BANK OF ENGLAND

Speech

Tail risks and contract design from a financial stability perspective

Paper given by

Paul Fisher, Executive Director for Markets, Member of the Monetary Policy Committee and Financial Policy Committee

Co-authored with Patrik Edsparr, Visiting Fellow, Bank of England

Clare College, Cambridge

1 September 2011

The authors would like to thank Ravi Mattu (PIMCO, Newport Beach), Morten Hviid (UEA) and colleagues at the Bank of England for their comments on a draft of this paper. The views represented in this paper are those of the authors and should not be taken as the policy of the Bank of England or to represent the views of any other person at the Bank.

The wider theme of this conference is about what we have learned from the recent crisis. There have been many lessons. Some are not new but just a re-learning of old lore: 'banks need to hold adequate capital'; 'real-estate prices can fall dramatically'; 'financial institutions need to avoid excessive risk taking'. The authorities are pursuing a long list of regulatory initiatives to address the externalities arising from risks in banks and markets, including Dodd-Frank in the US, the European Market Infrastructure Regulation (EMIR) in Europe, the Independent Commission on Banking (ICB) in the United Kingdom and the various Basel capital and liquidity rules internationally. And the Financial Stability Board has taken on a role in co-ordinating much of the other international effort. Academic research also has a large part to play in this process, in both identifying the issues and proposing or evaluating policy responses.

A lot of discussion is taking place around disentangling different forms of complexity in the market place. For example, regulators are promoting more transparency around funding and legal entity structure within complex financial institutions. And there is also a growing body of work that is exploring clearing houses and exchanges as important structural vehicles to mitigate the increased inter-connectedness driven by the growth of derivatives[1].

The approach taken in this paper is to look at some of the details of how risk-taking is executed and the underlying market structures. These micro-foundations can have a profound impact on systemic stability beyond the normal consideration of formal regulations. In particular, we focus on one aspect of market structure: contracts where, because of a failure to take into account how the financial system as a whole operates, the true value of the contract is different from what it was intended to be - by at least one of the counterparties who struck the contract. This can arise either because in states of the world in which a particular contract is designed to have value there is high correlation with other events or where, in that state of the world, full adherence to the legal structure would cause large unintended consequences in terms of signalling or reputational damage.

Tail risks with counterparty risk correlation

The first category of contracts falls under the broad headline of 'wrong-way' risk in tail events, and is purely statistical in nature. The particular issue is focussed on contracts that can be perceived of as providing insurance yet, in a stress scenario, when the contract is sufficiently in the money to be worth calling, there is a significant conditional probability that the insurance provider is no longer in business to pay out. In what follows, we look how such problems arise using practical examples and consider what lessons should be drawn.

[1]For example, the Squam Lake Working Group on Financial Regulation, July 2009, "Credit Default Swaps, Clearinghouses, and Exchanges".

When investor 'A' decides to buy protection from bank 'B' on exposure to a third company 'X,' she may or may not explicitly take into account that in the event of default by X, B may in fact also be in trouble. Whether or not she does so would depend on the reason for the trade. If A wants default protection then that default correlation certainly should be taken into account[2]. But if, for example, the goal is just to reduce the price volatility of holdings of a certain corporate bond issued by X, this default correlation may not be seen as that relevant.

As a practical example, take the case of **a pension fund managing its corporate bond portfolio**. Suppose the pension fund owns a large amount of a particular corporate bond and wants to reduce this holding, but the liquidity for outright sales is limited. The fund may take a view that it will have less price impact by gradually selling the bond in the cash market over a period of a few months but, at the same time, it would prefer to take off the market risk immediately. If that is the case, the fund may take out a standardised CDS (Credit Default Swap) contract with a bank, probably written on a basket of corporate bonds issued by the same corporate. There would be basis risk from any mismatch in the bonds, so the CDS is not an exact hedge, but in most scenarios, holding the bond and the CDS contract together will ensure a substantial reduction in the net value of the market position, compared with holding the bond alone. The fund could then unwind the CDS contract simultaneously with slowly selling the cash bond. In short, the superior liquidity of the more homogenous CDS market (which benefits from relatively elastic gross supply) enables the pension fund to reduce the execution costs of its portfolio management whilst not sacrificing its nimbleness in terms of market risk management.

