Eating our own cooking

Christopher C. Finger
chris.finger@riskmetrics.com

June 2005

Last month, we wrote about backtesting. We concluded with the speculation that future work in backtesting will focus not just on questions of model accuracy (*Do large losses occur with the expected frequency and timing?*), but on diagnostics (*If the model has failed, what part of the model is at fault?*) and explanations (*What do I tell my regulator?*). As if on order, last month provided us with market events that stressed our risk models, and consequently an opportunity to act on our last note's suggestion. In this note, we take stock of recent events in tranched credit indices, and diagnose the performance of our risk models during this period. Our purpose is twofold: first, to diagnose this specific event, and second, to develop a set of questions we might use in future cases.

The products

There are three credit derivative products relevant to our discussion: single-name credit default swaps (CDS), credit default swap index contracts, and tranches on the indices. The single-name CDS is the most simple: a protection seller (or risk taker) receives a regular premium in return for providing credit protection on a single issuer. The protection buyer pays the premium and, in the event that the issuer defaults, receives a payment of par less recovery for a particular referenced bond. Single-name CDS are quoted according to the fair premium (or spread) for an issuer at standard maturities. A CDS paying the fair spread has an initial present value of zero. The fair spread for each issuer is the risk factor we model for risk analysis.

Index contracts are similar to single name CDS. A protection seller earns a premium in return for providing credit protection; the protection, however, is not on a single issuer, but rather on a standard, equally weighted basket of issuers. The index contract is almost the same as a simple portfolio of single-name CDS, with a few minor differences. First, the index trades as a contract paying a specified premium. This means that the present value of the contract is not always zero, and that entering the contract may require some upfront payment.[1] Second, because the index is a contract in its own right, it has its own supply and demand, meaning that it can trade differently from the equivalent CDS portfolio.[2]

Indices exist for a number of markets, though we will concern ourselves with just one: Dow Jones CDX North American Investment Grade, or simply CDX from here on. The index contract does have a fixed spread, but it is typically quoted as a fair spread: the

[1] This also means that as defaults occur, the premium rate stays constant; in a simple portfolio of CDS, the premium rate would be reduced, for instance, if a default occurred on a name whose CDS was paying a higher than average spread.

[2] A third, more subtle distinction is that the legal definition of default used in the index can differ from that used in single-name CDS.

spread that in theory would make the index contract have zero present value. For risk analysis, it is most straightforward to model just the index (fair) spread as the risk factor; it is also possible to model both the spreads of the portfolio constituents along with the basis between the theoretical and actual index spread.

Tranches on indices are an application of basic credit portfolio structuring techniques to the standard index portfolios. They are in most respects identical in structure to many synthetic Collateralized Debt Obligations (CDOs), but are special in that they reference standard index portfolios and consequently trade more liquidly than most CDOs. A tranche on an index references some slice of the possible index portfolio loss. For example, in the 3-7% tranche on the CDX, the protection seller receives a premium in return for covering CDX portfolio losses in excess of 3%, but not more than 7%. Among the standard index tranches, the first loss (that is, 0-3%) tranche has a fixed quarterly premium[3] and trades according to a varying upfront payment. The other standard tranches[4] trade, like single-name CDS, according to a fair spread, and require no upfront payment.

In general, the price of a tranche depends first on the quality of the credits in the index portfolio. As with the index contracts, we may consider the index spread (roughly, the average quality of the portfolio) as a single risk factor, or we may consider the spread of the individual issuers as distinct risk factors. Beyond spreads, the other factors that drive tranche pricing depend on our choice of pricing model. Pricing models for tranches have for the most part standardized.[5] Under the standard model, tranche prices depend on the average correlation between the portfolio constituents. In a simple world, the pricing model, spreads, and a single correlation value would determine the price of all tranches on an index. This is not the case, however, and so each tranche is associated with a notion of implied correlation: the correlation level such that the standard model prices a particular tranche.[6]

Interpretations of implied correlation vary. On one hand, implied correlation can be seen as a fundamental indicator: the expected dependence between defaults that the market expresses through tranche prices. Consequently, changes in implied correlations represent shifts in market expectations of the relationship between future defaults. On the other hand, those cynical to the model see implied correlation as little more than a technical artifact of an incomplete pricing model, and point to the need for multiple implied correlations as evidence of the model's shortcomings.

