* credits defaulting out of a portfolio of five (2nd to default risk). We saw that higher default correlation decreases 1st to default risk. But we said that reference credit default probability was relatively more important in estimating 1st to default risk than default correlation. We also saw that default correlation first increases and then decreases 2nd to default risk. And we pointed out that reference credit default correlation is relatively more important in estimating 2nd to default risk than default probability.

Nth to default swaps and notes allow investors to take on increase credit risk, but remain exposed only to investment grade credits. The limited number of names in the portfolio allows protection sellers to thoroughly vet the names to which they are exposing themselves. And the static nature of the reference portfolio eliminates any surprise from an asset manager and allows the purchase of single name default swaps later to hedge credit risks. Nth to default swaps are a particularly attractive product in low spread environments. We expect to see greater use of nth to default swaps and notes to take on credit risk.

Appendix I. Beyond Pairwise Default Correlation

In this appendix we show default probability and default correlation pictorially, present the basic algebra of default correlation, and then delve into the deficiency of pairwise correlations in explaining default distributions.

Picturing Default Probability

Suppose we have two obligors, Credit A and Credit B, each with 10% default probability. The circles A and B in Exhibit 13 (next page) represent the 10% probability that A and B will default, respectively.

Exhibit 13. Credit A and Credit B Default Probability, Pictorially

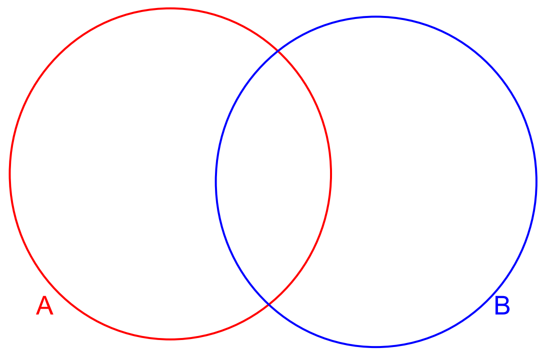

There are four possibilities depicted in Exhibit 13:

(1) both A and B default, as shown by the overlap of circles A and B;
(2) only A defaults, as shown by circle A that does not overlap with B;
(3) only B defaults, as shown by circle B that does not overlap with A;
(4) neither A or B default, as implied by the area outside both circles A and B.

Recall that we defined positive and negative default correlation by how one revises their assessment of the default probability of one credit once one finds out whether another credit has defaulted. If upon the default of one credit you revise the default probability of the second credit *upwards*, you implicitly think there is *positive* default correlation between the two credits. And if upon the default of one credit you revise the default probability of the second credit *downwards*, you implicitly think there is *negative* default correlation between the two credits.

Exhibit 13 is purposely drawn so that knowing whether one credit defaults does not cause us to revise our estimation of the default probability of the other credit. Exhibit 13 pictorially represents no or *zero default correlation* between Credits A and B, neither positive or negative default correlation. In other words, knowing that A has defaulted does not change our assessment of the probability that B will default.

Here's the explanation. Recall that the probability of A defaulting is 10% and the probability of B defaulting is 10%. Suppose A has defaulted. Now, pictorially, we are within the circle labeled A in Exhibit 13. No or zero correlation means that we do not change our estimation of Credit B's default probability just because Credit A has defaulted. We still think there is a 10% probability that B will default. Given that we are within circle A and circle A represents 10% probability, the probability that B will default must be 10% of circle A or 10% of 10% or 1%. The intersection of circles A and B depicts this 1% probability. This leads to a very simple general formula for calculating the probability that both A and B will default when there is no or zero default correlation.

Recall the phrase in the above paragraph that the overlap of A and B, or the space where both A and B have defaulted is "10% of 10% or 1%." What this means mathematically is the probability of both Credits A and B defaulting (the joint probability of default for Credits A and B) is 10% *

10% or 1%. Working from the specific to the general (which we label Equation 1), our notation gives us the following:

$$10\% \times 10\% = 1\%$$

$$P(A) * P(B) = P(A \text{ and } B) \qquad (1)$$

Where:

- P(A) = the probability of Credit A defaulting (10% in our example);
- P(B) = the probability of Credit B defaulting (10% in our example);
- P(A *and* B) = the probability of both Credits A and B defaulting, a.k.a., the *joint probability* of default for Credits A and B (1% in our example).

This is the general statistical formula for joint default probability assuming zero correlation.

Now that we have calculated the joint probability of A and B defaulting, we can assign probabilities to all the alternatives is Exhibit 13. We do this in Exhibit 14 (left). We assumed that the default probability of Credit A was 10%, which we represent by the circle labeled A in Exhibit 14. We have already determined that the joint probability of Credit A and Credit B defaulting, as represented by the intersection of the circles labeled A and B, is 1%. Therefore, the probability that Credit A *will* default and Credit B *will not* default, represented by the area within circle A but also outside circle B, is 9%. Likewise, the probability that Credit B will default and Credit A will not default is 9%. The probabilities than *either or both* Credit A and Credit B will default, the area within circles A and B, adds up to 19%. Therefore, the probabilities that neither Credit A nor Credit B will default, represented by the area outside circles A and B, is 81%.

Exhibit 14. Default Probabilities, Pictorially

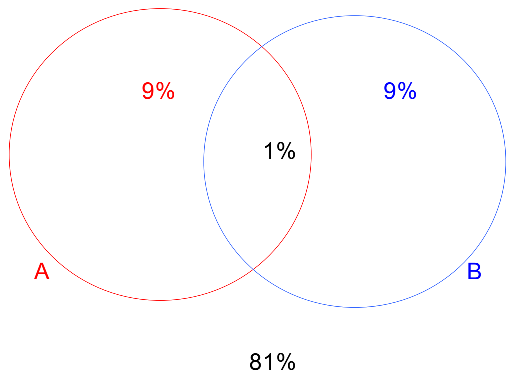

These results are also shown in Exhibit 15 (below) throwing some "nots," "ors," and "neithers" into the notation. P(A not B) means that A defaults and B does not default. P(A or B) means that either A or B defaults and includes the possibility that *both* A and B default. "Neither" means neither A or B defaults.

