Financial Stability Institute

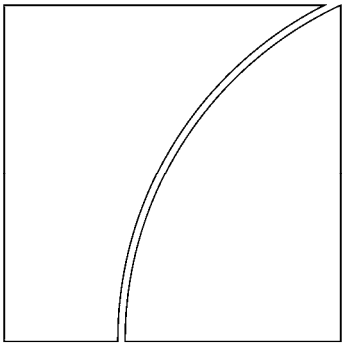

Occasional Paper

No 10

Liquidity transfer pricing: a guide to better practice

Joel Grant
Australian Prudential Regulation Authority

December 2011

BANK FOR INTERNATIONAL SETTLEMENTS

The views expressed in this paper are those of their author and not necessarily the views of the Financial Stability Institute, the Bank for International Settlements or the Australian Prudential Regulation Authority.

This publication is available on the BIS website (www.bis.org).

© *Financial Stability Institute 2011. Bank for International Settlements. All rights reserved. Brief excerpts may be reproduced or translated provided the source is cited.*

ISSN 1020-8461 (print)

ISSN 1020-9999 (online)

Contents

List of abbreviations

BCBS	Basel Committee on Banking Supervision
CEBS	Committee of European Banking Supervisors
CRMPG III	Counterparty Risk Management Policy Group III
EC	European Commission
FSB	Financial Stability Board
FTP	Funds Transfer Pricing
IIF	The Institute for International Finance
GFC	global financial crisis
LCR	Liquidity Coverage Ratio
LIBOR	London Interbank Offer Rate
LMIS	Liquidity Management Information Systems
LTP	Liquidity Transfer Pricing
NSFR	Net Stable Funding Ratio
SSG	Senior Supervisors Group
WGL	Working Group on Liquidity

Executive summary[1]

This paper identifies better practices for liquidity transfer pricing (LTP) by drawing on the responses to an international survey that covered 38 large banks from nine countries. The survey focused on the enhancements banks are making to their LTP processes.

Responses to the survey show that many LTP practices were largely deficient. Many banks lacked LTP policies, employed inconsistent LTP regimes, relied on off-line processes to manually update changes in funding costs, and had poor oversight of the LTP process. Probably the most striking example of poor practice was that some banks failed to attribute liquidity costs to assets and conversely liquidity credits to liabilities for some business activities. Others did attribute liquidity costs and benefits, albeit at one average rate. This approach failed to penalise longer-term funding commitments for assets and, conversely, reward longer-term funding benefits from liabilities, and failed to incorporate timely changes in banks' actual market cost of funds. Moreover, banks' liquidity cushions were too small to withstand prolonged market disruptions and were comprised of assets that were thought to be more liquid than they actually were. Overall, these shortcomings encouraged risky maturity transformation, without regard to the structural liquidity risk that was being generated.

---

[1] The author is grateful for comments received from APRA colleagues, in particular, John Laker, Charles Littrell, Katrina Ellis, Bruce Arnold, Neil Grummitt, Nick Palmer and members of APRA's Research Unit. I would also like to thank Jim Embersit from the Board of Governors of the Federal Reserve System, Kumar Tangri from the UK Financial Services Authority, and members of the Basel Committee on Banking Supervision's Working Group on Liquidity.

The author can be contacted at: joel.grant@apra.gov.au.

Better LTP practice requires each bank to produce and follow an LTP policy that defines the purpose of LTP and provides principles and/or rules to ensure LTP achieves its intended purpose. Banks should manage LTP centrally, such as in group treasury, with sufficient oversight provided by independent risk and financial control personnel. Treasury should have complete visibility of individual business balance sheets. To properly manage funding liquidity risk, banks should charge rates based on their marginal cost of funds and matched to the maturity of the product or business activity at origination. For amortising or non-maturing products, blended marginal rates should be applied. In regard to the sizing of liquidity cushions, banks should use the results of stress-testing and scenario analyses, which include idiosyncratic and market-wide disruptions, as well as a combination of the two. Assets held as part of banks' liquidity cushions should be of the highest quality to ensure liquidity can be generated when needed. Finally, business activities creating the need for banks to carry additional liquidity should be charged based on their expected usage of contingent liquidity.