In this example, the strategy seems reasonable. The probability that the corporate defaults in the few months the strategy takes to execute is probably pretty low, and the joint probability of the corporate and bank defaulting is even less. In essence, the CDS protection is not really being used to insure against default but to reduce market price volatility.

Let us now turn to a second example of **a bank hedging its corporate lending**. Banks lend to a broad range of companies. In many cases they would be expected to hold these loans on their books until maturity (as opposed to selling them off to other investors). In order to reduce its measured outstanding credit risk, a bank may choose to purchase off-setting CDS protection, probably on a portfolio of loans. By so doing, the bank can make greater use of its internal credit risk limits and may get regulatory capital relief. In most cases, another leveraged market participant (bank or hedge fund) would write the CDS. (In some cases the insurance provider may be an unlevered institution, for which most of the concerns below would be reduced.)

In the first example, the pension fund was concerned about short-term price swings. In this second example, the bank should be concerned to protect itself against clusters of defaults. But if the risk management of the

[2] When the Bank of England takes collateral, it is solely to guard against the risk of default and so all correlations in that state of the world are relevant to what collateral is taken and the haircuts applied. See Breeden and Whisker, 2010, "Collateral risk management at the Bank of England", *Quarterly Bulletin*, Q2, pp94-103.

ensuing portfolio of loans and CDS is not done with serious attention paid to the joint risk profile of the corporate borrowers and the providers of the CDS protection, the true net risk position might be worse than it appears. That could fool regulators as well as the bank's own risk management function. Any risk management framework that is driven by the correlation structure of short-term moves alone would most likely be proven too rosy in the event of clusters of defaults. The issue is that the correlation structure locally may be very different than that in the tail of the distribution. Of course, there are many practical problems in analyzing what could happen statistically in the tails, since these scenarios are, by definition, rare and the pay-offs may be non-linear[3].

Taking this second example a little further, assume that most big banks follow similar strategies. Most banks will then end up with a more diversified credit portfolio than if each bank was only holding the loans it originated itself, but they may also end up with similar portfolios of risk. In such a scenario, the system can easily absorb small shocks, but because of the diversification, a really big shock may threaten to bring down all of them instead of just the ones that happened to have originated the most problematic loans. In a very inter-connected financial system, it can be almost self-fulfilling that when the market starts to think some of the banks are in trouble, in fact all of them are[4].

A third example in this section is the dynamic hedging strategy of a bank's credit risk arising from its derivatives portfolio (**Credit/Counterparty Valuation Adjustment: CVA hedging).** This normally refers to the credit risk of a counterparty that is not directly offset by some form of collateral[5]. Suppose a corporate has issued a fixed rate bond because investor preference makes that the best value, but really it wants to have a floating rate liability. A bank will accommodate this by providing an interest rate swap. If interest rates go up, that swap may generate a net value owing to the bank. Receiving that value depends on the corporate not defaulting and so the bank will have a credit exposure to the corporate. But the corporate cannot be expected to post collateral – it will just regard its debt as having been converted to floating rate (and typically it won't hold any collateral it could post anyway). A bank that is active in these interest rate swaps will end up with a large number of such mark-to-market exposures. Collectively, the resulting exposures constitute a credit portfolio which changes with the mark-to-market values of the underlying derivatives contracts. To manage the credit risks embedded in this so-called CVA portfolio, most banks will hedge dynamically in the CDS market. This is usually a very effective risk management technique during periods when CDS markets function normally. But it is worth emphasizing that the embedded 'jump to default' risks in credit products, coupled with liquidity issues during times of extreme stress, can make it very difficult to execute a dynamic hedging strategy in stressed conditions.

[3] See Haldane and Webber, July 2008, "Risk reallocation", *Risk*, for a wider discussion of the problems in pricing tail risk and Shleifer, Gennaioli, Vishny, 2011, "Neglected Risks, Financial Innovation and Financial Fragility", *Journal of Financial Economics* for the possible consequences.

[4] See, for example, Haldane, April 2009, "Re-thinking the Financial Network", speech delivered at the Financial Student Association, Amsterdam, for a discussion on the inter-connectedness of the financial system.

[5] See the *Quarterly Bulletin*, 2010 Q2, Vol 50, p81.

