Securities and Exchange Commission
100 F Street, NE
Washington, DC 20549-1090

February 13, 2007

Re: File Number S7-04-07 Comments on SEC Proposed Rules and Oversight of NRSROs

The following comments are not meant to be a comprehensive review of the SEC's proposal. Rather, I highlight specific issues to promote further thought in the general direction and soundness of the SEC's next steps.

SEC Interpretation of the Role of the Rating Agencies

In order to comment on the Securities and Exchange Commission's report, it is first important to review the rating agencies' role, since it seems that even former employees of the SEC misunderstand it. In his report [Bank of America, N.A., successor by merger to NationsBank, N.A. et al. v. William R. Bartmann, et al.; Case No. C-J-2000-02274], Richard Breeden, former head of the SEC, opined that the rating agencies' inability to identify the fraud [on the part of Commercial Financial Services, Inc.] "provides further compelling evidence that…failure to discover it cannot fairly be attributed to negligent or otherwise commercially unreasonable conduct..." Mr. Breeden's opinion is inconsistent with the views of the industry and the rating agencies themselves. It is well known in the industry that the rating agencies are not auditors or investigators and are not responsible for unearthing fraud.

Rating agencies do not perform due diligence for investors; they merely provide an opinion. Independent organizations exist that will perform rigorous reviews and audit tests for placement agents, and their reviews go well beyond what rating agencies will do. Furthermore, in cases in which securitization legal disputes have arisen, rating agencies have successfully claimed

journalist-like privileges to avoid turning over notes and analyses produced during the rating process.

The rating agencies are Nationally Recognized Statistical Rating Organizations (NRSROs). One would expect that the rating agencies can be relied upon to apply generally recognized statistical concepts when assigning ratings based on the data presented to them, but that is not always the case. Notching, arbitrarily lowering the rating of securities rated by other rating agencies in pools of securitized assets, is one example. The NRSROs do not wish to recognize each others ratings. Far from an attempt to employ "nationally recognized statistical" methods, notching is simply arbitrary junk science. There is little accountability or independent outside auditing of the soundness or consistency of rating agency methodologies with respect to statistical principles.

The rating agencies' role with respect to rating securitizations has become a necessary tick on a checklist and serves to facilitate the investment banking process of bringing deals to market. The robustness of securitization ratings from an investor's point of view is unreliable across asset classes and at times even across deals. Furthermore, ratings of synthetic securitizations employing leverage are nearly meaningless due to the fact that the primary risks are market risks influenced by supply and demand forces, and credit ratings are wholly insufficient to evaluate that risk. Investment grade ratings – even "triple A" investment grade ratings such as that awarded Long Term Capital Management – should not lull the SEC, investors and other financial industry professionals into an unwarranted sense of safety.

Since at least one NRSRO rating is often a requirement for investors, to the extent that the SEC can help promote sound statistical methodology, it would be a welcome step in the right direction.

Comments on Ratings of Securitized Assets

I applaud the SEC's attempt to address notching, but I do not applaud its proposed alternative of allowing a rating agency to arbitrarily refuse to rate a securitization if it rates less than 85% of the market value of the underlying assets. The SEC proposes to replace one arbitrary statistically unsound practice with another arbitrary unsound practice. I do not have any alternative proposal at the moment, but I must also point out, that the SEC has not put forth a credible one.

While notching is a problem, there is much bigger problem that the SEC hasn't even addressed on the topic of rated securitizations. There has been explosive growth in synthetic securitizations swamping cash issuance. The ratings methodologies for some of these products are so sloppy, that I reject the ratings of all three "nationally recognized" rating agencies, and encourage others to do the same. Furthermore, tranches of securitizations rated using flawed methodologies are themselves used as collateral in other securitizations causing the errors to compound and spiral out of control. Ratings on these products are based on smoke and mirrors.

Beyond the issues mentioned above, ratings on leveraged synthetic credit products are often misleading. A recent example is the "triple A" rating achieved by products like the constant proportion debt obligation (CPDO), which is largely a leveraged bet on the credit quality and market spreads of indexes based on US and European investment-grade companies. Credit rating agencies are not well qualified to assess the market risk combined with leverage presented by these products. Besides credit risk, there is a substantial amount of financial engineering risk.

Late last year, ABN Amro launched constant proportion debt obligations, or CPDOs. After rating the ABN Amro CPDO transaction "triple A," Moody's was criticized by industry professionals, including me. Moody's then changed its rating

methodology applying a different standard for subsequent transactions. Since then Lehman Brothers, Calyon and others have launched their own look-alike products. Standard & Poor's indicated that new CPDO deals paying LIBOR + 100bp (early deals were targeting LIBOR + 200 bps) would be rated “triple A.”

CPDOs generate cash by taking in premiums for selling protection on investment grade US and European credit default swap indexes, the iTraxx Europe and the Dow Jones CDX. CPDOs employ leverage, usually a maximum of 15 times, in order to boost their income. Without these products, around $50 billion of contracts are traded on these indexes, but soon demand created by CPDOs will more than double that volume. New products are in the works that will reference other asset classes.

There are several challenges with this product since excess spread above what is needed to pay the promised coupon is needed to support the “triple A” rating. This excess reserve has to absorb defaults, if any. Default risk is mitigated, since deteriorating credits are kicked out of the indexes and the exposure is rolled over every six months. But there is also the potential of mark-to-market losses when the sold protection is rolled over into a new index series every six months. Furthermore, many investment banks charge around 2 bps each time this roll occurs. Since many CPDO managers will need to roll their trades at specific times, there is the potential for front-running by hedge funds and investment banks.

Potential losses are due to defaults or market value changes (when spreads widen). The high leverage puts investors’ principal at risk, since it acts as first-loss protection on the leveraged exposure to the indices. CPDOs have an extreme amount of mark-to-market and liquidity risk. The ratings volatility is likely to be very high. Potentially, investors may hold a 10-year maturity bond levered fifteen times to a five-year index. If a reference credit were to default prior to being replaced (recall the fast deterioration of names like Enron and Worldcom), however, the

recovery value on CPDO would probably be much lower than the recovery value on a traditional "triple A" rated bond, and therefore it presents more risk.