Exhibit 15. Default Probabilities, Notationally

P(A) = 10%
P(A and B) = 1%
P(A not B) = P(A) – P(A and B) = 10% - 1% = 9%
P(A or B) = P(A) + P(B) – P(A and B) = 10% + 10% - 1% = 19%
P(neither A or B) = 100% - P(A or B) = 100% - 19% = 81%

Exhibit 16. No Joint Probability

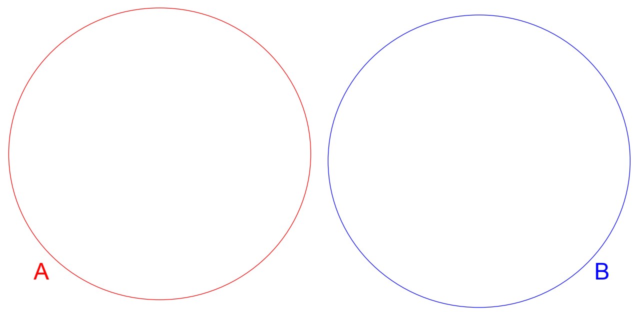

Picturing Default Correlation

We've pictorially covered scenarios of joint default, single default, and no-default probabilities in our two credit world *assuming zero default correlation*. Exhibit 14, showing moderate overlap of the "default circles" has been our map to these scenarios. There are, of course, other possibilities. There could be no overlap, or 0% joint default probability, between Credit A and Credit B, as depicted in Exhibit 16 (above). Or there could be *complete* overlap as depicted in Exhibit 17 (below). The joint default probability equals 10% because we assume that Credit A and B each have a 10% probability of default and in Exhibit 17 they are depicted as always defaulting together. (Note that we draw the circles in Exhibit 17 a little offset so you can see that there are two of them. In theory, they rest exactly on top of each other.)

Exhibit 17. Maximum Joint Probability

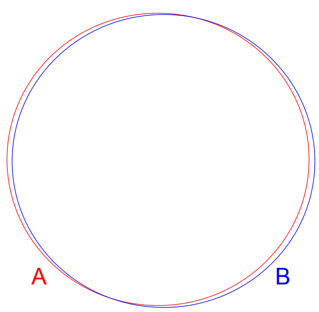

One will recall, hopefully, that Exhibit 16 depicts perfect *negative* default correlation since if one credit defaults we know the other will not. Exhibit 17 depicts perfect *positive* default correlation since if one credit defaults we know the other one will too. Unfortunately, our Equation 1 (previous page) does not take into account the situations depicted in Exhibits 16 and 17. That formula does not help us calculate joint default probability in either of these circumstances or in any circumstance other than zero default correlation. Which leads us to the next section in this appendix.

Calculating Default Correlation Mathematically

We are perhaps already further than most introductory statistics courses would go with respect to correlation. But with Venn diagrams under our belt, we can become more precise in understanding default correlation with a little high school algebra. What we are going to do in this section is mathematically define default correlation. The equation will allow us to compute default correlation between any two credits given their individual default probabilities and their joint default probability. Then, we are going to solve the same equation for joint default probability. The reworked equation will allow us to calculate the joint default probability of any two credits given their individual default probabilities and the default correlation between the two credits.

What we would like to have is a mathematical way to express the degree of overlap in the Venn diagrams or the joint default probability of the credits depicted in the Venn diagrams. From no overlap depicted in Exhibit 16 to "moderate" overlap depicted back in Exhibit 15, to complete overlap depicted in Exhibit 17. One way is refer to the joint probability of default. It's 0% in Exhibit 16, 1% in Exhibit 14, and 10% in Exhibit 17. All possible degrees of overlap could be described via the continuous scale of joint default probability running from 0% to 10%. However, this measure is tied up with the individual credit's probability of default. A 1% joint probability of default is very high default correlation if both credits have only a 1% probability of default to begin with. A 1% joint probability of default is very negative default correlation if both credits have a 50.5% probability of default to begin with. We would like a measure of overlap that does not depend on the default probabilities of the credits.

This is exactly what default correlation, a number running from –1 to +1, does. Default correlation is defined mathematically as:

$$\textbf{Default Correlation(A and B)} = \frac{\text{Covariance(A,B)}}{\text{Standard Deviation(A) * Standard Deviation(B)}} \quad (2)$$

What we are going to do now is to delve more into the right side of Equation (2) and better define default correlation between A and B.

In denominator of the formula, *standard deviation* is a measure of how much A can vary. A, in this case, is whether or not Credit A defaults. What this means intuitively is how certain or uncertain we are that A will default. We are very certain about whether A will default if A's default probability is 0% or 100%. Then we know with certainty whether or not A is going to default. At 50% default probability of default, we are most uncertain whether A is going to default.

The term for an event like default, where either the event happens or does not happen, and there is no in between, is *binomial*. And the mathematical formula for the standard deviation of a binomial variable is:

$$\textbf{Standard Deviation(A)} = \{P(A) * [1-P(A)]\}^{1/2} \quad (3)$$

In the example we have been working with, where the default probability of A is 10%, or P(A) = 10%, the standard deviation of A is:

$$\begin{aligned}\textbf{Standard Deviation(A)} &= \{P(A) * [1-P(A)]\}^{1/2}\\ &= \{10\% * 90\%\}^{1/2}\\ &= \{9\%\}^{1/2}\\ &= 30\%\end{aligned}$$

All the possible standard deviations of a binomial event, where the probability varies from 0% to 100%, are shown in Exhibit 18 (next page). Above 10% probability on the X axis we can see that

Exhibit 18. Standard Deviation and Default Probability

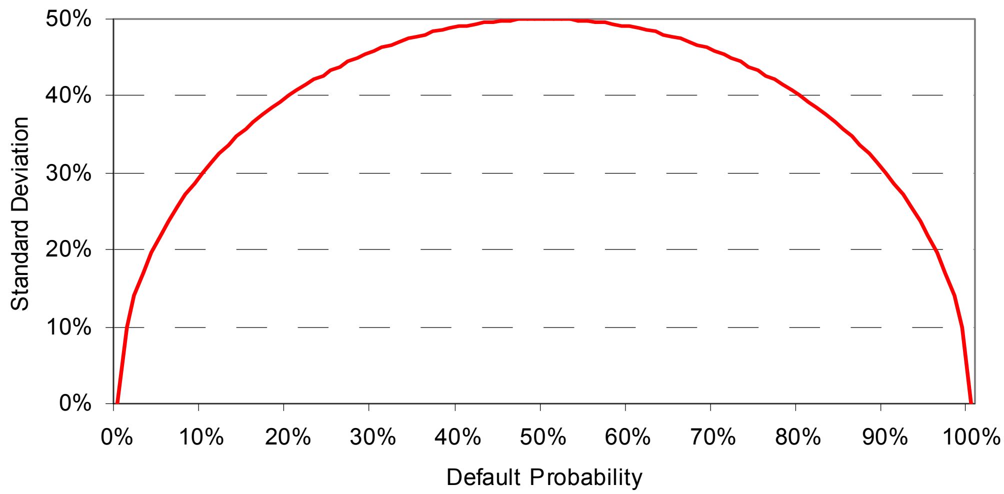

the standard deviation is indeed 30%. The chart also illustrates the statements we made before likening standard deviation to the uncertainty of whether or not the credit is going to default. At 0% and 100% default probability, where we are completely certain what is going to happen, standard deviation is 0. At 50% default probability, where we are least certain whether the credit is going to default, standard deviation is at its highest.

