



A subprimer on risk

Christopher C. Finger
chris.finger@riskmetrics.com

August 2007

Early in 2007, there were concerns about two issues that could wind up causing significant havoc. One was the potential unwind of the yen carry trade, which we covered in our last issue. The other was the weakness in subprime mortgages in the US.

At the risk of further cursing the market, it is fair to say that the yen carry trade scenario has not come about as yet. But clearly, the subprime issue has come to a head, and by now has impacted markets directly linked to subprime, such as securitized products and the equity of mortgage lenders, and others where the link is at best indirect, such as corporate credit, leveraged buyouts, and global equities.

A postmortem analysis of these market events is premature, since the situation is still quite fluid. A comprehensive analysis of subprime mortgages, the catalyst of our current excitement, is beyond the scope of our efforts here. Still, it is not a market we can ignore, and so we offer some thoughts here on what is particular about subprime, and what we might be able to learn after this storm has blown over.

What is the subprime market?

At its heart, the subprime market is about mortgages granted to borrowers of less than stellar credit quality; but poor credit quality is not the only complication in the market. A related complication is that

lenders have created numerous innovations to make mortgages attainable for these borrowers. Subprime mortgages may now likely charge an initial low fixed rate and an adjusting rate thereafter, or provide the option to make monthly payments that do not cover current interest (meaning that the mortgage principal actually goes up), or allow borrowers to disclose full details of their employment status.

Topping all of this off, different mortgage originators have different definitions of subprime. Thus unlike traditional agency-issued mortgage-backed securities, where most pools contains a well defined “conforming” set of mortgages, subprime pools may vary significantly depending on which institution has originated them and when. So poor credit quality is only the beginning; subprime pools are also notably heterogeneous, both within and across pools, in terms of quality, mortgage type, and information available about the borrowers.

But the complexity does not end there. Pools of subprime mortgages from a given originator are packaged into asset-backed securities (ABS). A subprime ABS deal will hold a subprime pool as assets, and issue a number of securities (or tranches) as liabilities. Among many other features of a typical deal, cash-flows are assigned according to the priority of the tranches, meaning the most senior tranches have less risk (and the most junior tranches more risk) than

the original subprime pool. The tranches, then, can achieve a wide range of credit ratings, and appeal to different categories of investors.

Finally, sets of comparably rated ABS from different issuers can form the collateral pool for a collateralized debt obligation (CDO), with further prioritization of losses, as well as diversification effects, leading to more tailoring of credit ratings.

In summary, we have collateral of generally poor credit quality which is structured into a first ABS, which is then pooled with other ABS and structured again into a CDO. So as difficult as it is to ascertain the quality of a given pool of subprime mortgages, it is more difficult still to follow the impact of events on the pool through to the final investment vehicle. Not surprisingly, the opacity of the link between the fundamentals and the investment results in a market with almost no secondary trading, and therefore no price discovery (or data to build risk models) until something goes wrong and investors are forced to sell.

Given the argument about heterogeneity, we would at least expect that some ABSs or CDOs have better fundamentals than others, and that there should be some differentiation in pricing within a class of investments that were issued with the expectation of similar credit quality. But the lack of transparency appears currently to trump the heterogeneity:

...CDO prices are more pessimistic than our results, and there is a lack of distinction between awful bonds and reasonably healthy ones.¹

¹Lucas and Murray (2007)

²It is worth noting that a similar pattern emerged in 2001 in CDOs backed by high yield corporate bonds, particularly bonds issued by telecommunication firms in the late 1990s. Though the default rate on these bonds picked up early in the year, it took many large investors until the late summer to fully ascertain and disclose the extent of their losses on their CDO holdings.

There may well be investments that are being unfairly punished and are thus underpriced, but this is little solace to someone who must post margin based on this (albeit unfair) valuation.

A final implication of the complexity of the structures and consequent lack of pricing information is that the market is critically dependent on the rating agencies. But even if the agencies do their job well, it takes a long time to link new information about mortgages to new expectations on the likelihood of a CDO tranche paying its full interest and coupon. And the agencies will not ever—nor is it their mission to—warn about the price risk coming from an overall flight from the asset class.