These products have played a significant role in driving spreads on credit derivative indices to levels so low that it is often possible to put on a profitable negative basis trade: for example, buy a bond and pay a premium on a credit derivative protecting against default while locking in a profit.

An unintended consequence of CPDOs is that fixed income pension funds that used bought protection on CDS (credit default swap) index products in 2006 were harmed. Many funds shorted a CDS index (bought credit protection) and took a long position in swap spreads to hedge their bond portfolio credit risk. Huge inflows from CPDO investors (protection sellers) caused these spreads to narrow for reasons unrelated to credit risk, thus destroying the utility of this hedge for pension funds. CDS indexes were originally hyped as hedge instruments for credit risk, but this new development turned that theory on its head.

"Triple A" tranches of other synthetic CDO deals are often not what they seem, and many do not technically deserve this rating. I attach an article I authored for *The Journal of Structured Finance (Winter 2006, Volume 11, Number 4)*. Please refer to the first three pages.

Sincerely,

Janet M. Tavakoli

Date: February 13, 2007

Janet M. Tavakoli
President
Tavakoli Structured Finance, Inc.

The Elusive Income of Synthetic CDOs

By Janet Tavakoli

Published in the *Journal of Structured Finance* [Winter 2006 Volume 11, Number 4]

If you get 256 people into a room and give them each a coin to flip, the odds are that half of them – 128 – will flip heads on the first try. That is the object, you tell them: to flip heads. Of those, 128 winners, 64 will flip heads on the next go-round as well. Twice running. Not bad. Thirty-two people will flip heads three times in a row, sixteen will flip heads four times in a row, eight will flip heads five times in a row, four will succeed six times in a row, two will rack up an incredible seven straight successes, and one-one out of all 256 in the crowd – will flip heads eight times in a row. What talent! What genius!

What nonsense. This man is no more or less likely than anyone else in the room to flip heads the ninth time. His chances are fifty-fifty. He is not a genius, he is a statistic in a probability formula. As is some other man in the crowd of 256 (who may actually be a genius) who, odds are, failed eight times running to flip heads.

EXPLOSIVE GROWTH

Investment grade synthetic collateralized debt obligations (CDOs) will continue to gain market share, because of the tremendous boost credit derivatives give to the CDO "arbitrage." Credit derivatives add flexibility, but investors must invest time in education, or they will invest for less relative value than they enjoy when they invest in more conventional instruments.

Historically, most stand alone credit derivatives trading desks could not generate much revenue for the banks that housed them. This was because much of the trading inquiry was for protection, the products were illiquid, there were a limited number of trading partners, the trading desks had limited credit lines, and the trading desks had difficulty matching off or hedging the positions.

Yet, the same trading desks that couldn't make money trading single name credit default swaps (CDS), have experienced recent explosive growth. There are several reasons, but in my view the main ones are as follows:

- Trading desks are unlikely to experience short-term losses despite massive net long positions in credit risk;
- Senior investment bank and bank managers, confused by misleading jargon, are unaware they have large one-sided positions in credit risk; and
- The recent changes in the way structured credit products are presented and quoted obscure the cash flows for both senior bank managers and investors.

HYPOTHETICAL CASH CDO

A simple cash collateralized debt obligation is based on a portfolio of corporate bonds. The bonds throw off coupon income and are redeemed at par at maturity. For simplification, we'll use a hypothetical deal in which all of the bonds mature in 5 years (In practice, cash CDOs have a target average life and target final maturity due to the varied maturities of the underlying bonds). At the 5-year final maturity, the bonds are redeemed at par. Figure 1. shows the basic CDO structure.

Figure 1.

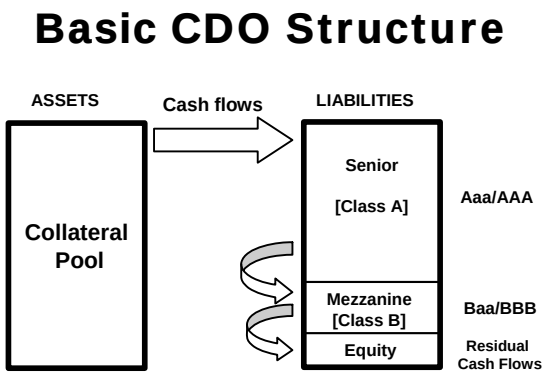

A CDO backed by the portfolio of bonds might be tranched into four classes of risk with the following ratings: a senior ("AAA") tranche, two mezzanine tranches (rated "A" and "BBB" respectively and shown in the figure as one "block" and one unrated first loss or equity tranche. Let's assume the "AAA" tranche trades at LIBOR + 50 in our hypothetical market. The investment bank arranging the deal will sell all of the tranches at market prices. The difference between the income from the portfolio and the cash owed to the investors, the liabilities, less the deal expenses (legal, rating agencies, structuring fees, and more) is known as the CDO "arbitrage". In particular, the investment bank arranger will normally pre-sell the equity tranche, the riskiest tranche. The implied internal rate of return at which this equity risk can be sold to an outside investor is a key determinant of the CDO arbitrage.

CASH FLOW MAGIC TRICK

When banks wanted to reduce the credit risk of their loan portfolios, they realized they had a problem. Loans were put on the books for very little income. For investment grade loans, the CDO "arbitrage" didn't work. Bankers invented a creative and somewhat deceptive solution. Bank regulators, including the Bank for International Settlements, allowed banks to get away with it. The solution was the Super-senior tranche.

If the investment grade deal were a cash deal, the "AAA" tranche would make up more than 85% of the CDO. That meant banks would have to pay out LIBOR + 50 bps on more than 85% of the deal (85% for the purposes of this paper). Since synthetic deals are unfunded, a "AAA" tranche investor would receive 50 bps on the notional amount of 85% of the deal. This was unworkable, until the creation of the super-senior tranche.

