Market Consistent Discounting

Interim report of the Market Consistent Valuation Working Party
to the Finance, Investment and Enterprise Risk Management Conference.
June 2008.

Working party members

Seamus Creedon (chair)
Iain Forrester
Parit Jakhria
Malcolm Kemp
Antoon Pelsser
Andrew D Smith
Colin Wilson

Abstract: This note is an extract from a larger forthcoming paper looking at appropriate discounting for market-consistent valuation of insurance and pension liabilities. Section 3 below summarises the working party's research on interest rate markets and the components of spreads, including expected default losses, default risk premiums, illiquidity premiums, inconvenience premiums and unexplained elements.

3 Capital markets

3.1 Government bond markets

3.1.1 Most governments borrow money by issuing bonds. The bonds are typically structured with regular coupons of equal amounts in the currency of issue, with a repayment of principal at the end of the period.

3.1.2 There are a number of alternative structures, including bonds with variable redemption dates, perpetual bonds and bonds whose coupons and / or principal are linked to inflation indices.

3.1.3 Government bonds for developed western economies are usually regarded as the lowest credit risk for their particular currency. This does not mean that government bonds are risk-free, but that default is very unlikely.

3.1.4 There are a number of ways in which government bonds can be economically in default while legally honouring their terms. This happens because governments write the law. For example, governments can introduce new withholding taxes on coupon or capital payments, can impose conversion of one bond for another, less valuable one, or net off payments against other debts they believe the bondholder owes. Some commentators regard inflation as a form of default on government bonds with payments fixed in monetary terms. Euro denominated bonds face uncertainty in the event that the Euro falls apart, or that an issuing government leaves the Euro zone and seeks to redenominated its national debt into a national currency. All of these measures are likely to harm the international reputation of the government concerned, leading to higher borrowing costs in future.

3.1.5 Nevertheless, defaults can and do occur, even on government bonds. Defaults are often associated with regime change or revolutions, with a new government repudiating promises made by a former regime. Regimes subject to economic embargos or sanctions often respond by suspending foreign debt payments, if not already in default. Defaults are more likely for developing countries, politically unstable regimes, regimes not enjoying good diplomatic relations with the debtholder's country and regimes with significant budget deficits. Debts denominated in foreign currency are often less secure than debts in domestic currency, because it may be more politically expedient to default selectively at the expense of foreign investment.

3.1.6 The effect of budget deficits is interesting. The effect of higher than previously expected budget deficits increasing a government's borrowing cost, is noticeable even in developed western economies. One explanation is that even government bonds contain a loading for default risk. Although the default event is remote, budget deficits bring it closer.

3.1.7 We now consider the issue of liquidity. Unfortunately, liquidity has several meanings. In the early literature, longer dated bonds were considered "illiquid" because the lender waits longer to get their money back. Early references to a "liquidity premium" would more accurately be called a "term premium" and form a partial explanation for the commonly seen pattern where bond yields increase with term. There are several possible explanations for the term premium, one of which is that short term price volatility poses a risk to bond investors, which is most severe for longer bonds. The term premium is then the market's reward for bearing that risk.

3.1.8 Yields on government bonds are explained by a number of factors besides the term of promised cash flows. More liquid bonds usually trade at a price premium. In this context, liquidity refers either to tight bid-offer spreads or to deep markets where large transactions have a small effect on the price. As illiquid bonds are inherently less attractive (other things being equal), markets talk of an "illiquidity premium", which of course refers to a premium on the yield, or, equivalently, a discount on the price. As in most markets, a bid price is a price at which a market maker will buy a bond, and the offer is the price at which a bond is for sale. The offer price is higher than the bid price.

3.1.9 The drivers of liquidity are subtle. Setting of bid and offer prices is a commercial decision by market makers, and is outside the direct control both of government and bond investors. Market makers face fixed costs of listing a particular bond. Volume traded is part of the pricing calculation; a market maker who expects to trade a particular bond in large volume may be prepared to offer a tighter percentage spread in order to cover the fixed costs or achieve a given profit target. Larger trades are usually in "benchmark" bonds, that is bond whose maturity is close to a key integer number of years, particularly if the bond is recently issued, a situation sometimes called "on the run". "Tap" issues, that is, issues likely topped up in future government auctions, are also more likely to attain benchmark status. This effect is self-reinforcing: liquid bonds are more useful as benchmarks because prices can be determined more accurately. The use of benchmark bonds for hedging also increases the traded volume, which in turn increases competition between market makers and low percentage spreads, sustaining the liquidity observed in the first place.

3.1.10 In addition, there is a reluctance among market makers to build up large positions in bonds because of the market risk exposure generated. Large bond issues usually enjoy deeper markets, and market makers are more likely to be confident about their ability to offload positions quickly in such bonds. On average, market makers hold stocks of smaller and less liquid bonds for longer periods. To recoup the risk borne, these are subject to wider bid-offer spreads.

3.1.11 These liquidity effects create the impression of an inverted demand curve. Usually a demand curve is downward sloping, that is, the higher the price, the less is demanded. In the case of gilts, this may go into reverse, as larger issues are generally more liquid, and so trade at lower yields relative to smaller, less liquid, issues.

3.1.12 The number of positive feedback loops in relation to liquidity gives rise to some odd price behaviours. We will later see that some of these effects also spill over into the repurchase market.

3.1.13 Government bonds are usually low cost from an administrative and regulatory point of view. For example, government bonds are given zero weight in bank capital calculations. Furthermore, investors in government bonds often see no need for credit risk management systems, given that these systems are typically geared to corporate debt.

3.1.14 This section has focused on conventional government bonds. Some governments also issue inflation-linked debt. Just as government bonds are a possible source of reference rates for liability valuation, so index linked bonds provide a possible source of reference inflation assumptions. Index linked bond markets are typically less liquid than the corresponding conventional government bonds.

3.2 Corporate bond markets

3.2.1 Corporations also turn to bond markets to raise cash. Like government bonds, corporate bonds usually consist of level coupons paid quarterly, semi-annually or annually, with redemption of principal at maturity.