The *covariance* of A and B is a measure of how far the actual joint probability of A and B is from the joint probability that would obtain if there was zero default correlation. Mathematically, this is simply *actual joint probability of A and B* minus the *joint probability of A and B assuming zero correlation*. Recall from Equation 1 above that the joint probability of A and B assuming zero correlation is P(A) * P(B). Therefore the covariance between A and B is: [5]

Covariance(A, B) = P(A and B) – P(A) * P(B) (4)

In our example, from our work around Exhibit 14, we worked out that the joint probability of default assuming zero default correlation is 1%. From Exhibit 16, we know that given perfect negative default correlation, the actual joint probability can be as small as 0%. From Exhibit 17, we know that given perfect positive default correlation, the actual joint probability can be as high as 10%. Exhibit 19 (next page) depicts the relationship between joint default probability and covariance graphically.

Recall that a few pages back in Equation 2 we presented this formula for default correlation.

[5] Covariance is more formally defined as the Expectation(A*B) – Expectation(A) * Expectation(B). When one defines default as 1 and no default as 0, Equation 4 is the result.

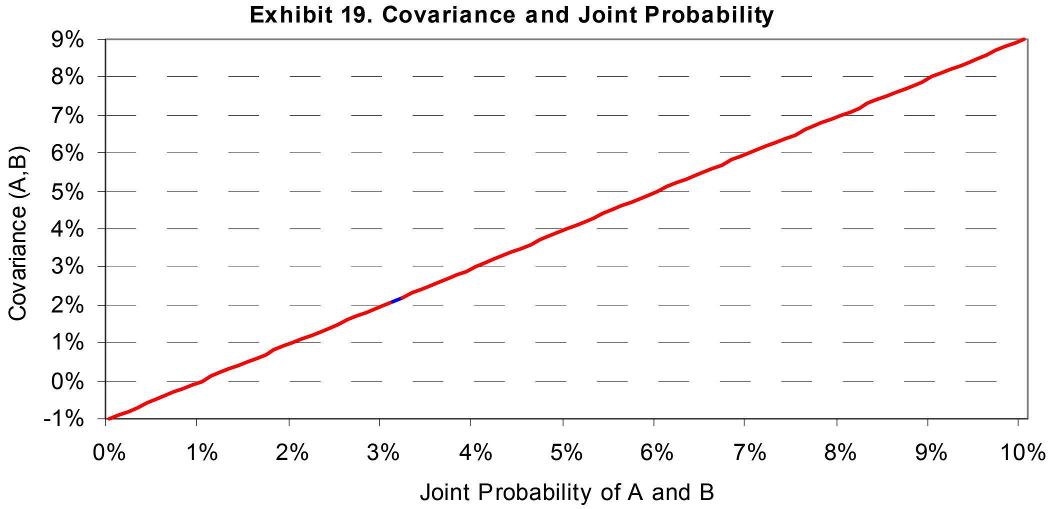

$$\textbf{Correlation(A and B)} = \frac{\text{Covariance(A,B)}}{\text{Standard Deviation(A) * Standard Deviation(B)}} \quad (2)$$

In Equations 3 and 4, we refined the numerator and denominator of Equation 2. Substituting Equations 3 and 4 into Equation 2 we have:

$$\textbf{Correlation(A and B)} = \frac{P(A\text{ and }B) - P(A) * P(B)}{\{P(A) * [1-P(A)]\}^{1/2} * \{P(B) * [1-P(B)]\}^{1/2}} \quad (5)$$

Now, finally, we can define mathematically the default correlation we saw visually in Exhibits 14, 16, and 17. In Exhibit 14, the joint default probability of A and B, P(A and B), was 1%, simply because we wanted to show the case where the default probability of one credit did not depend on whether another credit had defaulted. The product of A's and B's default probabilities, P(A) * P(B), is 10% * 10% or 1%. Moving to the denominator of Equation 5, the product or A's and B's standard deviations, $\{P(A) * [1-P(A)]\}^{1/2} * \{P(B) * [1-P(B)]\}^{1/2}$ is 9%. Putting this all together, we have

$$\textbf{Correlation(A and B)} = \frac{1\% - 1\%}{9\%} = 0.00$$

Similarly, for Exhibit 16, where joint default probability is 0%, default correlation is -0.11. And in Exhibit 17, where joint default probability is 10%, default correlation is +1.00. In our example, as joint default probability moves from 0% to 10%, default correlation decreases linearly from –0.11 to +1.00, as shown in Exhibit 20 (next page).

Exhibit 20. Default Correlation and Joint Default Probability

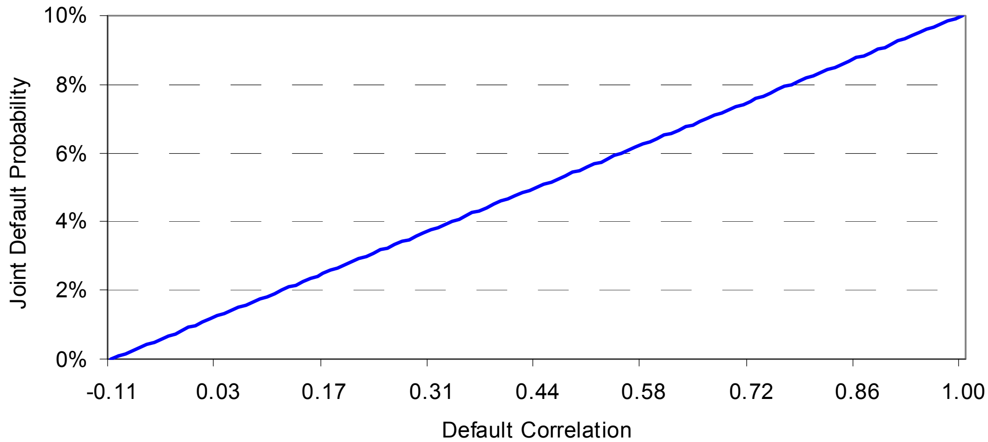

Theoretically, correlation can range from –1.00 to +1.00. But for binomial events like default, the range of possible default correlations is dictated by the default probabilities of the two credits. With 10% probability of default for both credits, the possible range of default correlation is reduced to the range from –0.11 to +1.00. If both credits do not have the same default probability, they can not have +1.00 default correlation. Only if the default probability of both credits were 50% would it be mathematically possible for default correlation to range fully from –1.00 to +1.00. Exhibit 21 (below) shows the relationship between default correlation and joint default probability when the individual default probability of both credits is 50%, when the individual default probability of both credits is 10%, and when the individual default probabilities of credits are 10% and 50%, respectively.[6]

Exhibit 21. Default Correlation, Joint Default Probability and Underlying Default Probability