In addition, and as already noted, CVA hedging is likely to be undertaken by buying CDS from another bank or a hedge fund. In this example we have tail-risk correlation not just with corporate and counterparty risk but with CDS market liquidity. If proper consideration to the correlation structure is not applied, the danger is clearly that the realized outcome in extreme scenarios would look very different from what was expected by senior management, regulators, etc., based on local analysis.

A fourth example from the recent crisis was the **'super senior' credit exposure from pools of US mortgage loans** reinsured with monoline insurance companies (MBIA, FGIC, AMBAC, etc) or other insurers such as AIG. Banks and other originators of leveraged credit risk (either in cash form of CDOs or CLOs, or in its synthetic forms) needed to warehouse excess 'super senior' tranches, which were considered virtually risk free, to support the origination business. After the volatility experienced in 2005, banks started to worry about controlling the mark-to-market volatility, even though the fundamental risk initially continued to be considered negligible. The idea came to lay off a significant portion of this negligible economic risk to reduce the short and medium term price volatility. Similar risk reduction processed largely worked for traditional corporate credit risk, but proved disastrous in the case of super-senior risk backed by US sub-prime mortgages. Insurance by the monolines proved practically worthless as they successively failed. And reinsurance with AIG would have proved worthless if AIG had not been rescued by the US Government. The problems affected other securities in unexpected ways - at one point, US municipal debt wrapped by monolines was trading at a discount to the same debt without the wrap.

Once a default event has occurred, the correlation problems with such contracts become obvious. The question is why they are not apparent *ex ante*. The normal motivation for insurance is risk aversion to large losses arising through tail events. If such an insurance contract is worth having, one must be able to envision the tail event happening and what the circumstances might be. Why was it so hard to anticipate what would happen if there were losses on large numbers of bonds insured by monoline insurers which had relatively small amounts of capital?

One reason why such risks may not be appreciated *ex ante* is over-reliance on local risk measures such as VaR (Value at Risk) or other historical average correlations, and not enough on severe stress tests over different horizons. For example, in the United States it was widely believed that house prices would not fall significantly nationally (it had never happened after the Great depression in the 1930s). We assert that, if one had stressed bank capital in the US with the explicit assumption that national residential real estate prices could go down 30 to 40%, and in some local markets more, most of the losses for US banks could have been correctly predicted (at least if the complexities of structured credit products were properly assessed[6]). The collapse of functioning inter-bank money markets and other 'run-on-the-bank' phenomena obviously added further costs, but the first-order effect was caused by residential house prices falling. The key, however, was to correctly predict what that kind of price drop would do to other market participants, and

[6] This caveat might seem an important element of hindsight. But one reason why the complexities were not addressed was the assumption – based on local correlations – that the fundamental risk was negligible.

market structures, as well as to the individual firm. Extreme stress tests have served a crucial purpose after the crisis to re-establish confidence in banks, and might usefully play a wider role in more aspects of on-going bank oversight.

An additional concern is that many of the transactions suffering from this risk may have been entered into for portfolio management reasons (such as the pension fund example) or regulatory capital relief and other 'window dressing' purposes rather than genuinely to protect against default. Net exposure numbers quoted on earnings calls, or stress losses reported to regulators, can all be made to look more palatable if the insurance is apparently in place, while the true risks in the event of default may not have changed much.

These issues raise general questions for market structure with tradeoffs that are far from trivial. Is it beneficial for, say, banks to lay off sovereign risk with other banks and other leveraged financial institutions such as hedge funds, or does the ensuing web of contingent exposures across leveraged institutions *increase* the systemic risks? It seems that many of these arrangements increase diversification in the case of small and medium size stress scenarios, but may actually exacerbate the systemic implications in the case of very large shocks[7]. These arrangements also highlight the delicacy of the appropriate regulatory structure. In many instances, subject to sufficient diversification, these tail risks might be better borne by so-called 'real money' investors, i.e. insurance companies, pension funds and traditional long-only bond and equity funds, as these investor categories operate with little or no leverage. In case of massive shocks, these investors could be better suited to absorb the losses, as there is limited leverage to trigger further spillovers through defaults or panic driven asset sales. Many of these investor groups are, however, subject to an extensive regulatory framework, making them unable to hold such risks in their portfolios, in many cases even as a very small fraction of total assets. There is a potentially difficult trade-off here between public policy objectives of appropriate investor protection and systemic stability considerations. That probably deserves to be debated more fully. It is clear, however, that highly leveraged institutions such as banks and hedge funds are not really suited to be the ultimate repositories of extreme tail risk[8].