3.2.2 Corporate bonds may be secured on specific assets, or secured on a floating pool of assets. Secured corporate bonds are sometimes called “debentures”.

3.2.3 Corporate bonds may be issued in tranches with differing priorities. These may be referred to as “senior”, “mezzanine” or “junior” debt. Debt may also be issued subordinate to other creditors. Examples include the issuance by insurers of subordinated debt, which ranks behind policyholder contractual payments, or contingent loans whose repayment is triggered by a certain level of statutory surplus.

3.2.4 As with government bonds, there are a variety of indexed corporate bond structures in issue. However, the issues are typically small and illiquid.

3.2.5 For several reason which we explore below, markets usually value corporate bonds at a price below that of a government bond with the same term and coupon. The yield on any bond is the internal rate of return than reproduces the price from the promised cash flows. As corporate bonds trade at lower prices, and price is a decreasing function of yields, the yields on corporate bonds tend to be higher than comparable government bonds. The difference is called the “spread”, or more fully, the spread of corporate bonds over government bonds. Corporate bond investors do not always earn the whole spread, but they will do so if the corporate is held to maturity and pays the promised cash flows in full.

3.2.6 Corporations, like governments, can default on their debt, but corporations typically do so more frequently. Governments of western, developed economies are usually regarded as more credit worthy even than the largest multinational corporations. On the other hand, some large corporations may enjoy stronger credit standing than some developing countries. And at the bottom end of the spectrum, the likes of Enron and Zimbabwe compete for the wooden spoon. In this paper we will decompose spreads into expected default losses plus a premium. The expected default losses are usually estimated from historic default studies. The premium then represents the expected return for the buy-to-hold investor, averaged across paying and defaulting future outcomes.

3.2.7 Several rating agencies opine on the credit worthiness of bonds, and while these ratings are not infallible, enough users are prepared to pay for this information to keep the rating agencies in business. Their role has been reinforced by regulatory formulas explicitly incorporating rating bands, for example the spread widening stress tests under Solvency II.

3.2.8 As corporate bonds are typically subject to greater default risk than government bonds, so the prices are typically more volatile. Fluctuating market views on future default losses represent a source of additional source of risk for bond investors. Investors require a compensation for this risk in terms of a higher expected return. This implies that the spread on a corporate bond, relative to a government bond, should not merely correspond to differences in expected default losses; the corporate bond should provide an additional premium in the yield, over and above expected losses due to default, to compensate investors for uncertainty in estimated default losses.

3.2.9 Like government bonds, corporate bonds are subject to varying degrees of liquidity, according to commercial decisions that market makers take. Investors prefer liquid bonds and corporations are keen to minimise borrowing costs, so they seek to make their bonds as liquid as possible. While issuers cannot control market makers' spreads, they can encourage trading by issuing in size and mimicking the coupon frequency and maturity of benchmark government bonds. Such bonds tend to be more actively traded because the existence of a government bond reference makes pricing more transparent. However, corporate bonds are intrinsically less liquid because of their idiosyncratic character, that is, no market maker in practice holds stock of all bonds in issue. We have already discussed how part of the bond's yield premium can be attributed to uncertainty in default losses. Another portion may be a form of compensation for illiquidity.

3.2.10 However, in the case of corporate bonds, there is a further element of liquidity risk which is less relevant for government bonds. This risk is information asymmetry. If some information, and for example, let us suppose it is bad news, becomes available to part of the market, that part will try to sell the bonds to market makers. As this information spreads, the price falls. Potential buyers take the price fall as a signal of information asymmetries, and defer purchase until the information is fully disseminated and market prices stabilise. Market makers are left with losses on positions they cannot unwind, and the result is a "crunch" where spreads explode and nobody has the confidence to trade close to previous prices. Investors are concerned about both types of liquidity risk, but the potential losses are far greater in a crunch scenario than the routine transaction costs in normal conditions. A liquidity crisis on one particular asset may be of little concern, since a diversified investor seeking cash flow can choose which asset to sell. The greater concern relates to systemic liquidity crises that affect many assets simultaneously.

3.2.11 It is plausible that a large part of the liquidity premium on corporate bonds is attributable to the crunch risk. However, we have been unable to find literature references to support this assertion. The empirical difficulty is that, while bid-offer spreads are observable in normal market conditions, crunch scenarios are hypothetical, are reinforced by feedback loops, but occur infrequently and so are difficult to calibrate or model. Any attempt to explain yield spreads by reference to crunch losses is dependent on an unavoidably speculative model of how crunches behave. However, we can guess at factors that might make a bond particularly vulnerable in a crunch situation. Information asymmetries are most likely to arise where financial reporting is limited, where a bond is unsecured, or where structures are complex or opaque. Asset backed securities, often making use of unlisted special purpose vehicles, logically trade at higher yields than straightforward bond issues from listed and well-rated issuers. Government bonds are usually untouched by credit crunches as the negligible default risk implies there is no relevant information which could be asymmetric. Indeed, government bond liquidity could even increase during crunch scenarios because of a flight to quality for new money. On the other hand, the reverse affect could apply if the government's response to a liquidity crisis is to offer cheap finance or guarantees at taxpayers' expense.

3.2.12 We wish to draw a distinction between the premium for default uncertainty and the illiquidity premium. The risks are difficult to separate, as both materialise in the form of a lower than planned bond price. To distinguish default uncertainty from liquidity risk, we need to attribute a loss of value between changing market default estimates and illiquidity. One approach is to decompose the fluctuations into movements in the mid market price, minus half the movements in bid-offer spread. The default uncertainty premium is compensation for the first of these, while the illiquidity premium is compensation for the latter. There is a large body of theory on the first element, relating the reward for risk to levels of risk aversion within the market. Market risk, measured in terms of sensitivities to market prices, combines linearly across portfolios. Therefore, any equilibrium model (of which CAPM is an example) must imply a linear relationship between market risk and the risk premium, with the constant of proportionality conventionally denoted as the "market price of risk". An assumed market price of risk allows us to relate the premium for default uncertainty as a multiple of the statistical uncertainty itself..