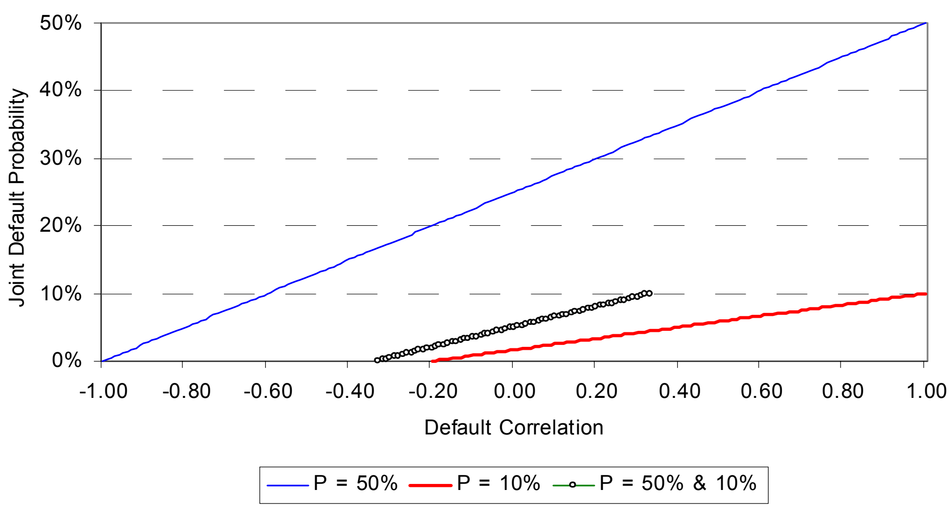

[6] These restrictions on the range of default correlation are a surprise to people used to dealing with the correlation of continuous variables. In fact, we've seen mistakes made on this point by people and firms that hold themselves out to be experts in default correlation, even among those who cite our 1995 paper as a reference.

Note that as described, default correlation in the case where the default probability of both credits is 50% ranges from +1 to -1. Also, note the slope of the lines. The same increase in default correlation has a bigger effect on the joint probability of default when individual default probabilities are 50% than when individual default probabilities are 10%, or 10% and 50%.

We will see in Appendix II that Equation 5 allows us to calculate historic default correlations from empirical default data. Rearranging Equation 5 to solve for the joint probability of default, we have a formula that will allow us to calculate the joint default probability of A and B given their individual probabilities of default and their default correlation.

P(A and B) =
$$\text{Correlation(A and B)} * \{P(A) * [1-P(A)]\}^{1/2} * \{P(B) * [1-P(B)]\}^{1/2} + P(A) * P(B) \qquad (6)$$

Default Correlation in a Ménage à Trois

As this point, readers who are familiar with the concept of correlation for continuous variables like stock returns or interest rates are apt to find some surprises. Perhaps you are used to looking at portfolio risk in the context of Harry Markowitz's portfolio theory and covariance-variance matrixes.[7] In that framework, if you know the standard deviation of each variable, and the correlation of each pair of variables, you can explain the behavior of the entire portfolio. Not so with a binomial variable like default. We illustrate the difference in this section.

Instead of the two credit world we have been focused on, let's suppose we have three credits, A, B, and C, each with a 10% probability of default. Let us suppose that the default correlation between each pair of credits is 0.00. As we have discussed before, around Exhibit 14, this means that the joint default probability between each pair of credits is 1%. We illustrate this situation in Exhibit 22 (below).

Exhibit 22. Pairwise Joint Default Probabilities

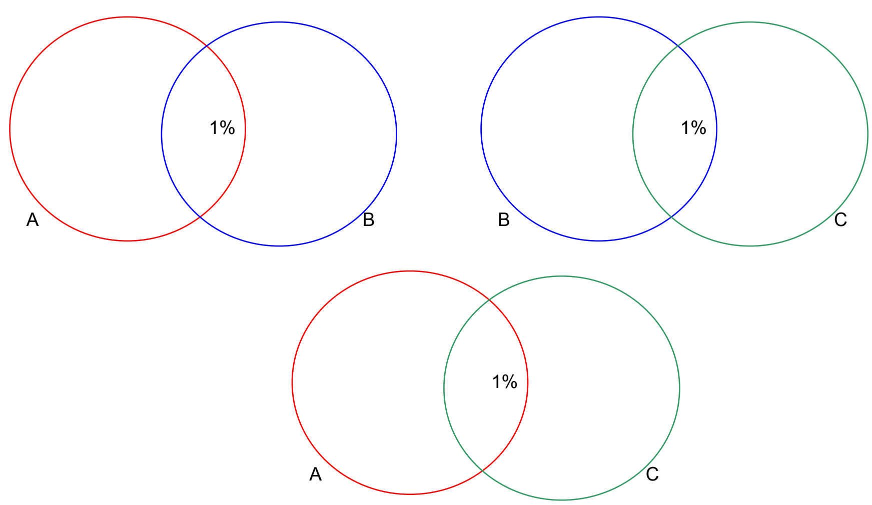

[7] H.M. Markowitz, "Portfolio Selection," *Journal of Finance*, March 1952. But really any introductory finance text.

Exhibit 23. Zero Pairwise and Positive Three-Way Default Correlation

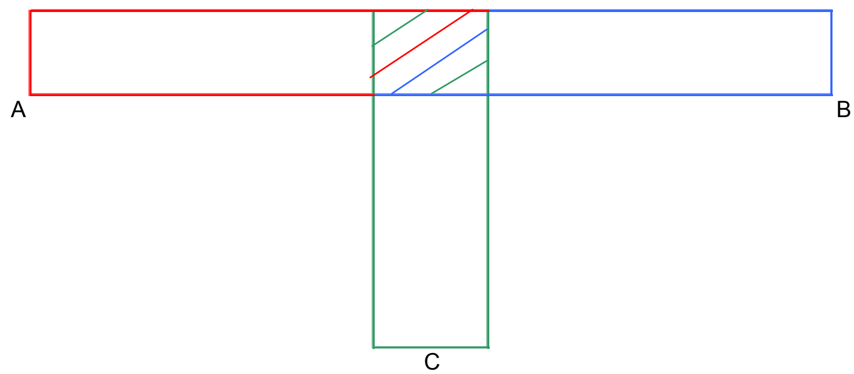

Now we are eager to understand the behavior of all three credits together. We seem to have a lot of information: each credit's default probability and the default correlation between each pair of credits. What does this tell us about how defaults will occur among all three credits? Not much, it turns out. Exhibit 23 and 24 show the extremes of possible outcomes. (Note the switch from circles to rectangles and ovals in these Venn diagrams to show the overlapping probabilities clearly.) In Exhibit 23, whenever two credits default, the third credit joins them in default. This sounds like positive default correlation: if you know that any two credits have defaulted, your estimate of the probability of the third credit defaulting is 100%. But it occurs while all *pairwise* default correlations are zero.

In Exhibit 23 (above), there is no situation where only *two* credits default. Exhibit 25 shows the probabilities of all possible default outcomes under the heading "Positive Three-Way Default Correlation." There is a 1% probability that all three credits default, 0% probability that two credits default, 27% probability that one credit will default and 72% probability that no credits will default.