What does all of the above tell us? The answer may seem trivial *ex post* but if these issues had been highlighted and subjected to more transparency, we believe that the ultimate outcome probably could have been different. A lot of the surprise seemed to come from the fact that virtually all risk management had been done within a 'local' framework, rather than genuinely extreme stress tests. If regulators, rating agencies or, for that matter, bond and equity investors had demanded analysis based on extreme stress tests, many of the repercussions in the system could have been identified. The main point is that a stress test has to be internally consistent. As an example, if you stress test a 30-40% fall in US residential real estate prices and that shows your institution would be in trouble, you need to make consistent assumptions

[7] A general idea explored by Viral Acharya, 2009, "A Theory of Systemic Risk and Design of Prudential Bank Regulation", *Journal of Financial Stability*, 5(3), in a somewhat different context. See Also Haldane, 2009, *Op Cit* and Nier, Yang, Yorulmazer and Alentorn, 2008, "Network models and financial stability", Bank of England Working Paper No 346.

[8] Nor, of course, should the public sector be expected to socialise losses from tail events!

about the deteriorating credit conditions of your counterparties, anyone that provides reinsurance and general market conditions.

At some level, there is always a stress scenario that forces bankruptcy. Perhaps the greatest sin in the years preceding the financial crisis was blatantly ignoring what really would happen if the 'unthinkable' drop in US house prices actually happened[9]. Some major institutions seemed to have had no grasp of how risky their exposures really were, despite spending millions on risk management and while being inspected by regulators on an ongoing basis. By addressing these issues, the stress tests implemented by the Federal Reserve in the United States, and by the European Banking Authority in Europe, could play a key role in re-establishing confidence in the banking system.

Reputational risks in tail events

The second category of contracts we want to examine can be put into more of a game theoretic framework, where the perceived pay-off in the tail event is offset by 'unforeseen' costs. In this section, we are focusing on behaviour that, at first glance, may seem irrational, but is best understood as being driven by classic time inconsistency. In particular a market participant may voluntarily choose not to enforce a contract that is 'in the money', if the reputational repercussions are perceived to cause more damage than whatever could be gained financially by enforcement.

During the crisis, banks moved a significant amount of their mortgage-based assets off-balance sheet into **Structured Investment Vehicles ('SIVs').** Most of the funding for these SIVs came from publicly sold asset-backed commercial paper ('ABCP') and other medium-term notes ('MTN'). There was a small amount of equity in the form of a junior 'capital note,' which split the excess return with the asset manager. From their invention by Citigroup in 1988 and until the crisis, there was sufficient spread between the assets and the funding to generate a reasonable return to attract investors to buy the capital notes. In addition, there was typically a partial liquidity back-stop facility provided by the originating bank. When all the capital notes had been sold, there was no economic risk retained by the originating bank except the back-stop facility, and only a management fee continued to accrue to the asset manager (typically, the originating bank).

When the crisis hit in autumn 2007, one of the first casualties was the closure of the ABCP markets and funding for the SIVs dried up. Beyond the liquidity lines, banks generally had very limited legal obligations to the SIVs (eg for credit losses on the assets) but most decided to accept responsibility for the assets and absorbed them back on balance sheet or at least fund them directly. This was obviously done for reputational reasons and was viewed as less costly to the franchise than walking away.

[9] Since the crisis, many firms are now considering the previously unthinkable eg the sustainability of sovereign debt, liquidity in core markets etc.

Part of the explanation for banks' behaviour was that these structures were not necessarily designed to be profit centres, but rather to enhance capital efficiency for the bank. There was essentially a massive maturity mismatch between the banks' underlying assets and their funding. And the SIV structure effectively enabled the bank to isolate that mismatch and hence support its other businesses with the embedded lending in the assets that were being put into the SIVs.

The perception was that the quality of the underlying assets was so good that the risk of a real loss (as opposed to some shorter-term mark-to-market fluctuations) was negligible, and that the returns accruing to the capital note holders and the asset manager were largely due to the funding arbitrage. Another way of explaining it would be that, because of the difference in liquidity, the 'buffers' in the structure, the capital notes and the liquidity back-stop provider were effectively earning a risk premium for absorbing this liquidity mismatch between assets and liabilities. This view of the world made most of the capital note holders believe that the risk of any loss was very limited.