3.2.13 We are unable to characterise the illiquidity premium in such an elegant way, because we are unable to find a numerical characterisation of illiquidity that is additive across portfolios. For example, suppose an insurer has 50% of "liquid" liabilities that policyholders can encash at short notice, and 50% "illiquid" liabilities with no possibility of surrender. Other things being equal, this insurer is indifferent to anything between 0% and 50% of its assets in illiquid securities.

3.2.14 The non-linearity of the effect of liquidity creates subtle clientele effects. The natural clientele for illiquid assets is the investor who is unlikely to have to sell, or at least unlikely to have to sell in stressed market conditions. Most of these investors see an overall benefit, in that the liquidity premium they receive more than compensates them for the contingent transaction costs of stressed disposal. Such investors rationally increase their holdings in illiquid assets until, at the margin, the two effects equalise, but, thanks to the non-linear combination of illiquidity within a portfolio there remains a valuable margin in the existing asset stock relative to expected disposal costs. While reallocating market risk between entities, according to the arbitrage arguments of Modigliani and Miller, creates no overall value, it is possible for an overall economic benefit to arise from appropriate location of illiquid assets. This gain may appear in the form of a free lunch for an entity with low liquidity requirements who switches from cash into illiquid investments.

3.2.15 The separation of default uncertainty premiums from illiquidity premiums applies in principle to any asset class, but is particularly visible for dated instruments such as bonds where investors can avoid incurring sale costs simply by holding to maturity.

3.2.16 Corporate bonds require more vigilant administration on the part of investors, compared to government bonds. A number of third party vendors offer to investors systems to monitor credit risk in a bond portfolio. These systems may also involve some form of portfolio optimisation – although investors may be suspicious in some cases that the optimisation is not in regard to investor wealth but spreads generated by trading. These systems may be expensive to implement, both in terms of staff time but also license fees for software and data feeds, which may include accounting information in relation to bond issuers as well as market prices. Investors often put in place procedures to maintain portfolio credit risk within specified bands. Trading to maintain those bands typically means transaction volumes or frequencies are higher than for a government bond portfolio, and also at wider percentage spreads. Corporate bond investors may incur legal fees to maximise their recovery from any defaulting bonds. The inhomogeneity of corporate bond markets, the dispersal of basic information such as coupons, terms and prices, and the relevance of issuer balance sheet data, means that data management and audit costs for a corporate bond portfolio can be many times higher than for government bonds, while market making is also more difficult. A further reason for different expenses may be asymmetric tax treatment between government and corporate bonds, although this is not significant in the UK. These components together are often called “convenience” (Smith, 2000). An example of this applies in reverse in relation to cash. Individuals hold a certain quantity of banknotes, without interest, because of the convenience of using cash for transactions.

3.2.17 There may be a further element of bond premium attributable to inconvenience. To the extent that inconvenience costs do not scale with bond holdings, there may be convenience clientele effects, with investors over given critical mass enjoying a free lunch from a convenience yield that exceeds their management costs.

3.2.18 There are, therefore, at least three reasons why corporate bonds are potentially less attractive than government bonds. Firstly, corporate bonds are more likely to default; secondly corporate bonds are less liquid especially during crunches, and thirdly there are greater management costs. In a market, these elements result in lower prices and higher yields for corporate bonds. The prices stop at a point where the higher yield on the corporate bond adequately compensates investors for the default, illiquidity and inconvenience.

3.2.19 It would be desirable to decompose spreads into four components: credit risk, liquidity, convenience and an unexplained element. The credit component may further be subdivided into the expected losses due to default plus the higher return demanded by investors for the non-diversifiable element of default risk.

3.2.20 It is possible to estimate the default risk element, using a model such as that developed by Merton (1973). This has been applied by Bank of England (2007). Their model constructs a theoretical bond price using Merton's model and assumptions about firms' asset volatility. Their calibration happens to produce prices higher than market prices, with a significant unexplained discount. As the Merton model makes no allowance for liquidity premiums, the authors suppose that the unexplained discount contains the premium for illiquidity. However, their approach has two conceptual difficulties. Firstly, there is no allowance for illiquidity effects in equity markets, only for the bonds. To the extend that bond illiquidity premiums are consistent with equity illiquidity premiums, this allowance for liquidity should appear as part of the assumed equity premium, and is therefore already counted in the explained part of the price. Secondly, the unexplained proportion may contain some illiquidity effects (to the extent this apply disproportionately to bonds rather than equities) but it also contains any extent to which Merton model (or rather, tha Bank of England's calibration) is inappropriate for bonds. For example, Merton's assumption of normal distributions may not be appropriate. It is impossible to separate the unexplained portion of bond prices between liquidity effects and the extent to which the Merton model is just not a good fit to the data. The bank of England report recognises this, and at no point does their report claim that the unexplained portion consists substantially of liquidity effects. However, there is a risk in adopting these Bank of England calculations uncritically as a recipe separating default uncertainty premiums from liquidity premiums.

3.2.21 In principle, we may be able to calculate independently the appropriate compensation for liquidity risk should equate to the value of the bid-offer spread, multiplied by the number of times a typical bond holder switches holdings. This requires assumptions with regard to transaction frequency and the likely spread on the date of future transactions – bearing in mind that these may be in the middle of financial turmoil, or may be in more normal conditions. It also requires a characterisation of a "typical" investor in illiquid assets.

3.2.22 The convenience element is also potentially quantifiable. This is important in an insurance context because of the need to avoid double counting. For example, an insurer who invests in corporate bond probably holds higher expense provision than a government bond investor. If we are to insist on common discount rates for the two insurers, we should logically also equalise the approach to expenses. Allocating the expense differences is not straightforward, since, apart from tax, many of the expenses are fixed rather than proportional to the size of a bond holding. Conversion to percentage terms requires not only accurate expense information but also information on typical portfolio sizes..