The synthetic deal above with a 5% "AAA" tranche under the super-senior tranche, would now have a super-senior tranche as 80% of the notional amount of the deal. Banks claimed this tranche had virtually no risk, and therefore, one shouldn't have to pay much of a premium to hedge the risk. Instead of paying 50bps, banks could now pay around 6 bps to an outside protection provider, or the bank could simply retain the risk.

Monolines have become the protection provider of choice. Standard and Poor's assesses one such financial guarantor a capital charge of 10 bps against "AAA" tranches. Even though the monoline's internal economic capital assessments are considerably higher than the rating agency's charge, super-senior deals clear return on capital hurdles. super-senior tranches on corporate underlyings have traded as low as 6 bps, especially on early CDO deals. This was a savings of 44 bps per annum. For a $5 billion five-year transaction, that amounted to $17.6 million per annum for five years.

Ironically, credit derivatives professionals have long criticized bank loan officers for putting on loans at ridiculously low levels. One credit derivative professional criticized bank loan practices: "We're in the business of taking on risk for the wrong fee for future business we're never going to get." The necessary improvement in the CLO arbitrage is what prompted creation of the super-senior tranche in the first place. Banks squeezed cash out of the "AAA" tranche to make up for the low income loans' cash shortfall so that the arbitrage for collateralized loan obligations would artificially appear more attractive. This was the only way to get the bulk of the credit risk off of the bank balance sheet.

The greatest triumph of illusion in 20th century finance is perfectly legal. Now that it suits banks and investment banks to put super-senior risk in trading books at much lower levels than the payment they formerly received for the same risk from the top tier of the "AAA" tranche, bank managers have been persuaded it is a great deal. Banks used the super-senior to facilitate moving risk off the bank balance sheet, but banks later used the super-senior to facilitate putting risk in trading books.

The negative basis trade is a new innovation designed to mitigate the risk of these large low risk positions. These are not so much trades as the build-up of a huge bank-financed high quality credit investment portfolio. In a "negative basis trade" at the super-senior level, the bank or investment bank trading book retains the very senior level in the CDO capital structure and buys protection from a monoline at a low cost in the form of a wrap. "Negative basis trade" is just a fancy name for "I want a big bonus for an old-fashioned low-tech carry trade." The trading book becomes a parking lot for a huge carry trade investment.

The excess cash hived off the super-senior tranche covers up a multitude of problems in the trading books. The synthetic income created from the low pricing of the super-senior tranche creates a huge slush fund of apparent new revenue for the trading book, which is partially reflected in the equity, first loss, cash flows; but is often reflected in larger fees retained by deal arrangers.

THE "AAA" DISAPPEARING ACT

The rating agencies do not quite know what to make of it. In the simple example above, I made up the numbers for the size of the "AAA" tranche and the super-senior tranche. Everyone else in the market makes up the numbers, too. There is no standard for the size of the "AAA" tranche required below the super-senior tranche. Even more disturbing, Standard and Poors (S&P) insists that they do not recognize the super-senior tranche. S&P does not recognize anything above "AAA." According to S&Ps definition, if the "AAA" is subordinated to anything else, it is no longer "AAA." Figure 2 shows an example of the "AAA" tranche as a horizontal slice of a tranched portfolio and as the

vertical slice of a synthetic deal using the same portfolio of underlying credits. The new "AAA" tranche of the synthetic deal is the first loss tranche of a formerly "AAA" tranche, and if defaults occur in the portfolio, it is more vulnerable to downgrade than the vertically sliced "AAA" tranche.

Figure 2.

Two Different "AAA" Tranches of the Same Portfolio

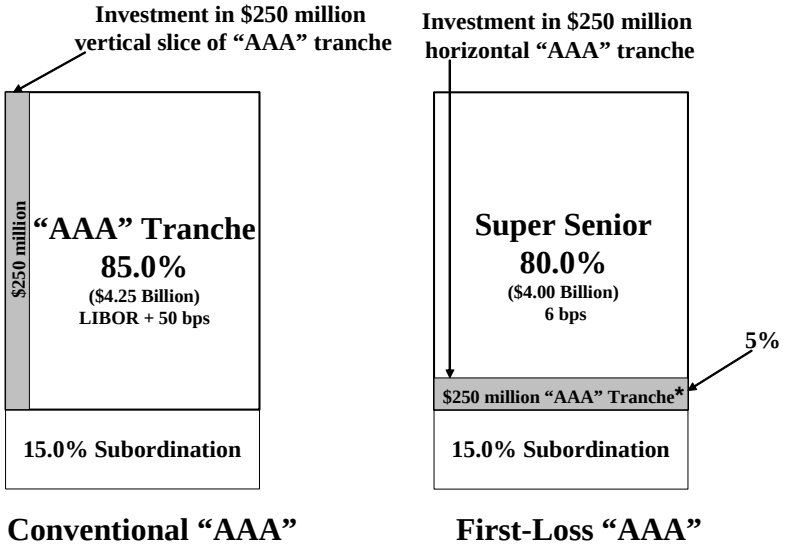

***Note: There is debate whether a "AAA" tranche can ever be subordinate.**

Investors should be aware that if they are sold the "AAA" tranche of a synthetic CDO, they should ask for more coupon income. They do not own a "AAA" tranche. What they own is the first loss position of what was formerly the "AAA" tranche. Yet, these first loss "AAA" tranches are being sold as legitimate "AAA" rated tranches. Investors should demand higher compensation for these "AAA" rated tranches.

To add to the confusion, an identical CDO tranche may earn the "AAA" rating from S&P, but only an "AA" rating from Moody's. If an investor will accept the S&P rating, then it should ask for even more coupon income to compensate for the subsequent liquidity risk. This is the perfect investment for a fund seeking a tranche nominally rated "AAA," but that provides the coupon income as if it were rated "AA." Less experienced investors will be at a disadvantage if they don't negotiate for more income.

The above is also true of cash asset backed CDO products, but the ratings and risk waters are particularly muddied for synthetic CDOs. Investors aren't alone in their potential confusion. Rating agencies and bank managers are also challenged by synthetic CDO products.