We now know all of that changed when sub-prime mortgage assets had become a significant portion of the assets in a typical SIV. The US housing market had started to slide and the lax underwriting standards started to show their impact on recovery rates. Suddenly, investors in the ABCP market came to realize that the risk of actual loss was very real and that a game of 'musical chairs' was unfolding in the roll-over process of the short-term debt. At this juncture, one would think that the optimal economic behaviour for the banks with outstanding SIVs, would be to let the SIVs unwind according to the legal construct in place, rather than accept responsibility. In fact, all the banks except one decided to collapse the structures and repurchase the securities[10].

From an investor protection standpoint, one may have drawn a sigh of relief, as any potential issues about misrepresented risk profile disappeared, but from a financial stability perspective, it was obviously disturbing. In a short space of time, billions of assets showed up on already over-extended bank balance sheets.

One may think it somewhat surprising that (almost) all the banks decided to absorb their SIVs, especially since it must have seemed likely that the SIV structure would not come back any time quickly as a viable funding structure. We believe that the main reason for doing this was that *not* doing so would have sent a distress signal to the market. In other words, if a bank chose not to absorb this problem, the perception would be that they simply could not afford to do it, thus telling the market that they were in even worse shape than previously feared. There may also have been an element of 'repeat game'. If one lets one's investors take the pain, then they may not return for future transactions.

[10] The one notable exception was Standard Chartered, which let its 'Whistlejacket' SIV unwind.

It seems plausible that if the SIV crisis had transpired during a time when most other markets had remained functioning, the decision making process might have been very different. 'Signalling' is at its most powerful when the level of uncertainty in the system is highest.[11]

We now want to turn to an example of a structure that did at least partially survive the crisis: **UK RMBS master trusts.** These are versatile structures that allow an issuer to transform a portfolio of assets into a very different investment product. In principle, the master trust purchases mortgages from the bank and issues tailored bullet-maturity securities. The assets and the liabilities can be very different, e.g. the demand for the securities may be stronger in US dollars than in sterling, so an FX swap converts the cash-flows into dollars. Or most of the mortgages may be floating rate, whereas investor demand is stronger for fixed-rate product, and an interest rate swap(s) converts the cash flows accordingly. Mortgages typically pay down somewhat randomly and relatively quickly, whereas investors prefer straight-bullet maturities. The master trust can accommodate this by allowing on-going reinvestments of the pay downs and conversely allowing the trust to put assets back to the bank if the cash-flows are not sufficient to make a particular principal repayment. The downside is clearly that the on-going linkages between the issuer and the trust remain very strong, and it is thus hard to identify where economic risks ultimately reside.

RMBS master trust securitisations are designed to be called on their bullet maturity date and that is how they were priced and traded. In one case that didn't happen: in November 2008 it was announced that the £40bn 'Granite' master trust had failed a non-asset trigger[12] which meant that Northern Rock would no longer provide new mortgages to the structure and it went into pass-through mode (early amortisation). Instead of owning bullet securities, the investors ended up holding potentially long-dated pass-through notes. Trading in the securities effectively ceased and, because of the maturity extension, the securities issued by 'Granite' fell precipitously in price – even though there had been no material credit losses. The subordinated securities would only be paid off after enough cash had come in to first pay off the senior bonds. As Northern Rock had already been nationalised, the dynamic around the decision to run-off 'Granite' was obviously quite different from what other banks were grappling with at the time. This event clearly 'burnt' a number of investors, many of whom have indicated that they will not return to purchase any master trust structure in future. In pure economic terms, it might have paid other banks also to not call RMBS when funding pressures were at their highest. But this would have destroyed the market and damaged relations with investors, possibly beyond repair.

Another example is the existence of **break clauses in derivatives contracts**. These clauses enable either counterparty to a derivatives trade to call for cash settlement of the mark-to-market of their contract at a particular future date. The main reason for this 'break clause' is that it reduces the contingent credit

[11] The issues arising from ABCP conduits were similar, although the structures were slightly different. The key difference from SIVs was that the underlying assets were largely purchased from third-parties and not necessarily in securities form.