3.2.23 We might hope to decompose the spreads using a multivariate model. For each bond, we develop estimates of the credit risk, liquidity and convenience. A least squares regression, with each bond constituting one data point, should in principle reveal the desired split. Unfortunately, this approach is doomed to failure. The reason is that both liquidity and convenience are related to default risk. The liquidity effect is dominated by asymmetries in relation to defaults in a crunch scenario. The inconvenience costs relate to processes for managing the risk of default. Without default risk, the liquidity and convenience effects are insignificant. This means that the three components: default, liquidity and convenience are statistically confounded and no statistical analysis can separate them. For example, we cannot extrapolate to a hypothetical bond with no default risk but positive illiquidity premium, because extrapolating the default risk to zero also collapses most of the illiquidity effects.

3.3 Inter-bank deposit markets

3.3.1 Banks routinely lend money to each other. This is usually in the form of time deposits, that is, cash deposits for a fixed time at a fixed rate. The depositor (or lender) provides cash to the recipient (or borrower). At the end of the term, the cash is returned with a previously agreed rate of interest, unless the contract is rolled over into a new contract.

3.3.2 As in any market, there are bid and offer prices. However, conventions for inter-bank deposits operate the other way around from conventional bonds. Here, a “bid” means a rate posted by banks looking to receive cash deposits. An “offer” means a rate asked by a bank with cash to invest. The bid rate is therefore lower than the offer rate, but if you convert the rate into an equivalent bond price then the offer price is lower than the bid price!

3.3.3 Like government and corporate bonds, bank deposits are subject to the risk of borrower default. The operation of the market is complicated by differences in credit standing between banks. A bank accepting deposits is indifferent to the credit status of the lender. However, riskier banks typically have to offer higher rates to attract deposits.

3.3.4 Lenders’ well-founded concern to operate credit limits and other underwriting criteria, prevents the operation of any exchange clearing system for inter-bank deposits as exists, for example, in the purchase and sale of traded bonds. As a result, there is no secondary market in bank deposits.

3.3.5 The lack of exchange trading complicates data collection. To collect data on the going rates for inter-bank deposits, we cannot simply download data from an exchange. Instead, some form of survey of banks is required.

3.3.6 One example of such as survey is the London Inter-bank Offer Rate. A panel of well rated banks provide information on rates for various currencies, and the British Bankers’ Association collates the data. The published rate for a given maturity is the arithmetic average of offer rates taken from each member of the panel, excluding the highest and lowest quotes. Other bodies compete with the BBA to be the benchmark rate, for example the European Banking Federation collates data for EURIBOR. These calculated rates are usually available free of charge from the web site of the compiling organisation and also online from systems such as Datastream, Bloomberg or Reuters.

3.3.7 There is a trade-off in the composition of the panel, between having an exclusive panel of only the best-rated banks, or broader market coverage while compromising on participants' credit quality. Nevertheless, from time to time there is concern over the panel's composition, and over whether data submitted to the panel is truly representative of actual trades. There may be no trades in a particular maturity for which a bank submits data, introducing a subjective element to the calculation. There may also be more serious issues of selective reporting. For example, a bank that is desperate for cash may be prepared to pay a high rate for funds. At the same time, if the market became aware of those high rates, this may signal the cash shortage, triggering a run on the bank. To avoid this difficulty, it is sometimes suggested that banks may be tempted to file false rates to the LIBOR calculation – a scurrilous suggestion which the data collators are naturally quick to deny.

3.3.8 One important variation on the standard contract is collateralisation. The lender is at a risk of borrower default, but that risk is mitigated if the borrower simultaneously passes assets to the control of the lender. These assets are called "collateral". This is the interbank equivalent to a secured corporate loan.

3.3.9 At first sight, we might expect cash to be useful as collateral. However, the purpose of a deposit is defeated if the borrower must return immediately to the lender any cash borrowed. Instead, the borrower uses other assets, usually traded investments. The transaction is then equivalent to an initial sale of the borrower's assets in exchange for cash, followed by a repurchase at the end of the deposit period, the purchase price being fixed in advance as the original sale price plus interest on the deposit. For this reason, collateralised deposits are also called "repurchase agreements", or "repos" and the rate paid on the deposit is a "repo rate". There is a similar market in "securities lending", although the securities lending market lacks the price transparency of the repo market, and a wider range of collateral may be accepted.

3.3.10 The use of collateral also provides another possible motivation for entering a repo transaction. The lender may wish to establish a short position in a certain asset, for example in order to benefit from falls in the asset price, or in order to offset a long position from another transaction. In that case, the lender seeks the assets he or she wishes to short, in the form of collateral, and then buys them for cash under a repo agreement. The lender immediately sells the collateral in the market. This is possible because the repo agreement does not oblige the lender to retain collateral throughout the deposit period, merely to return it at the end. Finally, immediately prior to the repo expiry, the lender buys back the collateral in the market and re-sells to the borrower at the previously agreed price, plus interest. In this way, the lender profits from any price fall in the collateral during the term of the trade, while losing money in the event that the collateral price rises in the market.

3.3.11 Central banks also participate in repo markets, but only as lenders, not as borrowers. This is the primary way by which central banks provide liquidity during crunches, rather than by unsecured lending as is commonly thought. The central bank defines the concept of "general collateral", which is a list of permitted assts or classes of assets against which the bank will lend money in repo transactions. Controversial central bank actions in 2007-8 to "provide additional liquidity" involved a broadening of the general collateral definition, including to various asset backed securities and structured credit instruments. While central banks are considered the "lender of last resort", some (for example, Buiter, 2006) argue that this extends central banks' role to become "market maker of last resort”. Stigma may attach to a bank who takes advantage of central bank repo transactions.

3.3.12 Default risk in repo transactions is limited. The lender worries that a borrower may not repay the deposit, but in that case the lender is entitled to retain the collateral. Borrower default, therefore, is only a concern to the extent that the deposit plus interest exceeds the market bid price of the collateral. This represents an effective mitigation of credit risk relative to unsecured deposits, and rates of interest on repo transactions are correspondingly lower than unsecured deposit rates.