Overall, better LTP practices will ensure that banks accrue less illiquid and correlated assets, use more stable sources of funding to meet the demands of their business activities, and carry a more sufficiently sized liquidity cushion to withstand unexpected idiosyncratic and/or market-wide disruptions. Banks, supervisors and other stakeholders are therefore encouraged to consider the better LTP practices that are identified in this paper.

1. Introduction

Internal transfer pricing is an extremely important management tool for banks. This paper observes that until the global financial crisis (GFC), many banks treated liquidity as a free good for transfer pricing purposes, and this was one cause for the very poor liquidity outcomes experienced during the GFC. Furthermore, although liquidity transfer pricing (LTP) practices are improving, there is little guidance publicly available to banks, regulators, and other stakeholders on what constitutes good practice. This paper makes a start on filling that gap.

LTP is a process that attributes the costs, benefits and risks of liquidity to respective business units within a bank.[2] LTP has gained considerable attention since the onset of the GFC with some reports linking poor LTP practices to the funding and liquidity issues witnessed at several banks (Senior Supervisors Group (SSG), 2008; 2009).

The purpose of LTP is to transfer liquidity costs and benefits from business units to a centrally managed pool. To achieve this, LTP charges users of funds (assets/loans) for the cost of liquidity, and credits providers of funds (liabilities/deposits) for the benefit of liquidity. LTP also recoups the cost of carrying a liquidity cushion by charging contingent commitments, such as lines of credit, based on their predicted (expected) use of liquidity. This is depicted in Figure 1 below. Banks with poor LTP practices typically under-price or (even worse) fail to price liquidity. Such banks are more likely to accrue illiquid assets and contingent exposures, and under-value stable sources of funding. This outcome applied to many banks and other financial institutions prior to the GFC.

[2] In this regard, LTP forms part of the funds transfer pricing (FTP) process.

Figure 1

A graphical representation of the LTP process

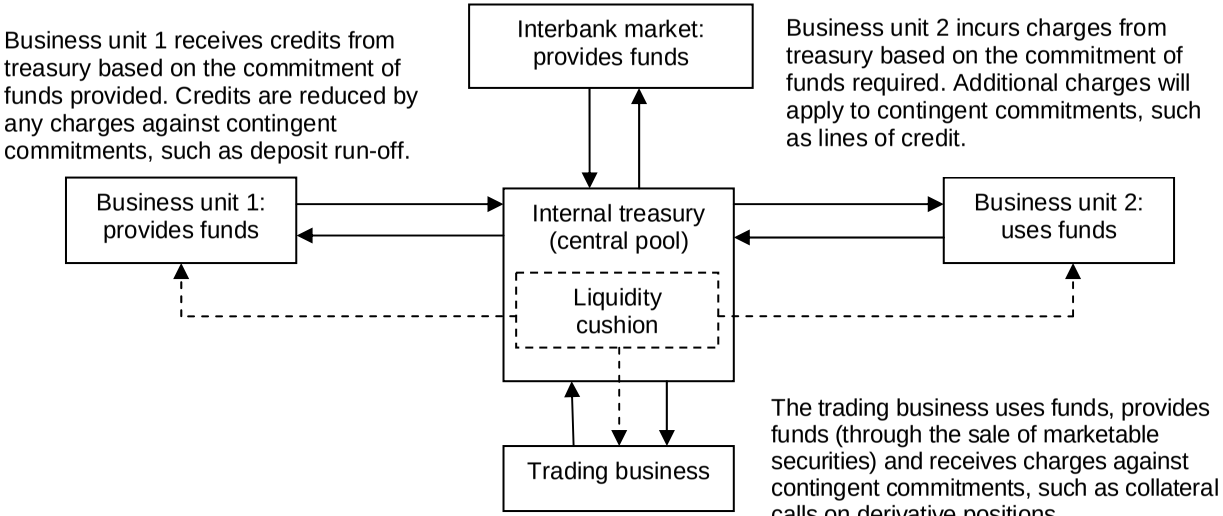

In the years preceding the GFC, liquidity was plentiful and cheap, and as we now know, unsustainably plentiful and cheap. Some of the larger and more creditworthy banks could obtain long-term funding at only the slightest margins above swap rates. Such ideal funding conditions proved fruitful for banks, widely encouraging leverage and maturity transformation, which underpinned their record profits. At the same time these conditions led many to believe that funding would always be available, and at permanently cheap rates. One consequence of this belief was that it provided little incentive for banks to devote attention to liquidity risk management. As a result, many banks failed to recognise the true nature of the liquidity risk embedded in their business activities.[3]