[12] After the 'current seller share' had fallen below its minimum on two successive distribution dates –the company's own investment in the vehicle had not been maintained at a high enough level. See also Northern Rock's 2008 Annual Report and accounts P25/6.

exposure. Think in terms of a simple fixed-for-floating swap for 10 years. (This transaction could be to convert a company's fixed-rate bond issue into a floating-rate liability for example.) The bank pays a fixed rate in exchange for receiving a floating rate (say, LIBOR) from the corporate that issued the bond. If no break clause exists, the bank's credit department would price the contingent credit exposure to the corporate over the full life of the contract. This could be very high: a substantial rise in short-term interest rates would make the bank effectively owed a significant amount of money by the corporate. The bank needs to factor in that with this 'gain,' there is an associated risk of default of the corporate. In reality, one can think of all these contingent credit exposures as the modern derivatives book's 'loan portfolio'. It is equivalent to a traditional loan book, except that it dynamically changes in line with market prices. The credit risk is obviously bigger when the creditworthiness of a company is lower and when the maturity is longer.

The bank's credit department will therefore calculate a credit spread that it will charge on top of the 'risk free' swap level (that would be charged between financial institutions that fully collateralize any mark-to-market swings). For a longer-dated swap, this charge may be very large. Introducing a break clause allowing either counterparty to terminate the contract effectively shortens the maturity of the swap, since the bank (or the company, in case it is owed money and the bank's credit appears more shaky) can demand payment at that time. In reality, most corporates do not expect such a break clause ever to be exercised.

The concern would be that, if the corporate expects the bank to ignore the break clause regardless of the circumstances, exercising the clause would be a surprise for which the firm may not have put appropriate contingencies in place - perhaps even making default more likely. Or that the market may start to interpret the enforcement of 'break clauses' as a sign of weakness on the bank's part. In other words, such behaviour could be seen as so antagonizing to customers that a bank would only do it if it is absolutely desperate. Yet if the bank cannot in fact enforce the break clause, its credit pricing and risk management has been faulty.

Another example, comprising elements of the break clause and the SIV problem, is the behaviour of managers of money market funds when faced with the prospect of '**breaking the buck'**.

A money market fund calculates its net asset value (NAV) on a daily basis. The NAV is its price per share, which reflects the total value of the fund's investment holdings. Traditional money market funds seek to maintain a constant NAV (CNAV funds) ie they invest with the explicit goal of maintaining a stable NAV of $1.00 per share. Investors like this because it avoids any market price fluctuations in their investments.

A CNAV money market fund is said to 'break the buck' when its NAV falls below $1.00 per share (or the equivalent in a fund's respective currency). This is a rare event. Prior to the crisis, it had not happened to a US money market fund since the Community Bankers US Government Fund broke the buck in 1994[13]. The conventional perception was that the principal should be 'safe'. When money funds came under pressure

[13] In Japan, money market funds lost much of their appeal after principal losses due to Enron-related investments, and investors largely returned to bank deposits.

during the crisis, more than 60 asset managers (among others, Wachovia and Legg Mason) unilaterally made up more than $12 billion of losses[14].

On 16 September 2008, the Prime Reserve Fund, the oldest U.S. money market fund, wrote off 3 cents of losses on Lehman Brothers. Not being able to make up the losses, the fund was left with a NAV of $0.97, triggering a 'run' in the money markets. On 19 September 2008, the U.S. Treasury announced the establishment of a temporary one-year guarantee program to protect investors in money market funds. The funds that were eligible had to pay a fee to participate in the program, and any fund that had already 'broke the buck' could not participate.

This example raises several interesting issues. In keeping with previous examples, it illustrates the extreme non-contractual lengths some of the asset managers were willing to go to in order to protect their reputations. It also shows how the equivalent of a 'bank run' can happen in the shadow banking system. And it shows the potential importance of public sector support – another complicating factor that makes the true risks in a tail event difficult to assess.

As a final example in this section, we have Deutsche Bank's decision in December 2008 ***not*** **to call a lower tier 2 subordinated debt issue** (a €1bn 2004/2014 bond with a call date of Jan 16 2009), despite the fact that it was customary for banks to call such bonds at the first possible date. The reason that this is interesting is that it was actually economically rational in isolation for Deutsche Bank not to call, and it was fully within its right not to do so, but still, investors became infuriated and the consequent short-term adverse reaction almost certainly cost Deutsche and its shareholders a multiple of the gain it received from not calling the bond. Despite that initial investor reaction, however, Deutsche was able to keep market access and to successfully issue further capital instruments in the months thereafter.