3.3.13 Paradoxically, under a repo transaction, the borrower is also exposed to a form of default risk from the lender. The risk is that the lender refuses (or for some reason is unable) to return the collateral in exchange for the deposit. This represents an economic loss to the borrower if the deposit amount plus interest is below the market offer price of the collateral. As the default risk works in both directions, any credit loading in repo rates is affected by the credit standing of both parties, and the correlation between credit events and price movements in the collateral. Credit risk is exacerbated, in both directions, if the collateral price is volatile.

3.3.14 As for unsecured deposits, repos are bilateral trades, often between banks but insurers, pension funds and some local authorities are also active in this area. In this case, there are two reasons that make exchange difficult: firstly, a lender’s desire to underwrite credit risk, and secondly to agree the collateral. The first can be partially overcome by over-collateralisation, and the second can be overcome for assets so widely pledged as collateral that a separate contract is worth setting up just for that particular collateral. It turns out that this is only worth doing (in the UK) for government bonds, that is, gilts. There is a government-backed gilt repo market operated by gilt edged market markers. Otherwise, for bilateral trades, data collection is similar to the situation for unsecured deposits; third parties form a panel of banks, secure enough to be regarded as good credit risk but inclusive enough for reasonable market coverage. Separately, central banks publish their own, usually higher rates. The rates are higher because the central bank must be open to all qualifying applicants, must accept any collateral on its list, and also because the central bank has no self-serving desire to transact; it is there only to support other market players. As an instrument of bank policy, central bank repo rates are typically expressed in multiples of 0.25% and jump discretely between different levels, behaviour very different to market rates that can take any value and fluctuate from day to day.

3.3.15 Differences in collateral complicate the data collection considerably. Although the central bank dictates a list of general collateral, in practice choice of collateral affects the repo rate paid. Lenders are more reluctant to lend against illiquid or volatile collateral, and so borrowers pay higher rates for repos on such capital. At the other extreme, there may be some assets subject to a temporary spike in demand. This might for example happen because a bond attains benchmark status as a result of moving into a particular maturity band, or because an asset is due for physical delivery in large quantity under the terms of a delivery contract. Where the squeeze is temporary, market participants may choose to acquire the asset under terms of a repo contract, and indeed may offer to lend at low, even negative, repo rates in order to acquire the asset. The borrower, often an investment manager, retains the desired exposure to the underlying asset, while also gaining the spread between the repo rate and deposit rates available elsewhere.

3.3.16 To collect repo rate data, the collation agency must therefore specify the collateral. The market standard, defining the "general collateral repo rate" or "GC repo" is to take the *highest* rate for any collateral on the central bank's general collateral list. Other collateral trading at a lower repo rate, is called "special on repo" and is ignored in the market GC repo rate calculation. Rates on repo transactions for other collateral, not on the central bank list, are also disregarded. The published rate is then the arithmetic average of the bid-side rates for the panel banks. Incidentally, "special" does not mean "exclusive"; on some days, special transactions may account for 90% of all trades.

3.3.17 The notion of liquidity does not extend easily from traded bonds to bilateral transactions. Either side may request the other to unwind a transaction mid term, but the terms of any such unwinding are subject to negotiation, likely unfavourable to the party seeking to unwind. Therefore, while making an inter-bank deposit is akin to buying a bond from that bank, in the presence of a secondary bond market, the bond purchase is the more liquid transaction. We therefore expect bank deposits to offer higher rates of interest than a traded bond of similar term and credit risk. However, liquidity is not necessarily a major concern as the typical transaction is of short duration, sometimes only a few days. In the case of overnight deposits, these are more liquid that bond transactions, because the investor gets the cash the next day anyway, just as soon as settlement from a bond sale.

3.3.18 Differences in data collection complicate the analysis of spreads between rates. LIBOR is an offer side mean of rates while GC REPO is a bid-side maximum rate, averaged between contributors. In addition, the panel of banks for the two data exercises may not be the same. So to explain differences between LIBOR and REPO, we must look to differences in the underlying credit risk, the effect of bid-offer spreads and the effect of taking a maximum rather than an average.

3.4 Swaps markets

3.4.1 As interbank deposits are usually of terms less than a year, these arrangements are frequently rolled over. That means that a borrower requests a new deposit from the original lender, or another lender, in order to repay the first deposit. Similar dynamics apply to bank overdrafts, where rollover is usually automatic. Rates are changed from time to time but the interest just rolls up.

3.4.2 An individual or business financed by short term borrowing therefore faces deteriorating cash flows should interest rates rise. To manage this risk, the business might take out an *interest rate swap*. The swap has two parties; the *payer* pays a series of fixed cash flows, and in return receives a series of variable cash flows linked to a published index such a LIBOR. At the same time, the *receiver* pays variable cash flows and receives fixed flows. The fixed flows are usually expressed as a percentage of a *notional* amount, and the variable flows are calculated using the LIBOR rate applied to the same notional amount.

3.4.3 Settlement is usually based on the difference between the fixed and variable rates, so that the actual cash flows under the contract are often a fraction of one percent of the stated notional, unless interest rate moves have been particularly large during the transaction's life. This serves to mitigate the risk of default in swap transaction. Credit risk is further reduced by the use of margining. For example, if interest rates increase this is to the benefit of the payer, who will be concerned about the receiver's ability to continue paying under the swap transaction. In that case, the receiver may, by prior arrangement, deposit cash to the payer to provide security against the present value of future payments. Likewise, the payer deposits cash with the receiver in the event of an interest rate fall. Some credit risk, however, remains, as one or other party may default on a margin payment, and such risks grow at times of systemic instability just when default protection should be especially important.

3.4.4 There are other kinds of swaps, but when we refer to "swaps" without further description, we mean these conventional interest rate swaps whose floating legs are defined by reference to 3 month LIBOR or 6 month LIBOR.

3.4.5 Swaps, like inter-bank deposits, are bilateral transactions that are not traded on exchanges. Unwinding a swap technically requires a negotiated agreement between two parties. However, there are other ways of achieving the same effect. For example, suppose a borrower enters a payer swap, and half way through the term, the interest rate hedge in no longer necessary. If the former swap counterparty is unwilling to exit the transaction, then the borrower can go to a third party and execute a receiver swap maturing on the same date as the original swap. In net terms, the variable cash flows cancel out. What remains are a series of fixed cash flows equal to any change in the market swap rates between the two transactions, less the effect of any bid-offer spread on the swap transactions.