Regardless of our disposition toward the pricing model, we do have to accept that the models are widely used in practice and provide an important connection between changes in the underlying asset (the index portfolio) and changes in the derivative security (the tranche). We adopt the posture, then, that implied correlation is a model parameter with some fundamental basis, but that shifts with the market's appetite for tranches. Changes in implied correlation indicate simply changes in tranche prices that we cannot explain by spread moves alone.

[3] 500bp of the amount currently being protected

[4] 3-7%, 7-10%, 10-15%, and 15-30%

[5] See Finger (2004) and references therein.

[6] Even this point is subject to discussion, as implementations of the standard model will differ as to precisely which tranches to infer correlation from. See Finger (2004) for further discussion and references.

Table 1: Hedged index tranches on May 2, 2005

Position	Ref. port. size ($M)	Tranche size ($M)	Upfront fee ($M)	Premium ($000)	DV01 ($000)	99% VaR ($000)
Sell 0-3% protection	100.0	3.0	1.34	150.0	-19.75	219.0
Buy 3-7% protection	142.5	5.7	0.00	-141.4	18.93	186.03
Total			1.34	8.64	-0.80	1.27

The trade

Among the most common applications of the tranche pricing model is to calculate hedge ratios. As we discussed above, the identification of risk factors for a tranche position is not completely straightforward, meaning that it is not evident what a hedge ratio is designed to hedge against. We could use the model to hedge against moves in the spread of any single issuer or against a combination of issuer spread moves; most typically, we use the model first to hedge against moves in the overall index.[7] For a specific spread move, the hedge ratio gives the amount of the index contract we need to buy to offset the resulting change in our tranche. Typically, the hedge ratio for the 0-3% tranche is around 12: for a tranche position protecting $3m of losses (that is, a position referencing a $100m portfolio), the index hedge would be a $36m position.

A hedged 0-3% tranche position is a way to earn the high premium of the risky tranche, and at the same time be immune (at least according to the chosen pricing model) to small changes in the index. Motivated by giving up less of their premium to the hedge, traders have observed that it can be advantageous to hedge not with the index contract, but with another tranche. Thus, a common trade is to sell protection on the 0-3% and buy protection on the 3-7% tranche such that the changes in value of the two tranches, in the case of a small index move, are offsetting.

We analyze this trade as of May 2, 2005. On this day, the CDX index fair spread was 59bp. The 0-3% tranche traded with a 44.5% upfront fee, and 500bp running premium; the 3-7% tranche traded with a 248bp running premium. The standard pricing model suggested a hedge ratio of about 1.9 between the two tranches. We detail the trade in Table 1. On one side, we sell protection on the 0-3% tranche referencing a $100M index portfolio; thus, the our tranche position is $3M. This brings in $1.34M upfront and, assuming no defaults occur, also brings an annual premium payment of $150k. To offset the DV01 (that is, the price move due to a 1bp shift in the index), we buy protection on the 3-7% tranche. Our tranche position is $5.7M (the hedge ratio 1.9 times the $3M position in the first tranche). This hedge costs us nothing upfront, and, again assuming no defaults occur, costs us $141k annually. Overall, the trade nets us $1.34M

[7] We actually have to be more specific here, as there are many different ways that the underlying issuers can move to produce the same index change. For our examples, we will specify an index move in which all of the issuer spreads move by the same number of basis points.

Figure 1: CDX returns and Value-at-Risk

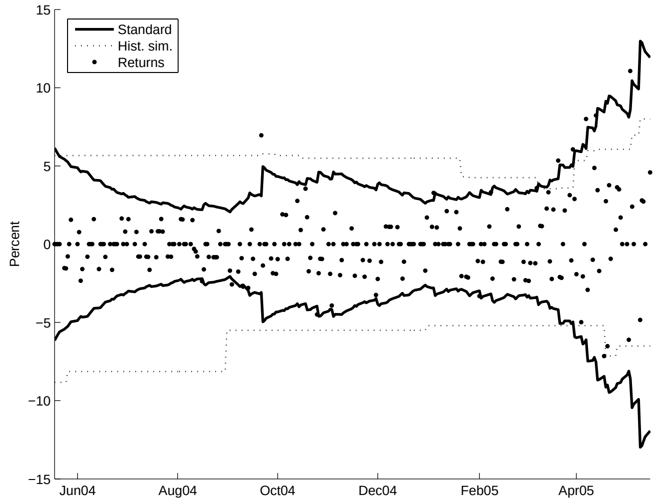

upfront, pays \$8640 (or 29bp on our 0-3% tranche position) annually, and in theory should have very little sensitivity to moves in the CDX index.[8]

Based on our standard[9] volatility estimate, we forecast the 99% worst case one week move in the index to be a widening to 71bp or tightening to 49bp. Our VaR forecasts for the two legs, as well as for the overall trade, are listed in Table 1. Note that under the assumption of constant implied correlation, there is a significant offset in risk, and the combined position has very low likelihood of losing money.