In Exhibit 24 (below), in contrast, there is no situation where all *three* credits default. In this case, if you know that two credits have defaulted, you are sure that the third will not default. Exhibit 25 (next page) shows the probabilities of all possible default outcomes under the heading "Negative Three-Way Default Correlation." There is a 0% probability that all three credits default, 3% probability that two credits default, 24% probability that one credit will default and 73% probability that no credits will default.

Exhibit 24. Zero Pairwise and Negative Three-Way Default Correlation

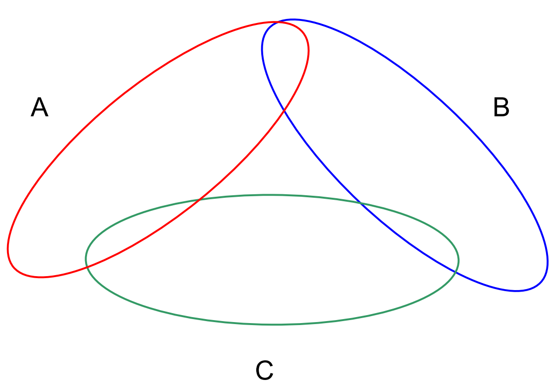

Note that the expected number of defaults in each three-way correlation scenario is the same. In the positive three-way correlation scenario, the

Exhibit 25. Three-Way Default Correlation and Resulting Default Probabilities

Number of Defaults	Zero Pairwise and Positive Three-Way Default Correlation	Zero Pairwise and Negative Three-Way Default Correlation
0	72%	70%
1	27%	24%
2	0%	3%
3	1%	0%

expected number of defaults in the portfolio is 27% * 1 + 1% * 3 or 0.3. In the negative three-way correlation scenario, the expected number of defaults in the portfolio is 24% * 1 + 3% * 2 or also 0.3. Note also that the probability of any two pairs of credits defaulting at the same time is 1%. In the negative three-way correlation scenario, defaults of pairs AB, BC, and AC each have a 1% chance of occurring, which adds up to a 3% chance of two credits defaulting. In the positive three-way correlation scenario, the probability of *all* three credits defaulting at the same time is 1%. This also means that defaults of pairs AB, BC, and AC each have a 1% chance of occurring. This is proof that pairwise default correlation is 0.00. But the sad truth of the matter, for those who thought differently, is that knowing pairwise default correlations does not tell you everything you would like to know about the behavior of this three credit portfolio. This is in direct contradiction to the experience many people have had with continuous variables.[8]

Pairwise and Three-Way Default Correlation in Nth to Default Swaps

In Exhibit 26 (below) we address pairwise and three-way default correlation in the context

Exhibit 26. Pairwise and Three-Way Default Correlation in a 1st to Default Swap

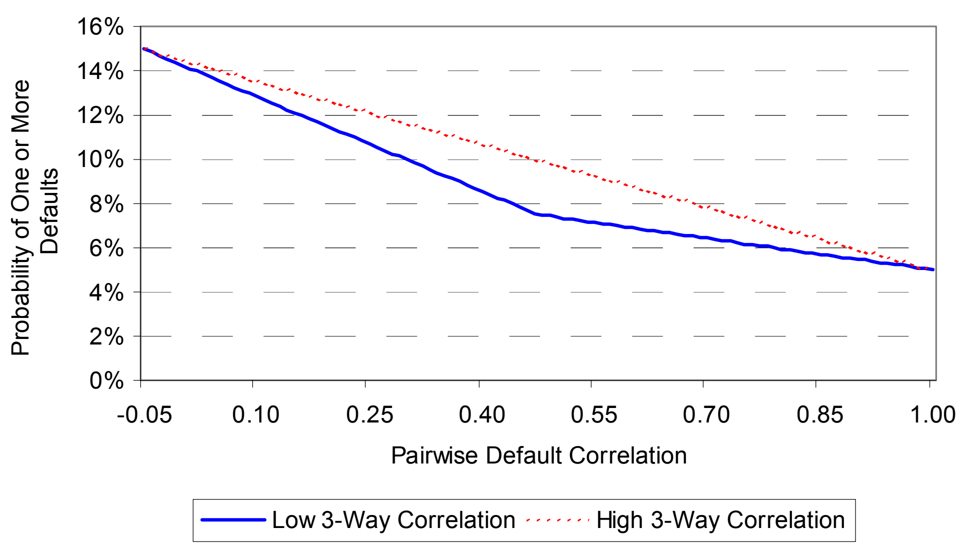

[8] This is a point we have never seen mentioned outside our 1995 paper.

Exhibit 27. Effect of Three-Way Default Correlation, Given Pairwise Default Correlation = +0.47

	Three-Way Default Correlation = -0.04	Three-Way Default Correlation = 0.70
Probability of 0 Defaults	92.45%	89.97%
Probability of 1 Defaults	0.11%	7.55%
Probability of 2 Defaults	7.44%	0.00%
Probability of 3 Defaults	0.00%	2.48%
Probability of 1 or More Defaults	7.55%	10.03%

of a 1st to default swap on three underlying credits. We assume each credit has a 5% default probability.

In the exhibit, we show the probability of at least one credit defaulting in the reference portfolio as pairwise default correlation increases from -.0.05 to 1.00 under two different assumptions about three-way correlation. The first is that three-way default correlation is as low as it can be, consistent with pairwise correlation. The second is that three-way default correlation is as high as it can be, again consistent with pairwise correlation.

It turns out that three-way default correlation makes a little bit of a difference in 1st to default risk, especially as pairwise default correlation increases. The probability that at least one credit will default is always greater when three-way default correlation is high. The biggest difference, 2.5%, comes when pairwise default correlation is 0.47. The effect of high three-way default correlation is to increase the probability that at least one reference credit will default. But wait, way back around Exhibit 10 we said that higher correlation reduces 1st to default risk, the probability that at least one credit will default. Exhibit 27 (above) shows what is going on.

When pairwise default correlation is 0.47 and three-way default correlation is –0.04, the probability of two defaults is 7.44% and the probability of three defaults is zero. What is happening here is that defaulting pairs of credits AB, BC, and AC each have a 2.48% default probability. When three-way default correlation is 0.70, all those defaulting pairs are scrunched into a sort of default ménage à trois ABC. To compensate for the loss of overall defaults, the probability of a single default increases. The situation is identical to that pictured in Exhibits 23 and 24.

We said without explanation that in creating Exhibit 26 we picked three-way correlation to be as high or as low as it can be “consistent with pairwise correlation.” This bears some explanation. Given a particular pairwise correlation value, the range of three-way correlation is restricted. For example, if pairwise correlation is 0.00, three-way correlation can range from –0.01 to 0.22. If pairwise correlation is 1.00, three-way correlation can only be 1.00. Exhibit 28 (next page) shows the range of three-way default correlation versus pairwise default correlation.