3.4.6 A borrower would naturally be a user of *payer* swaps. This is because the borrower has a variable rate interest liability. By entering a payer swap, the borrower pays a fixed rate to the swap counterparty, and receives a variable rate, ready to pay the variable rate on the overdraft.

3.4.7 In the same way, an investor may be a user of receiver swaps. Instead of buying a bond, the investor could roll over cash deposits at a bank and then take out a receiver swap to convert those variable payments into fixed coupons. One possible reason for doing this would be a desire to review at regular intervals the credit-worthiness of the bank receiving the deposits.

3.4.8 In both of these transactions, some *basis risk* remains. Basis risk arises when the interest rates on a swap do not correspond exactly to the interest rates on an underlying transaction that the swap is supposed to hedge. The operation of derivative markets is undermined if one of the parties can manipulate markets to change the promised cash flows. For cash settled derivatives such as swaps, the use of prices observed from an exchange or collated by an independent third party, reduce the scope for manipulation. The disadvantage is that the LIBOR rate underlying the swap is an average of market rates, with an inevitable time delay when the average is calculated. A borrower will not pay exactly LIBOR but rather the rate agreed on a bilateral basis with one lender. Similarly, an investor rolling over deposits does not earn exactly LIBOR.

3.4.9 Quite apart from market timing and averaging, the borrower faces further basis risks arising from credit risk. If the borrower's credit risk deteriorates, the bank may be unwilling to roll over an overdraft, or may do so only on worse terms. While the payer swap still pays LIBOR, the borrower experiences an additional variable cost linked to credit standing. The swap hedges market risk but not credit risk.

3.4.10 The investor faces a different problem. Averaging is one issue; bilateral negotiation may not produce the average level, but on average it does! The problem rather is of the bid- offer spread on deposits. An investor can set up his stall offering to lend money at a certain rate, but there is no guarantee someone else will take up the offer. To be sure of investing the money, the investor must take up someone else's bid rates. We might imagine a parallel universe in which interest rate swaps were settled relative to bid rates, offering a better hedge to investors (and a worse one to borrowers). Historically, the earliest users of swaps were borrowers, who were most interested in LIBOR and we are now stuck with this convention. As a result, the investor seeking to replicate a bond payoff worries not only about the averaging underlying LIBOR but also fluctuations in the bid-offer spread. This basis risk does not average out over time; even in benign market conditions the offer rate exceeds the bid rate. During a credit crunch, investors may receive a great deal less on the bid-rate deposits compared to the LIBOR due under a receiver swap.

3.4.11 Analysis of swaps is complex, because of their derivative nature. Our analysis of bonds suggests we should look for effects related to credit, liquidity and convenience. In the case of a derivative transaction, these measures are relevant both for the underlying rate and also for the derivatives themselves. It turns out that swaps of most terms are very liquid, with spreads much tighter than on government bonds; the credit risk is small because the netted cash flows are small compared to the notional amount, and the use of ISDA standard contracts makes these transactions convenient. So we might imagine that swap rates are as close as we could get to risk free. However, this is to overlook the much wider bid-offer spreads, the illiquidity and inconvenience of the deposits underlying the LIBOR calculation.

3.4.12 We can compare a corporate bond to a synthetic bond made from cash deposits plus a receiver swap. Both are subject to some credit risk – in the case of the synthetic bond, the main credit risk is the underlying deposit with a negligible contribution from the swap. Superficially, these are both vulnerable to a credit crunch. However, if the crunch blows over before the bond maturity, then the investor has lost nothing, while the synthetic investor has suffered directly from basis risk as widening bid-offer spreads hit coupon income. If a bank gets into difficulties, the synthetic investor may be able to move switch deposit to another bank; a "refresh" option not available to direct bond investors. On the other hand, if an investor does need to realise assets at an inopportune moment, this is likely more painful for the corporate bond holder than the synthetic bond.

3.4.13 Banks contributing to LIBOR panels are typically of AA or better credit standing. It is therefore sensible to compare the coupon available on synthetic bonds to bonds of the same term from AA issuers. The contrast is stark – the corporate bond has the higher yield, and this has been the case for most of the history of the swaps market (in the early 1990's the yields were closer). This suggests that the refresh feature is valuable. The direct bond investor is exposed to AA risk, but that may turn into AAA or B during the term of the bond. The synthetic investor is exposed to AA, regularly refreshed. Give the high quality AA starting point, the drift for the direct investor is more likely to be down than up, hence the advantage of refreshing. This is more likely a credit effect than a liquidity effect, because (as argued in the previous paragraph) liquidity applies differently for the two transactions, with no particular advantage to either over the other.

3.4.14 One increasing trend is the growing popularity of swaps based on overnight rates rather than on 3 or 6 month LIBOR. In sterling, these are called SONIA. Within the euro-zone there are two versions: EONIA and EURONIA, compiled by different agencies and slightly different for reasons of timing and panel compoision. These have the advantage of trading more frequently, so that published SONIA is an average of actual trades, while LIBOR is an estimate. On the other hand, SONIA swaps are administratively more tiresome because of the need to keep track of daily changing rates. And one surprising result is that the swap rates on SONIA or EURONIA are systematically lower than the corresponding swaps on LIBOR (respectively, EURIBOR), adjusting of course for the different compounding conventions. At first sight, this is puzzling. If both swaps are based on unsecured inter-bank deposits, why this spread between the corresponding swap rates?