In reality, over the next week, we lost \$186k on our 0-3% position, \$149k on our 3-7% position, and \$335k combined. Note that no actual defaults occurred; the losses were purely due to changes in tranche pricing. Neither of the individual losses are greater than the corresponding VaR forecasts, but the loss on the hedged position is over 250 times greater than our risk estimate.

What happened to our model?

Clearly, our (supposedly) hedged position has given us a VaR exception to analyze. One exception is of course not enough reason to reject a VaR model, but the magnitude of the exception is cause for concern.

Following the framework from last month's note, the first question we need to answer is whether our forecasts are adequate for those risk factors that we do in fact model. Over the week in question, the CDX moved from 59bp to 64bp, well within our 99% forecast bands. Furthermore, as we see in Figure 1, the VaR estimates for the CDX have performed reasonably well over the last year. There have been enough exceptions (ten) to barely qualify for the regulatory red zone. However, the exceptions that have occurred have been only slightly greater than the VaR forecast

[8]Of course, should defaults occur, we incur losses: assuming recovery rates of 40%, we pay out loss protection for the first six defaults on the index portfolio, and receive loss protection for the subsequent eight defaults. (Recall that the index contains 125 issuers.)

[9]Exponentially weighted moving average with a decay factor of 0.94, applied to logarithmic returns

Figure 2: Profit-and-loss for 0-3% tranche position

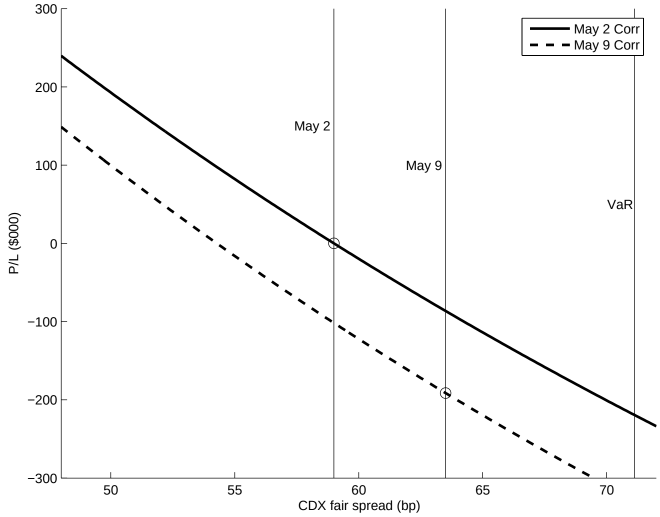

at the time, and the forecasts have reacted to periods of greater and less volatility, indicating adequate conditional coverage. For comparison, we plot the VaR bands from a historical simulation model as well; though these produce fewer exceptions, the exceptions that do occur are significantly greater than the somewhat stale forecasts at the time. In any case, the forecasts of the CDX risk factor itself, while perhaps needing minor adjustments, are not the cause of our large VaR exception.

Our second question is whether our pricing model, under our risk assumptions, adequately reflects changes in the positions. Consider Figure 2. The solid line represents the relationship between the CDX spread and the profit-and-loss on our 0-3% position, based on the assumption that implied correlation remain at its May 2 level; the dashed line represents the relationship assuming May 9 implied correlations. The vertical bars represent the actual CDX level on May 2 (59bp) and May 9 (64bp), as well as the forecasted worst case CDX level (71bp). The move down the solid curve between the May 2 and May 9 verticals is the loss on the 0-3% position that can be explained by our pricing model; the jump between the two curves is the loss due to moves in implied correlation, which are unexplained by our risk model. Thus, the explained loss makes up only about one half of the total loss. Coincidentally though, the total loss is still slightly less than our VaR forecast (mostly since the CDX moved by less than our risk forecast), and we do not observe a VaR exception on our 0-3% position.