SYNTHETIC INCOME AND SYNTHETIC CDOs

Bank managers have been persuaded that the super-senior tranche is virtually risk free. With respect to actual losses at this level, that is a pretty good assumption. The assumption is less useful when one considers mark-to-market risk at the trading book level. The defaults of Worldcom, Kmart, Enron, Adelphia, and others left arrangers' books with portions of former super-senior tranches that would only merit an "AA" rating.

An honest mark-to-market would require the arrangers to show a loss. But fortunately for traders, senior bank managers did not control the mark-to-market process. The traders who put the super-senior on the books worked out the mark-to-market methodology with the risk managers, or colluded to get market "evidence" of the trading level. Since super-senior tranches rarely traded, there was no way to independently verify prices. Creative trading book managers traded small slices of super-senior tranches at pre-agreed prices as evidence that the levels at which they marked their own positions was correct.

In an exception to FASB's policy on guarantees, monoline wraps on synthetic CDO tranches – or any products that must be marked to market – require mark-to-market treatment. Some feel that monoline pricing is one method for marking-to-market super-senior risk.

At the first loss level, synthetic CDOs have created other valuation challenges. When credit spreads were wide, the first loss or equity tranche of an investment grade CDO generated cash flow which provided more than a 60% internal rate of return ("IRR") - assuming Moody's base case losses. This was such a good deal that arrangers often kept a portion of the excess spread for themselves and sold the equity at the lowest feasible IRR to unwary investors. This still happens today to investors in the equity tranches of index CDOs and customized first loss CDOs.

It would have been ideal to retain the entire first loss risk on the arrangers' books, and some arrangers did this, instead. At those levels, the risk was a good bet.

Or so it seemed.

A large chunk of the excess income was simply the "found money" from the super-senior tranche. Arrangers were keeping the super-senior tranche and taking the income from this former "AAA" tranche (remember the 44 bps) and applying it to their equity "investment." The arrangers robbed themselves of income at the "AAA" level to pay themselves at the equity tranche level. At 60% IRR, the equity still looked great, but a large chunk of that IRR was due to synthetic income.

It is fine to agree to report income that way, but it is not fine to deceive ourselves about fundamental value. What had originally been a device to remove credit risk was now a device to put credit risk on the trading book and claim high real revenues. In reality, much of this cash flow stream should have either been considered credit risk reserves or synthetic income and not applied as trading desk real revenues. But that would have been no fun at bonus time. Since senior bank managers didn't question the practice, arrangers and other participants now declare revenue victory when the reality is the cash flows have simply been reshuffled.

SYNTHETIC EQUITY

Is Cash Kevorkian your investment banker? C.K. is always ready to assist you with the suicide of your fair share of a deal's cash flows. As they say in poker, if you don't know how to spot the sucker at the table, it is you. If you bought a bespoke tranche, also known as a single tranche collateralized debt obligation (STCDO), without first running an independent cash flow analysis, the probability is very high that you are the sucker at the table. This is particularly true if you bought the equity tranche.

Notice that the average credit default swap level in the previous example was around 80 bps. Yet, today the same deals have been getting done when credit spreads have collapsed to an average of 50 bps for the identical portfolio. How can this possibly be a good deal for the arranging banks?

Bank arrangers have been doing relatively unprofitable deals under very tight credit spread conditions. That isn't immediately obvious when one sells a single tranche and then hedges it with insufficient income. Credit spreads have often been very tight in the past couple of years. The return on the implied equity has been under 10% versus 60% and higher when credit spreads have been very wide. For single tranche hedges, this means that the income on the credit default swaps used to hedge the position is too low for the risk. If one considers the full notional amount of the CDO, it is easy to see arrangers took on too much risk for too little reward. Senior bank managers and investors have been fooled by the lack of transparency in these deals into thinking that they are doing quality business.

The CDX4, an investment grade CDS index, can be used as the underlying portfolio of an STCDO tailored to the desired tranche requirements of an investor. In a rational world, the investor would use a rating agency model to tranche a CDO using the underlying reference credits in the index. Then the investor would run the cash flows of the deal to make sure it is getting fair value in the deal. This is particularly important if the investor wants to buy the equity tranche. But this rarely happens.

Initially, the indexes were touted by everyone – including me – as a giant step forward in transparency for the CDS market in general and for synthetic CDOs in particular. We have managed to turn transparency to mud, because transparency meant that arrangers could not retain some of an unwary investor's cash flow value for themselves.

Instead of showing investors the potential internal rate of returns – along with clearly articulated assumptions and scenario cash flows - under different scenarios, arrangers quote correlation values and sell equity risk by offering an upfront cash payment and ongoing fixed payments, provided no defaults occur.

Unwary investors think this is a great deal. They are investing in a portfolio of investment grade names that they believe are unlikely to default. They receive upfront cash and ongoing payments. For the CDX4, as of the end of August 2005, for every $10 million pledged to cover losses of the equity tranche, the horizontal 3% notional slice, an investor would have received an upfront payment of $3.9 million plus 500 bps per annum. Many investors looked as this as free money, but the cash payments were too low relative to the risk and relative to what they should have asked for.

In early November 2005, investors could get paid an upfront premium of 80%+ (note this is not IRR) of the notional amount of first loss risk on the bottom 0-3% tranche of the investment grade index. It may appear like a good deal until one runs the implied cash flows, but this tranche actually provides the worst value of all of the tranches in the deal.

The mezzanine investor should also be wary. It is easy to fiddle around with the tranching of the mezzanine tranche. The way that index CDOs are quoted, obscures the implied rating of the various slices offered to investors. Investors should always independently run a rating agency model and determine the amount of subordination required for the rating quality one wants. The investor should then do a price comparison. I recommend Moody's, since it usually requires more subordination or protection for the investor. S&P uses a "black box" called The Evaluator. The assumptions used in The Evaluator are not transparent, so it is difficult to trust the ratings produced by this model.

If you are willing to take the equity risk of a deal, you should arrange your own deal and earn the excess spread on the entire notional amount of the deal. You should always choose your own portfolio, and this won't necessarily include all of the names in a pre-determined index. Eliminating restructuring as a credit event should be standard, however. The cash flow return for the risk will be much greater than that offered to you by any other arranger.