3.4.15 The answer lies in the "refreshing effect". Consider the positions of two synthetic bond investors, one rolling over semi-annual deposits with a LIBOR-based receiver swap, and the other rolling overnight rates with a SONIA swap. Both are exposed to the risk of bank default on their deposit. However, bank failures are seldom overnight affairs. A bank's difficulties typically feature in the papers for several weeks before the bank collapses. The synthetic bond holder based on term LIBOR is still exposed to the bank failure, unless he is fortunate enough to have one of the deposits mature during the window between bad news in the papers and the bank's collapse. The overnight synthetic bond holder, however, can switch deposits on a daily basis without penalty. An exposure remains to the bank that fails without warning, but the usual lag of a few weeks is more than enough time for the SONIA synthetic bond investor to switch elsewhere. Therefore, the LIBOR synthetic bond is exposed to greater credit risk. We expect it to pay a higher coupon, and indeed it does. We have already seen this effect in comparing quarterly or semi-annually refreshed synthetic bonds to unrefreshed bank debt. It is reassuring that the same effect continues, with yields falling further for more frequently refreshed trades.

3.4.16 In the context of conventional bonds, we looked into ways of constructing rates free of default risk. Here, the technically difficult issue was variation between bonds in liquidity and convenience effects. The conclusion was that if you extrapolate to freedom from default risk, you also eliminate spreads for liquidity and inconvenience, so you get back to yields on government bonds. Our examination of repo trades allows us to close a subtle loop in the argument – if risk free bonds cannot have liquidity spreads, why are there differences even between government bonds of different liquidity? The answer is that it is only the crunch element of liquidity risk that reduces with default risk. The bid-offer and market impact elements remain, even in gilt markets. We can explain those differences by looking to repo markets. It so happens that the most liquid gilts are more likely to go special on repo. The bearer of such gilts can therefore earn additional spreads from selling those gilts in a repo market, and reinvesting the cash on deposit at higher spreads. According to Duffie (1996), the present value of those spreads equates to the yield difference between more or less liquid bonds.

So you can look at benchmark liquid bonds and add some special repo spreads, or you can take reference yields from less liquid bonds. The special repo rate is another example of a market transaction whose price may give a guide to liquidity premiums. Either is a reasonable candidate for a “risk free” rate, and thankfully, the numbers are the same anyway.

3.4.17 We can ask the same question in relation to inter-bank transactions. LIBOR transactions are clearly not risk-free, but repo transactions come closer because of the effect of collateral, at least if we focus on times (unlike now) when central banks set demanding requirements for general collateral. The problem with both LIBOR and GC REPO rates are the short term of typical transactions. In the case of LIBOR, we have the swap curve as a means of extending to longer term. Sadly, swap markets based on GC REPO rates are not common, and current prices are difficult to find. However, SONIA / EURONIA swaps do exist, trading at a discount to LIBOR-based swaps, with that discount attributable at least in part to credit risk, because of the refreshing effect. Furthermore, historically the spread between REPO and overnight rates has been small and stable, even during credit crunches. Therefore, a possible approach to a risk-free rate, based on interbank transactions, is to use SONIA / EURONIA swaps, adjusted downwards for the small historic average spread between overnight unsecured rates and repo rates. This produces rates reassuringly close to government bond yields.

3.4.18 There are two kinds of bank failure. The date of failure can be previsible, that is, you know about the failure an instant before it happens, or it can be inaccessible, in which case it comes completely out of the blue. Structural models of default (such as Merton’s model) are often driven by a Brownian motion applied to a firm's assets, with default triggered when assets fall below a certain barrier level. The date of ruin is previsible because the asset paths are continuous; you can see the collision coming up. On the other hand, reduced form models (such as Jarrow-Lando-Turnbull) behave more as actuarial decrement models, in that the default probability varies slowly with time (for a given credit grade) and default comes out of the blue. If you take the limit of continuous LIBOR (ie beyond SONIA swaps) then the refreshing eliminates the previsible component of failure, while the risk of inaccessible failure remains. On the other hand, the collateral on a repo trade protects against loss both from previsible and inaccessible failures. Therefore, differences between overnight repo and SONIA should provide some insight into the split of bank credit risk between the previsible and inaccessible components.

3.5 Liquidity Transactions

3.5.1 We have considered liquidity issues for various forms of inter-bank lending. Most businesses face more substantial liquidity issues on the asset side. Businesses often invest in long term transactions – for an industrial company this might be plant or inventory, for a bank these are various sorts of loans. Usually, these transactions are impossible, or at least very expensive, to unwind. That is to say, the assets are illiquid. Furthermore, the value of the assets is tied to their owner; a loan is valuable in the hands of a bank that can collect interest payments and pursue unpaid debts; plant and inventory are useful to a manufacturing business. There is significant deadweight economic cost if these assets are separated from their originators, for example in a liquidation. For this reason, businesses are keen to secure continuity of funding, and to lock in the cost of any such funding to maximise the value of their assets. One way of doing this is to issue long term debt, but floating rate debt, rolled over quarterly, is generally seen as cheaper because lenders give credit for the option not to renew the debt. If the borrower is secure in their credit rating, they may consider using swaps to hedge their cost of borrowing, but they still face basis risk to the extent that LIBOR does not represent typical trades, or that the borrower is unable to borrow on normal terms, for example because of unusual transaction size. In other words, some liquidity risk remains even when the borrowing cost is hedged.

3.5.2 To remove these final elements of basis risk, a company might consider a sale and asset swap transaction. An insurer buys the illiquid asset, but then enters an asset swap where the return on that illiquid asset (measured on a suitable mark-to-model basis) is swapped with the originator for LIBOR plus a spread. The positive spread rewards the insurer for taking an illiquid asset onto its balance sheet, and the seller is prepared to pay the spread in return for security of funding. The asset originator is effectively renting liquidity from the insurer. Thus, the prices of such transactions may be regarded as an indication of a market price of illiquidity risk.

3.5.3 A crucial feature of this transaction is that, although the illiquid asset technically sits on the insurer's balance sheet, the originator retains an economic interest and can continue to add value to the asset by its own management. There may also be accounting benefits for the bank who accounts the asset at its mark-to-model (via the asset swap valuation), rather than the (lower) bid price for which the asset might be sold to a third party. The insurer, with matched assets and liabilities also avoids recognising the bid-offer spread.