Around Exhibit 11 we said that default correlation made more of a difference in 2nd to default swaps and now is the time to see whether this holds true for three-way correlation. Exhibit 29 (next page) shows the probability of two or more defaults out of three reference credits given that each reference credit has a 5% probability of default. Three-way default correlation makes a great deal of difference in 2nd to default risk.

Exhibit 28. Possible Range of Three-Way Default Correlation Given Pairwise Default Correlation

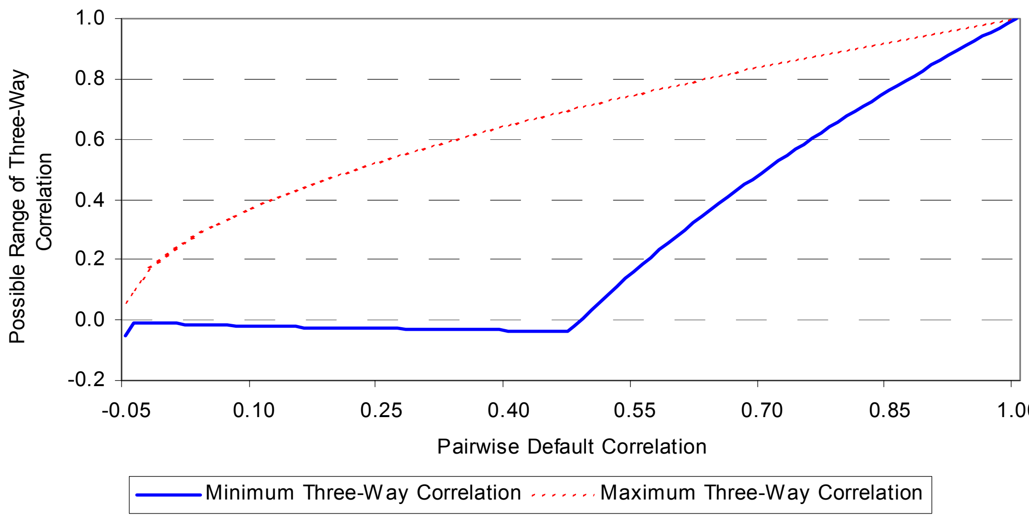

Exhibit 29. Pairwise and Three-Way Default Correlation in a 2nd to Default Swap

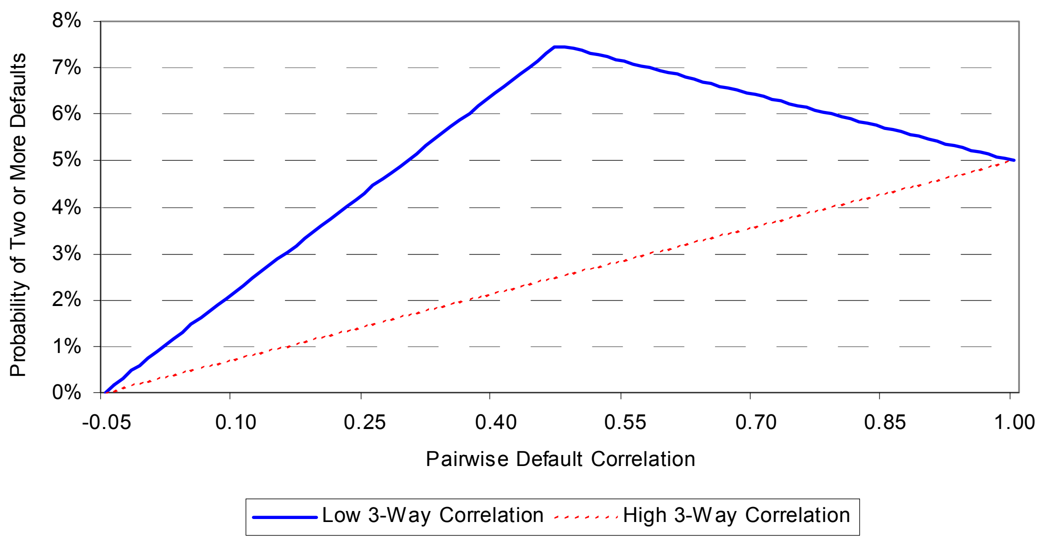

Default Correlation in a Crowd

We wanted to make sure that higher orders of default correlation were also important for large portfolios. So we considered a 100-credit portfolio where each credit has a 10% probability of default. We computed the probabilities of zero to 100 credits defaulting under three correlation scenarios:

- zero pairwise default correlation and zero higher correlations;

Exhibit 30. Default Probabilities in a 100 Credit Portfolio

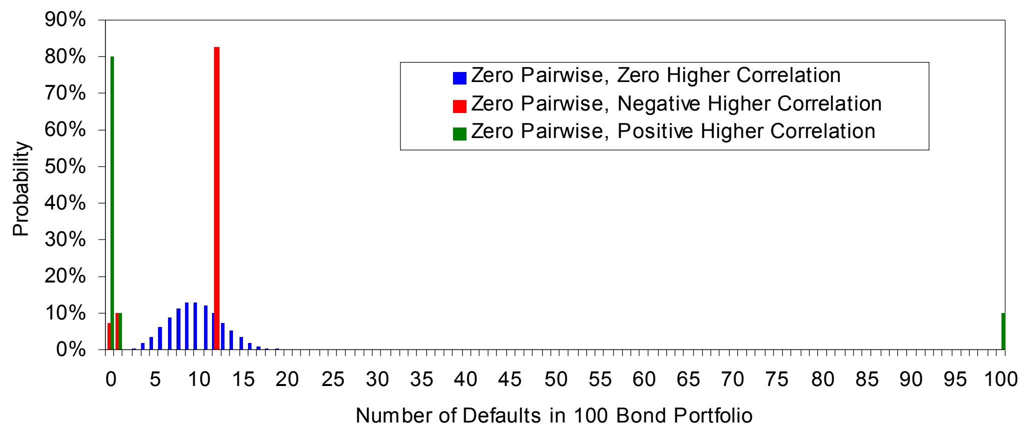

- zero pairwise default correlation and maximum negative higher correlations;
- zero pairwise default correlation and maximum positive higher correlations.

The results are shown in Exhibit 30 (above) and show the extreme distribution of the positive higher correlation portfolio and the very stable distribution of the negative higher correlation portfolio relative to the zero higher correlation portfolio. Again, our conclusion is that pairwise default correlations do not give us all the information we need to understand the behavior of a portfolio.

Intuitively, increasing higher-level default correlation seems logical. Assuming that positive pairwise default correlation exists, the first default in the portfolio will cause us to revise our estimation of the default probability of remaining credits in the portfolio upwards. It seems logical that if a *second* credit defaults, we would want to again revise our estimation of the default probabilities of remaining credits upwards. This is the effect of higher order positive default correlation. As more and more credits default, we think it more likely that remaining credits will also default.