One very savvy hedge fund retains all of the equity and shorts the mezzanine tranche without using a bank arranger. The hedge fund realizes it can play with the mezzanine subordination levels. It also realizes that retaining all of the excess spread for its own benefit is a much better deal than the fixed payments offered for equity risk by bank arrangers.

As a general rule, never accept fixed payments from an arranger if you are the buyer of equity risk.

There are no hard and fast rules for how much implied IRR is enough, but here are some suggestions based on the past 3 years of credit spread fluctuations. If one isn't getting more than a 25% base case IRR, it is a good idea to wait for a wider credit spread environment. If you see a base cases higher than that, examine the IRR under different scenarios. You may wish to assign probabilities to the various scenarios to calculate a probable return, and also examine the IRRs under the various scenarios to make sure it is always acceptable under every scenario that concerns you. In recent history, base case IRR's of more than 50% have been feasible, and held up well under stress scenarios. Notice that a wider credit spread does not necessarily mean more intrinsic credit risk.

SINGLE TRANCHE CDOs

The next evolutionary change was the single tranche CDO. These trading books grew in size, because arrangers didn't invest in the resources to effectively market all of the CDO tranches. Instead, arrangers typically sold only the mezzanine risk. As a result, arrangers often unwittingly held equity risk – as well as the risk of the higher rated tranches – in their trading books.

If one looks at the analogy with a cash CDO, it is easy to see the flaw in the hedge practices. If this were a cash deal, the arranger would sell the mezzanine tranche of the CDO and "delta hedge" by buying each bond named in the portfolio, but only a fractional

amount of the total original reference portfolio. If the deal were $5 billion in size (difficult to achieve in the bond market), the mezzanine tranche might be only $250 million in size, and the corresponding hedge would be to buy around $1 billion in bonds, assuming the market risk were already perfectly hedged, typical of recent delta hedge ratios for investment grade CDOs calculated using correlation models.

Investment bank arrangers often do not know how to quantify the value of synthetic single tranche CDO trading books. Arrangers wish to report the size of the synthetic CDO market by only counting the sold tranches. But the risk of the arrangers position is equivalent **to** that of the remaining tranches of the full notional amount of the synthetic CDO. For instance, if the arranger sells a $250 million mezzanine tranche of a $5 billion CDO, the remaining risk position has the impact of the remaining $4.75 billion synthetic CDO in terms of value at risk.

When deal arrangers sold each tranche of the synthetic CDO and then sold protection in the credit derivatives market on each portfolio position, they were fully hedged. Since single tranche trading books usually sell only the mezzanine (the middle) tranche and then delta hedge, they are riskier than fully hedged positions. Theoretically, one should make more money than the fully hedged position, because one has more risk. That hasn't been the market experience very often, however. The single tranche books often make even less money for greater risk, because traders make additional credit bets or make a bet on hedge ratios.

When traders lose money on the bets relative to a fully hedged position, the trading book is making insufficient income relative to the risk. That should show up in lower bonuses. It should, but it often does not, because senior management has no idea what is happening in the trading book relative to customary CDO business.

Risk is measured as value at risk (in simplified terms). The sold tranches of synthetic CDOs are notional slices of a much larger CDO, and the equivalent risk of the unsold CDO tranches remains on the arrangers books. The sold tranche only bears a horizontal slice of the risk of the entire CDO. The typical hedge for the sold tranche is to sell credit default swaps that are a vertical slice of the risk of the notional amount sold. The hedge ratio is a best guess.

It is very easy for traders to obscure the risk/reward economics of these deals. These are the single most gamed transactions both from the investor's and ironically, from the investment bank/arranger's point of view.

The cleanest and most conservative (from a risk perspective) hedge for a synthetic CDO is to sell protection in the credit default swap market on the full notional amount of the deal while simultaneously buying protection from the CDO investors by selling every tranche of the deal. In contrast, the STCDO always leaves the investment bank/arranger's trading book with residual credit risk and a huge position to hedge.

One would think that greater risk means greater reward. One might reason that fully hedging the deal cuts into the deal profitability. Often just the opposite is true, given the state of the art of hedging STCDO deals.

If this is true, why do banks do STCDOs? One of the main reasons is that arrangers have difficulty finding investors for all of the tranches. Arrangers created STCDOs so they could still do business without developing a deeper investor base for these transactions.

The following is an example of a STCDO brought to market when credit spreads were wider than today's levels. If this deal using the same reference portfolio and notional amount had been brought to market as a fully hedged CDO, the locked in revenue of the deal would have been approximately $11.3 million. For example, if the deal arranger locks in the bid/offer spread on 200 names, and if the bid/offer spread is 5 bps (For many off-the-run names, the bid/off spread is greater than 5 bps and some well traded names may have a tighter bid ask spread), the arranger can lock in $2.5 million per annum, for a typical deal with a 5-year maturity. In the no-default scenario, this has a present value of about $11.3 million using a 4% discount rate over the life of the deal.

For our mezzanine STCDO, the arranger sold the mezzanine tranche, but now the arranger was long the credit risk of the credit default protection sold to delta hedge the short mezzanine position. The arranger is long the credit risk of a large slice of the equity tranche (often unbeknownst to its own management), long the credit risk of the tranches that are senior to the particular mezzanine class being shorted as illustrated in Figure 3.

Figure 3.