Valuation troubles and illiquidity

The troubles with subprime mortgages are well documented: loose credit standards, small (or no) downpayments, flat (or falling) home prices. This means that structured investments in such mortgages are under stress, but as bad as the stress is the slow release of information. Mortgage delinquencies have been increasing since 2006, and yet only in the last few months have we seen the real strain on valuation for investors, or any real movement by the rating agencies to downgrade subprime-backed bonds.²

Difficulties in valuation go hand in hand with problems of illiquidity. For any investor, illiquidity presents the difficulty that new information is not reflected in prices in a timely manner, and when it is reflected, it is almost always through a significant

price jump. Beyond the obvious price risk to the investor, this can impact funding arrangements where illiquid securities are used as collateral, as a sudden price jump can necessitate the investor posting more collateral. For hedge funds, illiquidity poses problems as well in reporting accurate NAVs and at worst, presents hedge fund managers with too much discretion to manage the returns they report.

In a 2001 paper, Lo argued that one method to assess the liquidity risk exposure of a hedge fund is to examine its returns for autocorrelation. A security that trades in an efficient market should show no autocorrelation: the current price should incorporate all information about the past, and the return in one period should have no bearing on the return in a subsequent one. If returns are correlated from one period to the next, this is a possible indicator of poor pricing information in the securities that are held.

Lo examines a Box-Ljung statistic on hedge fund returns to test the null hypothesis that there is no autocorrelation. He demonstrates that a pool of mutual funds shows no significant statistical evidence against this hypothesis, consistent with the intuition that these funds should be primarily invested in highly liquid securities. Among the hedge funds Lo tests, only those labeled “Risk arbitrage” appear consistent with the hypothesis of no autocorrelation, again consistent with our intuition that such funds would be mostly invested in highly liquid securities.

To be fair, the story with hedge fund returns is a bit more subtle than with individual securities. It is certainly possible that an arbitrage fund trades in liquid securities, but takes advantage of arbitrage strategies that persist for some amount of time. Such a fund would tend to show runs of better returns, while a

strategy was profitable, followed by runs of worse returns, while the strategy stopped working and they were searching for something new. This fund would produce autocorrelated returns, but not necessarily indicate that the fund had issues with liquidity.

In a subsequent analysis, Getmansky, Lo, and Makarov (2004) further investigate the interaction of liquidity, serial correlation, and deliberate smoothing of returns. Key among their findings is that illiquidity is the most common implication of serial correlation in hedge fund returns. We thus utilize the Box-Ljung statistic to investigate the presence of illiquidity problems in funds during the subprime crisis.

In search of autocorrelation

In Table 1, we present the p-values for the Box-Ljung test on a variety of return series.³ The first two series no doubt represent liquid assets, and their p-values indicate no inconsistencies with the assumption of zero autocorrelation. The next series, the broad Lehman ABS index may well contain any number of illiquid securities but as a portfolio, it does not show any signs of autocorrelation either.

We next construct a hypothetical illiquid fund using the ABS index. Suppose we trade the securities in the ABS index, but that we only receive new price information once every three months. Most likely, we can accurately recognize interest income, but we only update the clean prices of our bonds sporadically. To mimic this situation, we decompose the index total returns into interest income and price return. For our hypothetical fund, we use the correct interest income, but compute the price return by only

³The p-value is the minimum level of significance at which our null hypothesis (no autocorrelation in returns) can be rejected.

Table 1: Box-Ljung statistics, October 2003 through February 2007

Return series	p-value (%)
JPY-USD FX	29.5
S&P 500 Total Return Index	20.6
Lehman ABS Total Return Index	30.4
Lehman ABS, lagged prices	2.8
HFR Fixed Income funds, 25 percentile	41.3
HFR Fixed Income funds, 50 percentile	8.9
HFR Fixed Income funds, 75 percentile	0.8
Bear Stearns High-Grade Structured Credit Strategies	10^{-13}
Galena Street Fund	10^{-5}
Basis Pac-Rim Opportunity Fund	24.3

updating the prices every third month. The resulting return series produces a p-value of 2.8%, meaning that at a 5% significance level, we would reject the null hypothesis of zero autocorrelation. Thus, behavior we might expect from an investor in illiquid bonds can manifest as autocorrelation in returns.