For our 3-7% position, the story is similar: the unexplained loss is significant, though not enough to produce an overall loss greater than our VaR forecast. Critically, though, the unexplained loss is positive for both legs of our trade, with dramatic effects on the overall picture. Figure 3 depicts the same relationships for our overall hedged position. As we expect, under the assumption of constant implied correlation, there is very little possibility for loss. However, the jump to May 9 correlation levels produces a significant unexplained loss, accounting for almost all of the total.

In the end, that there was no VaR exception with either leg of the trade is actually misleading. By ex-

Figure 3: Profit-and-loss for hedged tranche position

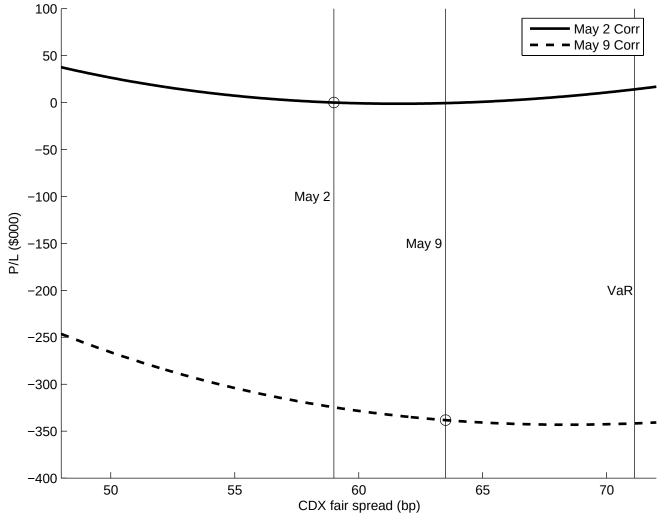

amining the explained and unexplained losses, we see that it was only by chance that the unexplained loss on either leg was not sufficient to overwhelm our VaR forecast. Most critically, the losses on the two legs were the wrong kind of losses. The trade was constructed based on the premise that losses on one leg would be offset by gains on the other. Losses explained by our pricing model have this property, but the unexplained losses do not. Thus, what ultimately caused the significant loss on the overall trade was that both legs experienced significant unexplained losses simultaneously.

Could we have done better?

The last question we must ask is whether we could have done better. To this end, there are three approaches to consider: a richer pricing model, a model of the evolution of the implied correlation parameters, and a direct (model-free) treatment of the tranche prices themselves.

Throughout our example, we have employed a pricing model in which the dependence of the tranche prices on the issuer spreads is only through the level of the index. This approximation makes the model much more straightforward and speeds up our risk calculations, but is an approximation nonetheless. For the 0-3% tranche in particular, the tranche price is more sensitive to the poorest quality names in the index, and a large move in one of these (indicating that the tranche is on the brink of realizing actual losses) has a greater effect than an average move spread out across all of the names. This is of particular interest for the week in question: though the index only rose by four basis points, Ford and General Motors widened by 135bp and 82bp, respectively. In the end, however, if we continue to leave the implied correlation levels constant, the granular model only explains 30% of the loss on the 0-3% tranche position, and only 10% of the loss on the overall trade. Clearly, more is missing than a granular treatment of spreads.

Our second approach is to model the implied corre-

[10] In this example, we model the 0-3% and 0-7% base correlations. We assume that logarithmic changes in the implied correla-

Figure 4: Relationship between changes in 0-3% and 3-7% CDX tranches

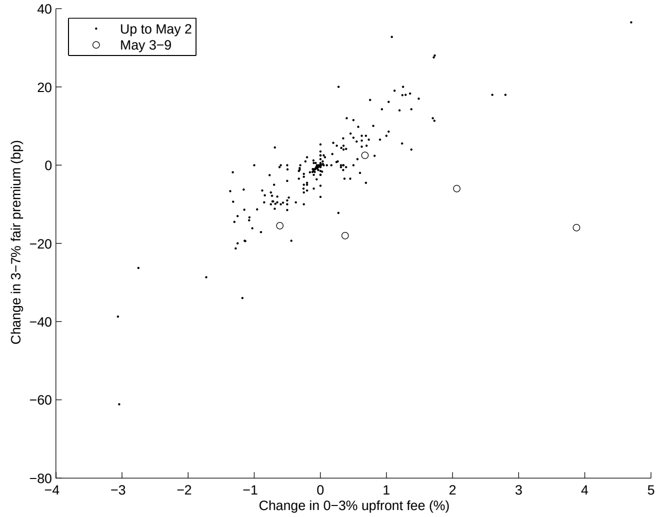

lation parameters as risk factors in their own right.[10] This is in priciple similar to the addition of a vega risk factor in a Value-at-Risk analysis of options. To a staunch proponent of the pricing models, this amounts to modeling an additional fundamental factor; to a cynic, to nothing more than modeling a technical factor caused by model error. Under either view, the empirical results are interesting.