3.5.4 The remaining loose end is the investment strategy of the insurer. The insurer technically owns an illiquid asset, but the economic effect is swapped back to the originator, leaving an asset earning LIBOR + a spread. Such a LIBOR asset may not be desirable within the insurer's own investment strategy. To overcome this problem, the insurer may obtain the cash to buy the illiquid asset by repo trades on other assets they already hold. In this way, the insurer retains exposure to the original asset, while receiving a stream of LIBOR – GC REPO + spread in return for the liquidity service

3.5.5 Further research is required into these transactions, to determine whether their prices can provide a reliable indication of the value of liquidity. There are many aspects affecting these transactions, and with so many unknowns, there inevitable guesswork in defining precisely which risks the spreads may be compensating. For example, if the transaction cannot be unwound and involves in effect a long term collateralised loan to a bank, then the spreads should contain some allowance for the joint default of both the bank and the collateral.

3.5.6 As yet, no widely used derivative pricing model refers explicitly to illiquidity premia; any account taken is implicit in the rates used as inputs, or in any discretionary spread a trader adds to a particular quote. We might be tempted to think that illiquidity premiums are like risk premiums, in the sense that they cancel out in the hedge construction and therefore do not enter option prices. This impression is reinforced by the absence, currently, of trades that price only liquidity. The lack of trading removes a mechanism that might otherwise drive consensus on the value of liquidity. For example, two investors in a corporate bond yielding 300bp over gilts may have different views as to how that premium splits between compensation for beta risk and illiquidity risk.

3.5.7 Despite limited success to date, numerical separation of beta from liquidity effects may be possible in principle. We have already noted sale-and-swap transactions, whose prices contain an identifiable liquidity element Unfortunately, the illiquidity in question relates to a loan or to industrial plant; even if we could observe market prices for such transactions, there are too many unknowns for us to deduce the illiquidity premium in bonds.

3.5.8 Special repo transactions provide a more promising source of liquidity premium information on gilts. At a simpler level, asset securitisations, property investment trusts and exchange traded funds, provide a way of repackaging illiquid assets to create more liquid assets. Such transactions could therefore be considered as acting as a prism, separating the price of market risk from the illiquidity element. If these markets develop further, they could, in principle, eventually lead to an observable market price of illiquidity risk.

3.6 Other Derivative Markets

3.6.1 Virtually all derivative trades are affected by the time value of money. Therefore, there is a need for a reference curve to price any derivative transaction. A common choice of reference curve is based on swaps relative to 3-month or 6-month LIBOR. In this section, we discuss various derivatives and consider evidence in relation to the underlying curve used.

3.6.2 Let us first consider debt issued by banks. Does this trade at swap yields? Absolutely not; bank debt trades at AA yields, according to the bank's credit rating. This is higher than the yields on swaps of comparable term, because of the way swaps refresh. Technically, these bonds do not count as derivatives. However, if we are to support the notion of a common reference curve for all derivatives, this implies there must be something special about derivative pricing which does not apply to corporate debt.

3.6.3 Unfortunately, even some derivatives fail the test. For example, EURONIA swaps fail. If a single reference curve could explain swap rates of any refreshing frequency, than the swap rates would not depend on refreshing frequency at all – which is clearly contrary to observed market prices.

3.6.4 Another interesting test is to look at currency forward prices. In theory, there is an interest rate parity argument, where the difference between forward and spot prices for a currency pair can be re-expressed in terms of differences in spot rates. Empirically this works well. As any finance text book explains, if this relation did not work, there would be a clear arbitrage from buying discount bonds in one currency, selling in another and making up the difference with a forward trade.

3.6.5 Forward parity works fine, that is, until one of the currencies is the Japanese Yen and the time is the late 1990's. Yen LIBOR was consistently 10-20bps higher than it ought to be in order to satisfy forward parity. This is the famous "Yen premium", which also applies to a number of emerging market currencies.

3.6.6 Why then, did arbitrageurs not borrow in euros, deposit in yen and close the arbitrage by selling Yen forward? The answer is that, to earn Yen LIBOR you have to deposit with a typical member of the Yen LIBOR panel – and these are different banks from the panel contributing to EURIBOR. During the late 1990's, many of the Yen LIBOR panel were Japanese banks, the subject of rumours of undisclosed losses in loan portfolios – which increased the perceived credit risk and so the cost of funds for those banks. In other words, to execute the theoretical arbitrage, you have to take credit risk on Japanese banks of questionable strength - which, of course, means that it is not an arbitrage at all.

3.6.7 Some evidence in support of a swap-based reference curve, comes from the credit default swap market. A credit default swap provides for a series of fixed "insurance payments", in return for a put option to sell an underlying bond (or one selected from a list of bonds) for its face value on the event of default. The fixed payments can be interpreted as a spread between the risky bond and a reference rate. Comparing this spread to the yield on the corporate bond itself, we can reconstruct the reference rate, which usually comes out to be close to swaps. However, the analysis of credit default swaps is complex, and undoubtedly reflects the credit standing of the contract counterparties as well as the risk of the underlying bond. Analysis is further complicated by the embedded interest rate option, arising because the bond is redeemed at par, and by the option to select from a list the bond to be subject to option.

3.6.8 When it comes to more familiar domestic derivatives, the use of LIBOR and Swap reference curves is widespread. It is true that the prices of equity put and call options are routinely calculated using models based on the swap curve. However, this does not exclude that the same price could not also be derived from another curve, for example government bonds. The key to understanding this point is the effect of collateralisation. Writers of options, particularly long dated options, typically post collateral with the option buyer. The collateral is usually cash, on which the option buyer pays LIBOR. From the perspective of a swap reference curve, the option is priced off swaps and the LIBOR interest on the collateral is cost-neutral. From the perspective of a government bond reference curve, the option is priced off government bonds, but the premium comes in two parts: firstly the explicit option premium, and secondly the obligation of the option buyer to pay interest on collateral at a spread over the risk-free rate. Neither perspective is wrong; both can explain the option price. It is not misleading to suggest a swap curve is used to price the option. It is misleading to suggest that no other reference curve could imply the same option price.