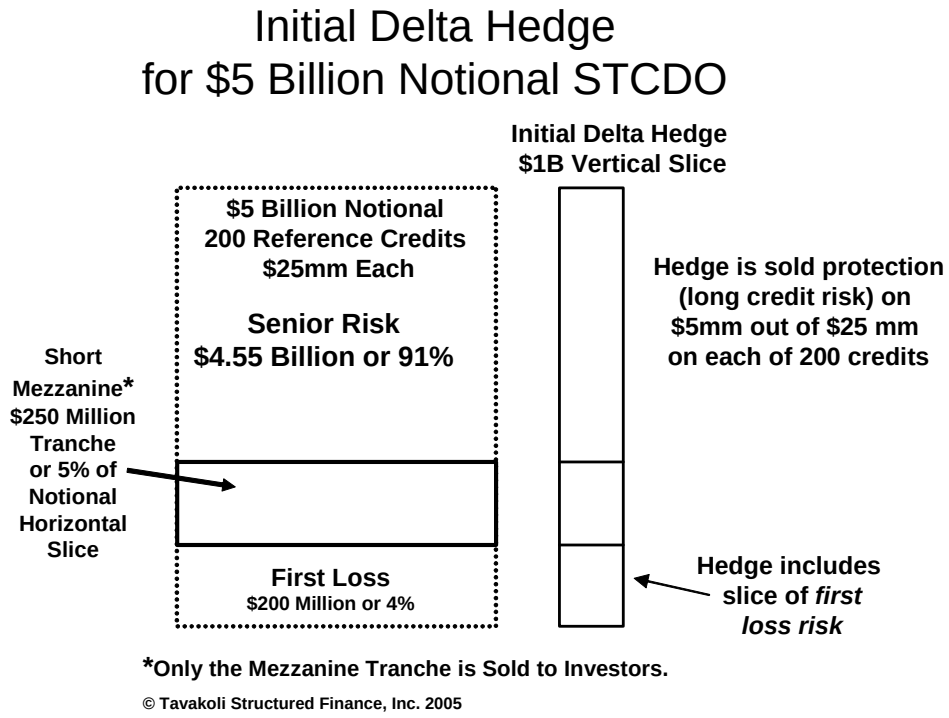

For a $5 billion notional CDO, the arranger created a mezzanine risk STCDO for an investor with a notional size of $250 million, and an approximate credit risk equivalent to a tranche rated "AA." The STCDO arranger needed to pay the stated coupon on the mezzanine tranche. Instead of selling protection on the entire $5 billon notional amount, the arranger initially only sold protection on a notional amount of $1 billion. Note that the hedge amount must be recalculated over time, but this is the initial delta hedge amount.

The average offer spread on the portfolio was 80 bps. The stated coupon on the $250 million mezzanine tranche is 120 bps per annum. The arranger earned the following:

$1,000,000,000 x .0080 = $8,000,000 per annum.

After tranche expenses, the net amount was $ 1 billion x 80 bps - $250 million x 120 bps = $5 million per annum or $22.3 million over the 5-year life of the deal assuming a 4% discount rate.

Note that if there are no defaults, the arranger does better than it otherwise would have done with the fully hedged deal. At first glance, one might think the arranger is making more money doing the STCDO than doing a fully hedged deal, but this isn't necessary the case. This revenue must cover other deal expenses plus reserves for potential losses. In other words, the higher apparent revenue comes with greater risk.

For instance, if one of the 200 names, or reference entities, in the reference portfolio defaulted immediately, the arranger must settle on the protection it sold. Let's assume a potentially generous recovery rate of 40% in the event of default. The original notional amount was $5 billion, and the arranger sold protection on approximately 1/5 of that amount, or $1 billion. In the event of default of one of the $25 million notional reference entities, the arranger would have to pay a cash amount of (1/5 x $25 million) x (1 – 0.4) = $3,000,000.

When the STCDO comes to market, it isn't obvious that the arranger has the first loss risk. If losses are tracked by the individual deal, however, it is easy to see that the arranger has underwritten the equity of the entire CDO, albeit losses will be limited to the notional amounts of the individual names on which the arranger has sold credit default protection. Risk adjusted income is a key metric, but there is a lot of confusion in how to quantify the risk.

In this case, the arranger has the first loss risk on the $1 billion in credit default protection it sold, but it also has the income on only the $1 billion it sold, the risk/reward is theoretically equivalent to the entire equity tranche. Note that $1 billion is only the initial hedge amount. Over time the hedge ratio will have to be readjusted.

But you ask – what about the short mezzanine position? Doesn't correlation trading ensure that the short position will change in value similar to the long credit risk position? No. Not by a long shot. Correlation trading is a guess. The "delta" of the price of the mezzanine tranche with respect to credit spread changes is largely an unknown, despite what arrangers claim. Often credit spreads on underlying names can gap out, but the mezzanine pricing for rated tranches can remain constant. This has often happened in the market. The models that try to quantify this relationship are fairly useless.

One of the key problems is that the short mezzanine position has different risk characteristics than the long position of sold protection. If credits are downgraded or if defaults occur which do not threaten the principal of the mezzanine tranche, it is difficult to monetize in the increase in value of this short position.

In May of 2005, S&P's downgrades of GM and Ford to below investment grade highlighted the weaknesses of the models. Traders blamed losses on unanticipated changes in correlation, but this excuse just obscures the fact that the trading books are making one-sided credit bets.

When treasury books take losses, traders own up that their hedges reflected a view on the direction of interest rate moves, and they lost their bet. In the credit markets, we want to blame correlation or correlation smiles as correlation laughs at our efforts to model it. Later in this article, I will comment on the false analogy of the correlation smile, sometimes called correlation "convexity." rather than own up to the fact that these trading books are covertly accumulating income by piling on credit risk – often at the wrong prices - the same behavior for which we formerly criticized loan officers.

The advantage to all of this smoke and mirrors is that one can fiddle around with the model assumptions. One can change the hedge ratio and create a lot of synthetic income at will. Senior managers and risk managers will rarely successfully challenge this slight of hand. We'll explore this in more detail later.

INVISIBLE HEDGE FUNDS

In our example, whether the arranger makes more money over time than they would have done if they had fully hedged their position is a matter of how well they manage their residual risk. Whether or not it was a good deal in the first place, depends on the implied pricing of the equity risk.

Notice that the initial retained equity risk on the $1 billion hedge position is virtually invisible. For banks this is problematic. If reported, this equity exposure should attract a dollar for dollar deduction against regulatory capital. This equity risk is not always reported, however. As the head of structuring at a Canadian bank said: "If my boss knew what I was really doing, he'd fire me."