The theme we have encountered in these examples is that many of the decisions taken during times of extreme pressure seem to be affected by (apparent) considerations outside of the immediate financial contract – at least, rejecting options which were perceived to be acceptable when the original contract was struck. Or that the financial consequences of particular decisions were more adverse for the decision maker (or more generally the markets) than anticipated. In that sense they are classic examples of time inconsistency. We have tried to illustrate how acting narrowly rationally when it is perceived to be outside the realm of 'what is normally done', can carry massive risks. The problem that this creates is that when we are dealing with tail events and disaster insurance, it is hard to determine *ex ante* what such insurance contracts (or options delivering insurance) are worth. How should investors and regulators evaluate these contracts? If the reputation and signalling implications of exercising are too damaging in case of extreme

[14] See Brewster and Chung, 17 September 2008, "Fear of money market funds 'breaking the buck' ", *The Financial Times*, and Moody's Investor Service, August 9.2010, "Sponsor Support Key to Money Market Funds".

stress, it would be naive to attribute full value to them when determining the risk profile and hence the appropriate capital buffer. It is also worrying from a financial stability perspective if these considerations force possible systemic implications to be more extreme rather than less.

Stress tests and contract design considerations

So what have we learned? The first set of examples above highlighted the need to critically evaluate contingent exposures: without true stress correlations, tail events will not be captured properly. We believe more emphasis should be put on this by investors, analysts and regulators alike.

Inevitably that means severe stress tests involving 'jump to default' risk and the need to consider counterparty and market liquidity risks, not just issuer risks in these circumstances. Some banks have told us that they think they should not be required to hold capital and liquidity to deal with such extreme tail events –leaving the public sector to be the capital provider of last resort. But that leads directly to moral hazard and excessive risk-taking. Tail events seem to happen far more often than people assume[15] and if the risks were properly acknowledged at the outset, many structures would be avoided or risks re-structured so as to limit losses in the event of tail risks. That has obvious implications for financial stability.

A crucial component of this analysis is the proper design of stress tests. Obviously, scenarios have to be rather draconian in order to serve the purpose of challenging the 'unthinkable', but at the same time there are difficult decisions to be made, for example, in deciding how much bank capital (contingent or not) banks should hold – and what the probability is of it being wiped out. We believe that it is better to have a collection of *ex ante* determined stress scenarios that illustrate banks' potential weaknesses publicly, even if the actual regulatory capital is not sufficient in all those scenarios. In other words, stress tests should not always be a check list 'pass' or 'fail' (after all, to make a bank fail or pass a stress test is just a question of scaling the test). Comparable, tail event stress tests could be an important piece in the information set that investors and regulators analyse to determine the relative value and risk profile of the institution.

The second set of examples above illustrated that contract design features which are exclusively relevant in the extreme tails are perhaps best avoided. The reason is simply that such features, which do not seem to matter in most states of the world, end up being treated as if they will never matter. In other words, market participants need to make sure that an exercise of an option is seen as just that, a normal exercise of an option. No more and no less. Sophisticated market participants should not be genuinely surprised (or feign surprise) by another market participant trying to optimise its behaviour consistent with a contract. The implications of a contract should be clear and the structures should be as transparent as possible.

[15] How many 'once-in-a-lifetime' events happen to us every year?

As the debate about adequate capital levels for banks rages on, we feel that it is important to keep these lessons in mind. Whatever is decided in terms of capital requirements, the numbers should be calculated with tail event analysis in mind.

A new example: contingent capital securities

One present issue is that, given banks will need to hold much more capital in future, how could that best be achieved? Some have called for very high levels of equity capital in banks[16]; others have been focusing on requirements to raise new equity before it is too late, e.g. mandatory rights issues[17]. We find contingent capital securities - so-called 'co-co's - to be a potentially attractive proposition for a number of reasons.