Appendix II. Empirical Default Correlations and Their Problems

With enough data, and some assumptions, we can calculate *historic* default correlations. These correlations might be useful in predicting the performance of portfolios. We are first going to show how these calculations are made, and then circle around in the next section to discuss the appropriateness of the underlying assumptions that go into calculating them. We go back to Equation 5, which we derived in Appendix I and show again here:

$$\textbf{Correlation(A and B)} = \frac{P(A \text{ and } B) - P(A) * P(B)}{\{P(A) * [1-P(A)]\}^{1/2} * \{P(B) * [1-P(B)]\}^{1/2}} \quad (5)$$

To compute, say, the default correlation of two B-rated companies over one year, we set P(A) and P(B) equal to the historic one-year default rate for B-rated companies. The remaining variable on the right side of Equation 5 is the joint probability of default, P(A and B). We compute P(A and B) by first counting the number of companies rated B at the beginning of a year that subsequently defaulted over that particular year. We then calculate all possible *pairs* of such defaulting B-rated companies. If X is the number of B-rated companies defaulting in a year, the possible pairs are [X * (X –1)]/2.

We next calculate all possible pairs of B-rated companies, whether or not they defaulted, using the same formula, [Y * (Y –1)]/2, where Y is the number of B-rated companies *available* to default. The joint default probability of B-rated companies in a particular year is:

$$\frac{[X * (X - 1)] / 2}{[Y * (Y - 1)] / 2}$$

The average of this statistic is taken over available years in the dataset. Having all the terms on the right hand of Equation 5, we can solve for the default correlation between Credits A and B. In a similar manner, it is possible to calculate default correlations over longer periods and between groups of credits of different ratings, for example the default correlation between Aa and Ba credits over five years.

In our 1995 paper on default correlation[9], we computed default correlations between all combinations of Moody's rating categories for time periods from one to ten years. The data we used included 24 years of default data covering the years 1970 through 1993, including industrial companies, utilities, financial institutions, sovereign issuers, and structured finance entities. To our knowledge, an update of the historic default correlations presented in that paper has never been published, so we reproduce our old results in Exhibits 31A (page 30) and B (page 31).
Our conclusions from studying the results in Exhibit 31A and B are:

- default correlations increase as ratings decrease;
- default correlations initially increase with time and then decreases with time.

We guess that default correlation increases as rating decreases because lower-rated companies are relatively more susceptible to problems in the *general economy* while higher-rated companies are relatively more susceptible to *company-specific* problems. Low-rated companies, being closer to default already, are more likely to be pushed into default because of an economic downturn. As

[9] "Default Correlation and Credit Analysis," *Journal of Fixed Income*, March 1995.

Exhibit 31A. Historic Default Correlations

	Aaa	Aa	A	Baa	Ba	B
One-Year Default Correaltions * 100						
Aaa	0					
Aa	0	0				
A	0	0	0			
Baa	0	0	0	0		
Ba	0	0	0	0	2	
B	0	1	0	1	4	7
Two-Year Default Correaltions * 100						
Aaa	0					
Aa	0	0				
A	0	0	0			
Baa	0	0	0	0		
Ba	0	1	1	1	6	
B	0	1	1	2	10	16
Three-Year Default Correaltions * 100						
Aaa	0					
Aa	0	0				
A	0	1	1			
Baa	0	0	0	0		
Ba	0	2	2	1	9	
B	0	2	3	3	17	22
Four-Year Default Correaltions * 100						
Aaa	0					
Aa	0	0				
A	0	1	1			
Baa	0	1	1	0		
Ba	0	2	3	3	13	
B	0	2	4	5	22	27
Five-Year Default Correaltions * 100						
Aaa	0					
Aa	0	0				
A	0	1	1			
Baa	0	1	1	0		
Ba	0	3	4	3	15	
B	0	4	6	7	25	29

economic conditions affect all low-rated credits simultaneously, defaults among these credits are likely to be correlated. In contrast, defaults of highly-rated companies, besides being rare, are typically the result of company-specific problems. As these problems are by definition isolated to individual credits, they do not produce default correlation.

With respect to default correlation increasing and then decreasing with the time period studied, we note that default correlations peak at five and six year periods for rating pairs Baa/Ba, Baa/B, Ba/Ba, and Ba/B. However, default correlations peak at nine years for rating pairs A/B and B/B.

Exhibit 31B. Historic Default Correlations

	Aaa	Aa	A	Baa	Ba	B
Six-Year Default Correaltions * 100						
Aaa	0					
Aa	1	1				
A	1	1	1			
Baa	0	1	1	0		
Ba	1	3	4	3	15	
B	1	4	7	7	25	29
Seven-Year Default Correaltions * 100						
Aaa	0					
Aa	1	0				
A	1	1	2			
Baa	0	1	1	0		
Ba	2	2	4	3	13	
B	3	3	9	8	24	30
Eight-Year Default Correaltions * 100						
Aaa	1					
Aa	1	0				
A	1	1	2			
Baa	1	1	1	0		
Ba	3	3	5	2	10	
B	6	5	11	7	23	37
Nine-Year Default Correaltions * 100						
Aaa	1					
Aa	1	0				
A	2	1	2			
Baa	2	1	1	0		
Ba	4	3	5	2	8	
B	8	6	12	6	20	39
Ten-Year Default Correaltions * 100						
Aaa	1					
Aa	2	1				
A	2	2	2			
Baa	2	1	1	0		
Ba	4	3	4	2	8	
B	9	8	9	6	17	38

We note that over arbitrarily short time periods, defaults are necessarily uncorrelated. Imagine a database whose column headings are the names of credits and whose rows represent time intervals, perhaps four-year intervals. The entry in a particular cell is 1 if the credit defaulted in that time interval and 0 otherwise. Ones in the same row indicate that credits defaulted together in that time interval and indicate the presence of positive default correlation. But if a shorter time period is used, fewer ones will appear in the same row, lowering perceived default correlation. At some arbitrarily short period of time, no more than one 1 will appear in a row and there will be no evidence of positive default correlation.

The decrease in default correlation that occurs in most rating categories over longer time periods may be caused by the relationship of the time period being studied to the average business cycle.