To add to the confusion, traders tend to make additional lopsided bets to exploit a particular viewpoint on the way they hedge the risk. The traders don't have a lot of downside in making a bet, but have quite a bit of upside.

If they aren't rigorously tracked, STCDO cash flows effectively obscure the effect of bets. If credit bets win, the traders have a call on the upside in the form of a higher bonus related to the higher profitability (the bet payoff). Anyone who owns such a call option has an affinity for greater volatility, so the interests of the arranger's managers and traders are misaligned. This isn't intentional, it is the nature of this business, and managers must fully understand this natural tension.

If the arranger isn't making more on STCDOs than the arranger would have made on a fully hedged CDO, then the positions weren't managed in a way that compensated the arranger for the additional risk. That means their traders were unsuccessful in exploiting the risk/reward potential (or they made unsuccessful bets), and the arranger would have been better off fully hedging their positions.

EXTRAORDINARY POPULAR DELUSIONS AND THE MADNESS OF CORRELATION

The 21st century has started off with a spectacular mass financial delusion in the fastest growing segment of the financial markets. One generally considers three stochastic variables when calculating credit losses: default probability, recovery rates, and correlation. Of these three, correlation is the *least* important. Yet correlation trading has spread through the psyche of the financial markets like a highly infectious thought virus. So far, there have been few fatalities, but several victims have fallen ill, and the disease is

rapidly spreading. Financial engineers who have built a career around market default correlation assumptions are likely to add cardiac arrest to their list of their symptoms.

To illustrate the madness, consider the following problem. Suppose I told you that I ran a single name CDS trading book and had sold credit default protection on $1 billion notional of 200 diversified investment-grade corporate names. How would you suggest that I hedge the risk? One might suggest various approaches: Leave the position unhedged (like an old-fashioned bank loan portfolio); Short the riskiest credits as indicated by a Moody's KMV model or CreditSights (or buy credit default protection for the riskiest names); Short an investment grade credit index in a calculated ratio. These are only a few examples among several possibilities.

The possibility that would probably come to your mind dead last is to create a $5 billion notional CDO using the 200 names as the underlying collateral, and then short the mezzanine tranche to hedge the risk. As time goes on, change the notional amount of the $1 billion sold credit default swap position to adjust for new market conditions. The reason this would be the last to leap to mind is because the hedge is nonsense. Yet, when we create and sell the mezzanine tranche of a $5 billion notional deal, the hedge of choice is a delta hedge which requires us to sell credit default protection on a pro-rata amount of each of the underlying 200 names amounting to $1 billion notional of sold credit default protection.

Correlation traders have adopted the language of the conventional options markets. This is a false analogy. For conventional option pricing the volatility of the underlying is a function of the price of the underlying, the strike price of the option, the time to option expiration, interest rates, and the price of the option. If we consider an interest rate option, the price of the underlying instruments can be calculated with certainty for various interest rate environments.

Market implied default correlation is much more nebulous. Market implied default correlation is a function of the tranche being considered (the hard-to-pinpoint-price and the attachment points), the loss given default assumptions (as usually estimated from flawed averages of CDS premiums and historical loss given default), and asset default correlation assumptions (equal pairwise default correlations and a form of copula model – Gaussian, Student's t, or Archimedean).

All of this effort is a waste of time and resources. In the first place, the market is estimating *asset* correlations instead of the necessary *default* correlations. The overwhelming flaw in the methodology, however, is that the only reason default correlation exists at all is if we pretend that default probability does not vary, but of course, it does.

This suggests that hedging should focus on default probability and recovery rates. Instead of hiring quantitative analysts to write inaccurate models that solve the wrong problem, a fundamental corporate balance sheet analysis that reveals better information about the probability of default and loss given default is a more sensible approach. That requires a different skill set and a different focus than is currently in use on the majority of trading desks today.

A major improvement in the financial markets would be to spend zero time, money, and resources on correlation and spend all of one's time, money, and resources on better estimates of default probabilities and recovery rates.

DELTA FACTS AND FICTION

I recommended taking down the equity risk on the entire notional amount and skipping the part about delta hedging, but that is just a recommendation. If you want to delta hedge, go right ahead. Knock yourself out. Just don't call it science.

Since the data used by the models is so debatable, it is no surprise that the results are often ridiculous. The delta hedge ratio is defined as the change in the mark-to-market of a tranche for the change in mark-to-market of credit.

Arrangers typically calculate the delta by shifting the spreads on each individual name in the CDO by one to ten bps and then calculating the mark-to-market change of the tranche. It is a nice theory, but as mentioned previously, it doesn't work that way in the market.

Arrangers use the term "spread convexity" to explain departures from the model for large spread moves. This requires an adjustment to the hedge ratio.

The term "spread convexity" implies the predictability is as good as in the treasury market in which we define duration and convexity. The terms are similar, but the dynamics of these two markets are very different. In the treasury market one can precisely recalculate the price of a treasury bond for a change in yields. The duration and convexity of the price curve for different bonds vary in precisely predictable ways. Treasury models are reliable and the results are quantifiable in advance. In contrast, delta hedging models for STCDOs are not reliable, and a back of the envelope guess of the hedge ratio is as good as a model.

In the case of coins and dice, you can learn a lot about them by flipping and throwing thousands of times and recording the results. A Monte Carlo simulation uses a computer to throw a whole lot of random inputs into a model. It is like shaking a newly made chair to see how stable it is. Correlation models appear clever but have little to do with reality. The result is a chair that collapses beneath you as soon as you sit on it. The models are highly unstable. Small changes to model inputs result in huge hedge ratio changes.

This is very bad news for correlation models. In the credit markets we are guessing at the relationships, and we are guessing about most of the data. One huge guess is about the relative spread movements. We also guess somewhat about the subordination required to get the credit quality we desire. Then we guess about the correlations. Even the rating agencies do not agree on how to treat correlation. Finally, we guess on the recovery rate assumptions.

The experts at one investment bank agree that the delta of the entire CDO should be 100% of the deal's notional. Over time if there are no defaults, the delta of the equity tranche will increase to 100% and the delta of the tranches above it will decrease to 0%. But when they calculate the individual deltas of the tranches in a synthetic CDO using their model, however, the deltas do not add up to 100% unless they force the model to give them that result. By that I mean they have to fudge, i.e. fake, the results.