We then examine returns on the 165 funds in the HFR database that follow one of the four fixed income strategies (fixed income arbitrage, mortgage-backed securities, high yield, and fixed income diversified) and that reported return data for the period from October 2003 through February 2007.⁴ The median p-value across these funds is 8.9%, while the 75th percentile is 0.8%, meaning that at least a quarter of the funds fail our test for zero autocorrelation at 1% significance.

Most interesting are the outliers. Of the five fixed income funds whose p-values are smaller than 10^{-7} , two have closed: the Bear Stearns High-Grade Structured Credit Strategies fund (p-value of 10^{-15} ,

closed in June) and the Galena Street fund (p-value of 10^{-7} , closed in July). The other three funds continue to report solid returns, though all reported lower (but still positive) returns for June, the latest data available. At very least, the test for illiquidity has shown to be a useful filtering mechanism.

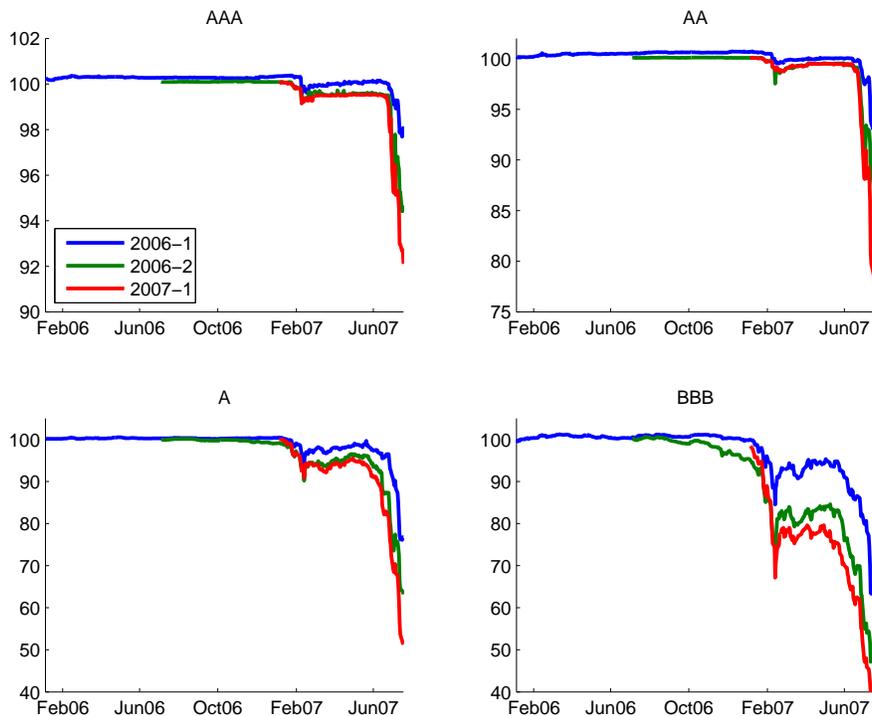
Interestingly, another fund that has announced significant losses due to the subprime crisis, Basis Capital, does not show signs of illiquidity based on our statistics. Alas, just trading in illiquid securities is not the only way for a fund to fail.

Enter the ABX

We have argued both intuitively and empirically that illiquidity, in particular in the form of difficulties valuing securities, is a critical component of the current subprime troubles. But if the problem is simply that there are no prices to observe, and therefore nothing off which to base risk estimates, then

⁴October 2003 is the earliest date for which the Bear Stearns High-Grade Structured Credit Strategies fund reported returns. We cut off the data in February 2007 so as to not include the market upheaval related to the subprime crisis.

Figure 1: ABX quoted prices, grouped by credit rating



what are we left with? Though it is far from a perfect solution, we should at least examine the one set of subprime-related instruments that does trade with some frequency, the ABX index contracts.