One has to evaluate contingent capital on how it would perform in a crisis. First, the funding for extra capital is already in place. In other words, no one has to scramble to execute a contingency plan, or convince someone to come up with new cash in a difficult situation. Obviously, conversion is consistent with a bank being under stress, but if conversion was always automatic, not discretionary, that could help to avoid signalling even wider problems.

Second, and related to the examples in this paper, because the contingent conversion feature is explicit, no investor should be able to say that 'equity conversion' could never have been anticipated. As new mechanisms are being explored, we believe that it is crucial that investors correctly assess and price the probability of conversion into common equity.

Third, sufficient level of contingent capital should enable a well-run bank to still operate with significant leverage and thus earn a healthy return. In the more extreme suggestions for the amount of contingent capital, one would very substantially change the return profile of common equity for banks, perhaps creating an entirely new investor category.

There are two lessons from this paper which should be applied to contingent capital instruments. First it has been suggested to us by market contacts that the trigger point in existing contingent instruments is such that many investors have bought them on the assumption that these contingent capital securities never will be called (or worse, that there will be official support before that point). If that were to be the foundation for this market, we believe that the very purpose of contingent capital may be subverted, creating a risk to financial stability in a crisis situation. The whole point of contingent capital securities should be that the

[16] Admati, DeMarzo, Hellwig and Pfleiderer, 2010, "Fallacies, Irrelevant Facts, and Myths in the Discussion of Capital Regulation: Why Bank Equity is Not Expensive," Stanford GSB research Paper 2065.
[17] Duffie, 2010, "Contractual Methods for Out-of-Court Restructuring of Systemically Important Financial Institutions", in *Ending Government Bailouts as We Know Them*, Scott and Taylor (eds). Hoover Press.

recapitalisation is triggered without any grand repercussions, making it easier than raising fresh capital in the market. If triggering the conversion were to cause the sort of damage reported in our examples above, then the market could be severely disrupted just when it was most needed. The contract design must therefore reflect a need for the trigger to be as smooth as possible.

One potential mitigant would be to have a range of trigger levels, so no single level becomes a focal point for indicating severe distress. It would be helpful if actual conversions happened reasonably frequently and not only in a rare crisis, demystifying the whole process. One would also need to make sure that the maturity profile was such that no dominant amount of contingent capital was rolled at any one point in time. There could be an on-going process of monitoring, rather than particular trigger dates that may cause a focus for a 'run' during times of stress.

Stress tests could also be one of the tools to force conversion of contingent capital securities by bank regulators. In addition to the accounting based capital trigger, capital requirements could be expressed in terms of what scenarios a bank has to be able to withstand, with capital securities triggered when the existing common equity is not sufficient.

One argument that has been raised against contingent capital, has been the lack of a natural investor base, reflecting market segmentation between 'credit', 'rates' and 'equity' investors. We believe strongly that market segmentation is an inefficiency which should not be taken as a given; there are already regular convertible debt securities with an investor base. New market segments can be created if the risk-return trade-off is appropriate. The new issues (Credit Suisse and Rabobank), have so far been very well received and we are aware of plans for funds being created to invest in these instruments. We see every reason to believe that new vehicles for owning contingent capital instruments can be created and existing mandates modified to allow them, as long as the market sees these securities as an asset class that will stay and reach critical scale, and subject to the normal market discipline of risk and return.

This plays into a second lesson from the crisis. It is obviously crucial from a financial stability standpoint that contingent capital securities do not end up largely in the hands of other highly leveraged financial institutions, where losses could cause further spill-overs and thus generate financial instability. We believe that regulators could play a constructive role in allowing a broad range of 'real money' investors to own sensible amounts of this systemic risk.

Conclusion

In this paper, we have given a number of examples taken from the financial crisis, where financial contracts undertaken for hedging purposes, or with optional features designed to protect the issuer or the investor, have not had the value expected at precisely the time when they needed to. These now look like classic time inconsistency problems coupled with partial equilibrium analysis. To avoid these pitfalls, a more

extensive use of extreme and holistic stress tests could be used to assess individual counterparties as well as systemic risks. It would also be worthwhile to debate further where these tail risks should be held, as leveraged financial institutions pose larger systemic risks than more traditional unlevered investor categories. Such considerations need to be carefully weighed against other investor protection concerns. We believe that a properly designed stress testing framework could provide a basis for a balanced discussion around future risk concentrations and systemic vulnerabilities.