Figure 4. shows a summary of their results. Note that the deltas add up to 118% of the notional of the hypothetical $1 billion notional CDO.

Figure 4.

Synthetic CDO Delta Hedge Summary

Tranche	Delta as % of Notional	Notional Size of Delta Hedge	Tranche Notional	Leverage Delta / Tranche Notional	Tranche Size (% Total)
Super Senior	22.36%	223,600,000	880,000,000	0.3	88.00%
Senior	18.88%	188,840,040	40,000,000	4.7	4.00%
Mezzanine	30.62%	306,200,000	40,000,000	7.7	4.00%
Equity	46.37%	463,700,000	40,000,000	11.6	4.00%
Total	**118.23%**	**1,182,340,040**	1,000,000,000	NA	100.00%

What is the correct delta on the mezzanine tranche, the 4% - 8% horizontal slice of this custom tailored synthetic single tranche CDO? Who can say for certain? We could use 7.7x leverage, but would 7.3x be more accurate? If I fiddle with the model, I may be able to get a hedge ratio of 7.9x or even 8.1x. Notice that for the same trade, I can fiddle around quite a bit, and boost my income while claiming the identical risk profile.

Increasing the hedge ratio creates an income slush fund for the same reported risk. Risk managers often rely on the recommendations of the much more highly paid traders for bonuses and future jobs. This tends to make risk managers open to suggestion and very accommodating. Senior managers are often oblivious to the nuances of these "models." At bonus time, there is an incentive to create synthetic income to aid in the negotiation for real bonus dollars.

DEBATE OVER THE SIZE OF THE SYNTHETIC CDO MARKET

Normally I would start an article by describing the size of the market I intend to discuss. The synthetic CDO market is different, however. When the first STCDOs appeared, the market practice was to report the full notional amount of the CDO reference portfolio. The reported notional amounts soon ballooned beyond what had been reported for fully hedged deals.

The huge notional amounts created a potential problem. Senior bank managers and bank regulators could become uncomfortable with the huge reported notionals. Even more important, they might become uncomfortable with the huge increase in risk taken on by traders who formerly claimed they couldn't make much money trading single tranche CDOs.

There was only one thing to do. Instead of reporting the entire deal notional amounts, arrangers only reported the amount of the tranches they actually sold.

This new method of reporting provides little meaningful content. Notice that if one adds up the notional amounts, it doesn't reveal anything about the degree of tranching, the relative risk taken on given the underlying portfolio, or the potential risk as credit conditions change. The amount could refer to the equity tranche, the "BBB" tranche, the "AA" tranche, or a blend of the CDO tranches. These are very different risks and reporting the notional amounts this way isn't very useful.

The only way to back out the risk of these deals is to start with the full notional amount. The only reliable benchmark is the fully hedged risk/reward of a hypothetical fully hedged deal. Furthermore, subsequent tranches of the entire deal can be issued, leaving no tranches on the trading book. I believe the only sensible way to report the statistics of synthetic CDOs is to report the entire notional amounts of the deals based on the entire notional amounts of the reference portfolios.

If the single tranche notional amounts aggregated from dealers and reported by *CreditFlux* are correct, then the underlying notional amounts of the CDOs are very large. As of the end of July 2005, bespoke tranches were reported to be $169 billion. If this represents tranches ranging on approximately average from 5% - 10% of the total notional amount, then the full notional of the deals ranged from approximately $1.7 trillion to $3.4 trillion.

I've heard disagreement and dissatisfaction from arrangers about my position on this issue. I've heard arguments involving implied correlations, dynamic hedging, spread convexity, and other obfuscating jargon. Can you spend implied correlation? Can you spend hedge ratios? Can you spend spread convexity? No. But you can spend cash.

Insist on a cash flow analysis that shows the timing, the magnitude and the probability of receipt of the cash flows. The only way to do this is to start with a first principles analysis of the deal, which requires knowledge of the full notional amount of the CDO. I haven't heard any sensible alternatives to reporting full notional amounts, but I would be delighted if someone could provide an alternative that didn't involve magical thinking.

SYNTHETIC INCOME AND MORAL HAZARD

If I had a large bonus in my sights and mischief on my mind, it would be very easy to manipulate the hedge ratio so that I would sell more credit default protection and increase my income without a perceived increase in risk. This is source of the extreme moral hazard of the delta hedging strategy.

It would be even simpler if one could do away with the inconvenient sold single tranche and simply sell credit default protection outright. But most investment managers would not allow this type of one-sided bet. The bet would simply be a use of the bank's credit line to generate risky income. Yet by manipulating my hedge ratios, I can create a large net long position in credit risk that is virtually invisible.

The easiest way for an arranger to do this is to sell the mezzanine tranche of a synthetic CDO and then delta hedge. Given the arguments about correlation data, and given the dissimilar nature of the hedge to the risk being hedged, the outcome is a coin

flip. In this case, the coin is biased, because the hedger has control of the hedge ratio and therefore control of the income and of the perceived risk. The bet is that credit spreads will not gap out and that defaults will not touch the trading book until after bonus time, or even better, until one has a big contract at another investment bank.

I would like to thank Arturo Cifuentes, Ph.D., Managing Director, R.W. Pressprich & Co. and Robert Selvaggio, Ph.D., Managing Director, Ambac Assurance Corporation for their comments and suggestions.

Janet Tavakoli is the president of **Tavakoli Structured Finance**, a Chicago-based firm that provides consulting and expert-witness services. Ms. Tavakoli has more than 20 years of experience in senior investment banking positions, trading, structuring and marketing structured financial products. She is a former adjunct professor of derivatives at the University of Chicago's Graduate School of Business, and she has recently authored a pair of books on the credit markets: *Credit Derivatives & Synthetic Structures* (John Wiley & Sons, 2nd edition, 2001) and *Collateralized Debt Obligations & Structured Finance* (John Wiley & Sons, 2003).