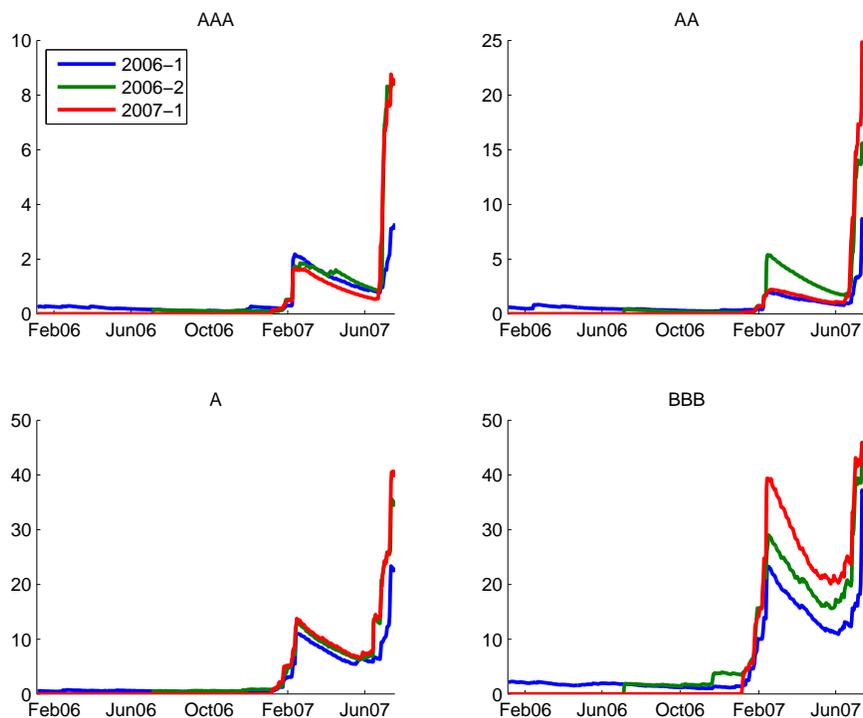
The launch of the ABX indices in 2006 was the natural extension of the credit derivatives index technology from corporate credit to subprime mortgages. Every six months, a series of ABX index contracts is specified. A series is defined by a reference basket of twenty home equity ABS deals; each deal is required, among other things, to include securities rated AAA, AA, A, BBB, and BBB-. That a series is composed of specific deals, rather than classes of securities from specific issuers (as with the corporate credit indices) is an important point. Each index thus reflects a vintage—the state of subprime mortgages originated at a specific time.

⁵The series initiated in the beginning of 2006, containing mortgages originated in late 2005

A look at the state of the ABX pools confirms our comments about heterogeneity. Remittance reports for the individual ABS deals disclose the proportion of mortgages that are either delinquent or in foreclosure. As of July, in the 2006-1 Series,⁵ this proportion ranges from 9.1% to 24.6% across the twenty ABS deals. For the 2006-2 Series, whose mortgages are roughly six months newer than those of 2006-1, the proportion ranges from 10.2% to 24.2%, and for the 2007-1 Series, from 9.2% to 23.6%. As varied (and high) as these numbers are, we are reminded often that most of the mortgages underlying these contracts will see their interest reset from a low fixed to market-determined adjustable rate in the next year, with more delinquencies surely to follow.

Within an index series, individual contracts reference securities of like ratings from the basket of

Figure 2: Estimated volatilities (in %) on ABX quoted prices, annualized



deals. In a specific index contract, a buyer of protection commits to paying a fixed premium on the notional amount being protected. The seller of protection pays both principal writedowns and interest shortfalls from the reference pool of securities. To the extent that interest shortfalls are paid back to a security later, the buyer of protection then reimburses the seller. Further, and different from the standard corporate credit derivatives, the notional of the contract decreases over time as the underlying securities pay back their principal.

Most importantly, as with the credit indices, a community of dealers is committed to providing liquidity in at least the most recent series of contracts. Thus, this is one instrument in the subprime space for which we can reliably observe prices. We plot the price history for the ABX contracts in Figure 1.

These figures appear regularly now in any press coverage of the subprime market.

We also display the volatility⁶ in Figure 2. The volatility estimates shed light on the distinctions between series (or vintages), as well as between the market events of February and those of today. For the poorest quality contracts (BBB-rated and below of the more recent vintages), the market volatility today is comparable to that of late winter. In contrast, for the better quality contracts, recent events have produced significantly higher levels of volatility.