A simple look at the histories of implied correlations looks promising: a comparison of returns to VaR forecasts yields a modest number of exceptions (eight) for the implied correlation. We estimate the covariance matrix for the joint returns on the CDX fair spread and implied correlations for the 0-3% and 3-7% tranches, and perform Monte Carlo simulations on our example trade. Our Value-at-Risk forecasts on the individual tranches change little from those in Table 1. For the overall trade, however, we see a significant change, with a Value-at-Risk moving to $97k from $1270. Though a large improvement, this forecast is still only one quarter of the actual loss.

As a final effort to improve our Value-at-Risk forecast, we consider a different approach, abandoning tranche pricing models altogether and estimating volatility and correlation of changes in the tranche prices themselves. Though appealing in its simplicity and possibly accurate for overall risk forecasts, we should warn that such approaches are ultimately less useful, as they do not permit any decomposition of risk (for example, into risk due to spread and correlation moves) nor stress tests using fundamental market factors. In any event, this "throw out the models" approach ultimately fails us as well. In Figure 4, we plot the daily price changes in the 0-3% and 3-7% tranches. For the year prior to May 2, the correlation between the two price changes was over 85%; three of the next five days, however, were clear outliers to this relationship. It is doubtful that any model based on the data up to May 2 could have forecast the dislocation that occurred over the next week.

tions are normally distributed. In theory, this assumption could be problematic, in that it permits correlations greater than 100%; in practice, in this exercise, under Monte Carlo simulations, this event never occurred.

So what did we learn?

Ultimately, though some improvements may have been possible, there appears to be little a historically based risk model could have done to forecast the magnitude of loss our trade experienced the week of May 2-9. Indeed, the conclusion of most of the dealer community is that the dislocation between the 0-3% and 3-7% tranches was due to unusual market technicals. Evidently, there were many market participants who had built significant positions in trades similar to ours. Concerned about the impending downgrade of Ford and General Motors, or seeking to close out positions before realizing greater losses, a marked number of investors tried to exit these trades, creating a simultaneous demand to buy 0-3% protection and sell 3-7% protection. The resulting demand effects on a still growing market caused the price moves we have already discussed.

It is tempting to conclude this discussion with the technicals explanation, and dismiss the market events as a surprise that no model could have predicted. As important, however, as the process of diagnosing the model is what we learn from this case where the model fundamentally missed our risk.

While this is a case where the statistics failed us, it is also one where expert knowledge and market insight could have helped. The knowledge that others had put on similar trades to ours, or the mere fact that our position was in significant size, should have prompted us to look at the trade more closely. And importantly, a closer look at the trade should have included not just an analysis of the model, but of the trader's motivation in constructing this particular set of positions in the first place.

The apparent safety of the trade relied on the assumption that the risk to spread moves was the only risk in the tranches. This may have been an adequate assumption for a single tranche position, but our trade was constructed to eliminate the risk in the only factor we were modeling. Recognizing that spreads are not the unique determinant of tranche pricing, we should pose the question of what else could go wrong. Though in previous notes, we have maligned stress testing as overused, this is a case where stress tests would have been a logical (and useful) way to at least avoid being surprised.

Further reading

- Finger, C. (2004). Issues in the pricing of synthetic CDOs, *RiskMetrics Journal*, **5**(1): 21–39.
- Finger, C. (2005). Back to backtesting, Research Monthly, RiskMetrics Group, May.
- Lucas, D. (2005). CDO Insight, UBS, May.
- Muench, D. et al (2005). Correlation talks, JP Morgan, May 10.
- Mahadevan, S. et al (2005). Credit Derivatives Insights, Morgan Stanley, May 13.
- Rosenberg, J. (2005). Situation Room: Correlation Contagion, Bank of America, May 10.