Exhibit 32. Simulated Annual Default Rates Under Different Default Correlations

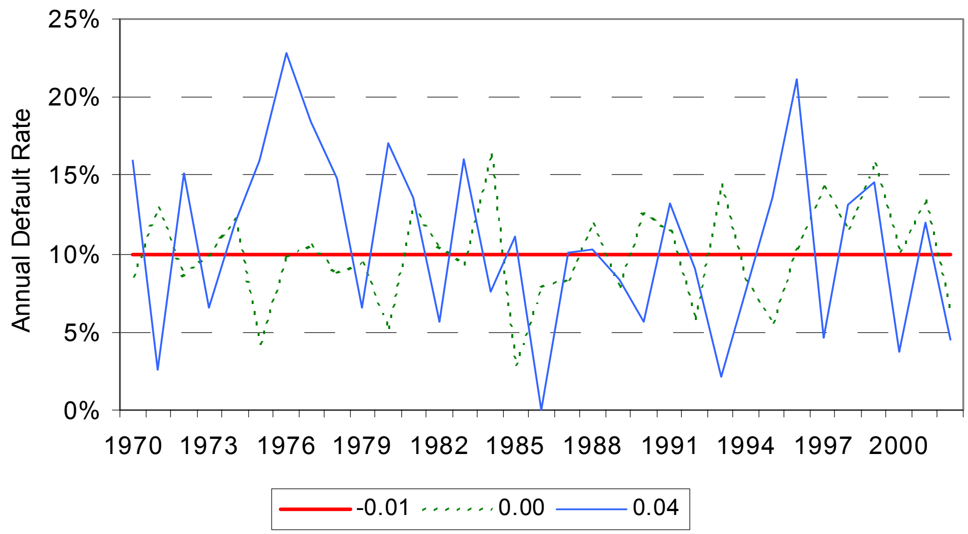

If the time period studied covers the entire ebb and flow of the business cycle, defaults caused by general economic conditions average out over the period, thus lowering default correlation. We think that default correlation is maximized when the time period tested most closely approximates the length of an economic recession or expansion.

Just as the pairwise default correlations in Exhibit 31 can be calculated, so too can higher order default correlations. But rather than work out the mathematics of how this would be done, we instead turn to a discussion of the reliability of empirically-observed default correlations.

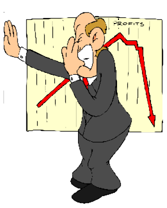

Problems with Historical Default Correlations

Implicit in the previous section on empirical default correlation is the idea that wide swings in default rates are an indicative of positive default correlation while small swings or steady default rates are indicative of low or even negative default correlation. We illustrate this concept explicitly in Exhibit 32 (above).

Exhibit 32 depicts simulated default rates for three 100-credit portfolios assuming 10% default probability for each credit and pairwise default correlations of -0.01, 0.00, and 0.04, respectively. The red line, steady at exactly 10%, is produced with perfect negative default correlation, in this case –0.01. The green line that ranges between 3% and 16% was produced with 0.00 default correlation. Finally, the most volatile series, the blue line, which varies between 0% and 23%, was produced with default correlation of 0.04. This shows that a little bit of default correlation can cause great swings in default probabilities.

There is, however, an alternative interpretation of the positive default correlation series. Instead of defaults being correlated, it could be the fact that default probability changes over time. For example, the 23% default rate in 1976 could have been the result of the true probability of default increasing,

Exhibit 33. Time Correlated Default Rates

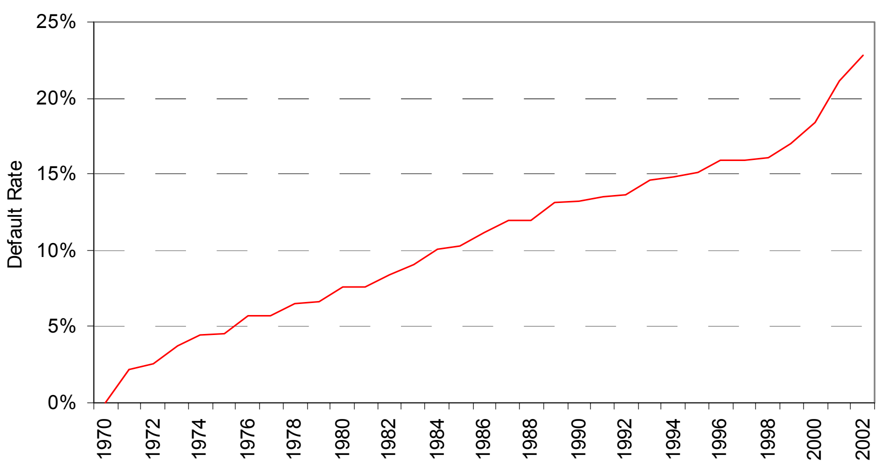

for some reason, into the low 20s. Given that revised default probability, defaults in 1976 might have been zero correlated.[10]

This simple and plausible story about fluctuating default probabilities puts into question all the work deriving empirical default correlations by rating category done in the previous section. In fact, it puts into question all consideration of default correlation. We cannot be sure whether the variability in default rates from year to year or over longer periods is due to default correlation or changing default probability.

Pragmatic scrutiny of credit ratings and the credit rating process suggests to us that ratings are more *relative* than *absolute* measures of default probability and that default probabilities for different rating categories change year-to-year. It is a hard enough job to arrange credits in an industry in order of credit quality. It seems to us very difficult to assess credit quality against an absolute measure like default probability and then calibrate these measures across different industries. In fact, the rating agencies themselves say that ratings are relative measures of credit quality.[11]

If ratings are relative measures of credit quality, or if for any reason the probabilities of default for different rating categories change over time, this would mean that the historically-derived default correlations presented in Exhibit 31 overstate true default correlation. But they are already low to begin with, especially in the Baa and A rating categories and over three to five year periods that correspond to the terms of most n^{th} to default swaps.

Here's another perspective on the interaction between default probability and default correlation. In Exhibit 33 (above), we have rearranged the annual default rates of the positively correlated series in Exhibit 32 so that the default rates are in strict order from lowest to highest. In our calculation of default correlation, the order of default rates does not make a difference to default correlation. This series would still have default correlation of 0.04.

---

[10] Way back in Exhibit 7, we showed the variability in annual corporate default rates since 1920 and we said that variability was evidence of default correlation. But it also seems logical that credit analysts in 1934 and 1952 would have had vastly different expectations of annual defaults.

[11] Jerome Fons, Richard Cantor, and Christopher Mahoney, *Understanding Moody's Corporate Bond Ratings and Rating Process*, Moody's Investors Service, May 2002.

Over the time period, it is true that the average annual default rate is 10%. But who, looking at this time series, could not think up a simple rule to better explain and predict default rates: defaults next year will be what they were last year. Yet, our method of calculating default correlation would not pick up the "memory," or *time series correlation*, of default rates.

A final confounding issue regarding empirical default correlation is the possibility that default correlation itself can vary. It could be the case that two different default correlations exist among assets depending upon whether the economy is in "normal" or "crisis" times. In crisis times, default correlation might take on more extreme values. Then the problem of measuring and using default correlation to predict default distributions becomes even harder as one struggles to identify the probabilities of normal and crisis times and the default correlation in each condition. ♦

Douglas Lucas
Director, CDO Research
Office: 212-713-3440
Mobile: 917-968-8487
Email: douglas.lucas@ubsw.com

Analyst Certification

Each research analyst primarily responsible for the content of this research report, in whole or in part, certifies that with respect to each security or issuer that the analyst covered in this report: (1) all of the views expressed accurately reflect his or her personal views about those securities or issuers; and (2) no part of his or her compensation was, is or will be, directly or indirectly, related to the specific recommendations or views expressed by that research analyst in the research report.