Across vintages, we see similar behavior for the contracts covering the poorest ratings. This indicates that even if the pools indeed vary by vintage, they are all expected to produce large enough losses to significantly impact BBB-rated tranches of ABS

⁶Calculated using an exponentially weighted moving average with decay factor of 0.97

deals. At the best quality, there does appear to be significant tiering, with volatilities in the 2006-2 and 2007-1 Series about three times as high as in the 2006-1 Series, indicating a perception that the tail event losses required to impact the AAA-tranches are much more likely for the more recent vintages. Even so, it is worth reflecting on the absurdity of annualized volatilities on baskets of AAA-rated securities of even 3% (for the 06-1 Series) and as high as 9%, and on baskets of BBB-rated securities that are as high as many individual equity volatilities.

More than a pretty face

Beyond providing writers with material for figures in a market without prices, what else might we use the ABX for? The index is marketed as an efficient way to speculate on ABS and as a mechanism to hedge a broad portfolio. We should also consider its applicability for valuation and risk analysis.

For valuation, it would be nice to rely on the ABX as a proxy with which to value all subprime ABS; at very least, we should use it to keep us honest. Returning to the Bear Stearns High-Grade Structured Credit Strategies fund, we see that the fund reported a return of about 1.5% in February, a month where spreads on most of the ABX contracts tripled, and then a loss of about 3.5% in March, when ABX spreads tightened. One possible explanation for this pattern is that the fund was short subprime, but it is fair to say that had they maintained such a position, they would probably not have made headlines. Another explanation is that they were long illiquid subprime-backed bonds, and had hedged their positions by going short through ABX contracts. If they had taken the mark-to-market gain on the liquid

ABX position, but not observed any price changes on the illiquid positions, they could have arrived at such a return pattern, while still being positioned for further trouble.

From two monthly returns, it is impossible to know what the explanation was, and the point here is not to bury this specific fund any further. The point is that as risk managers or investors, we should question such patterns: if a valuation or self-reported return does not square with our understanding of the positions and the moves in the market, we should push for an explanation.

For risk, we may certainly go through the exercise of mapping a broad portfolio to the set of ABX prices we observe, and calculating portfolio risk through this proxy. Cynically, it is unlikely that such an approach would be proven wrong, as our whole problem stems from not observing prices on the broad portfolio. However, the approach should fail eventually, since the asset class is heterogeneous, and at some point pricing on individual securities will reflect their true value based on their specific underlying mortgages. But in a market that is driven more by technicals (the aversion for the asset class as a whole) than fundamentals (the quality of the specific pools), it is prudent to assess the risk under the assumption that all of the class moves as one.

When it's over

When the current market upheaval subsides, there will certainly be hand wringing (in fact, there is plenty already) over the fact that most of the troubles in the world of subprime mortgages can be attributed to the perception that housing values would forever appreciate, and that any borrower who had

trouble servicing a mortgage could always sell his house and pay the mortgage back.

But beyond lessons about the home price bubble, we should also take away some insights about investing in illiquid securities. If there is no trading in a security, no rating change, and no change in its indicative price, it is easy to slip into the misconception that the security carries little risk. But a risk manager can do more than compute poor volatility estimates based on a string of false returns.

Where fund returns are available, there are useful statistical indicators for liquidity issues. And there is the common sense that funds generally should do well when the market they trade does well; this not being the case is not likely itself a sign of fraud, but should be a signal to ask for explanation.

For those holding the illiquid securities, there are mark-to-model approaches that seek to value and forecast risk given the specifics of the individual securities. But this is not an easy task, and moreover can ignore the fundamental source of risk when an

entire asset class is punished. As easy as it is to dismiss a one-size-fits-all approach of modeling based on a small set of liquid prices, it is much better than doing nothing, and in the atmosphere today, with technicals dominating fundamentals, perhaps even the best choice.

Further reading

- Getmansky, M., Lo, A., and Makarov, I. (2004). An econometric model of serial correlation and illiquidity in hedge fund returns. *Journal of Financial Economics*, 74(3): 529–609.
- Lo, A. (2001). Risk management for hedge funds: Introduction and overview. *Financial Analysts Journal*, November/December, 16–33.
- Lucas, D. and Murray, T. (2007). ABS CDO collateral losses, Version 2.0. *UBS CDO Insight*, August 9, 1.