One principle of liquidity risk management that lacked attention was LTP. In 2009, a group of prudential regulators conducted an international survey to assess the progress banks are making to enhance LTP. The survey covered 38 banks from nine countries. Total assets of the banks surveyed ranged from less than US$250 billion to greater than US$1 trillion.

The survey responses revealed that many of the LTP practices employed by banks were short of good practice. This paper extracts the lessons learned from the survey, and makes a first attempt at establishing better LTP practice. For confidentiality reasons, however, it is not possible to quote or reference directly from any of the survey responses.

[3] This claim is supported by the Basel Committee on Banking Supervision (BCBS), which reported that many of the basic yet fundamental principles of liquidity risk management were neglected by banks. For more information, see *Liquidity Risk: Management and Supervisory Challenges*, BCBS, (February 2008).

1.1 A summary of the major lessons learned

1.1.1 Governance of the LTP process

Most banks included in the survey lacked an LTP policy. As such, LTP was not defined nor were there any rules or principles in regard to how LTP should operate. Typically, this outcome meant that liquidity generators (such as retail branches raising deposits) were underpaid for their liquidity creation, and liquidity users (such as lending, investment, and trading portfolios) received free or unduly cheap liquidity.

Where banks in the survey were operating with decentralised funding centres, most had inconsistent LTP regimes. In addition, these banks relied on manual off-line processes to intervene and to update relevant funding costs, and were more prone to arbitrage between business units and internal treasuries.

For many of the banks in the survey with large trading businesses, internal treasuries often lacked visibility over individual business balance sheets, limiting their understanding of individual funding requirements and contingent liquidity exposures. Most of the time this resulted in treasuries charging all trading businesses based on their net funding requirement, with no add-ons for the implicit risk of a blow-out in liquidity needs.

Oversight of the LTP process at nearly all banks that participated in the survey was poor to nonexistent, especially by risk and financial control functions. This was one of the factors that resulted in the accumulation of highly illiquid (and often correlated) assets and the excessive reliance upon short-term (often overnight) funding.

Liquidity Management Information Systems (LMIS) employed by most of the banks surveyed were simplistic and inflexible. Many of the systems were unable to attribute the costs, benefits, and risks of liquidity appropriately to respective businesses, and at a sufficiently granular level. This resulted in

product mispricing, which distorted profit and performance assessments.

For a large proportion of banks included in the survey, their LTP process failed to account for the costs, benefits and risks of liquidity in the pricing and performance assessment of various products and business units. As a result, profit measures used as a basis for determining business unit performance and executive remuneration were distorted. Profit pools, for example, which are generally used to determine short-term incentives (bonuses) for employees, were derived from a simple percentage of accrued revenues without any regard for the liquidity risk taken to generate such profits.[4] This encouraged revenue and risk maximisation rather than risk-adjusted earnings.

1.1.2 The application of LTP

Probably the most striking example of poor LTP practice was how some of the banks that were surveyed treated liquidity as a "free" good, completely ignoring the costs, benefits and risks of liquidity. These banks neglected to charge or credit respective businesses, products and/or transactions accordingly. This was particularly the case for much of the contingent or unfunded business that was written. Examples included trading and investment banking activities, lines of credit, the need to prepare for collateral calls, and variable-rate (adjustable-rate) products including home mortgages.

Most of the banks surveyed recognised the need to attribute the costs, benefits and risks of liquidity to respective businesses. However, a large majority of these banks employed a pooled average cost of funds approach to derive

[4] Bonus pools often neglected other risks, not just liquidity, and the cost of capital employed to generate such profits. This is the subject of another paper.

the costs and benefits of liquidity. This resulted in short- and long-term assets receiving the same charge for the cost of liquidity and, conversely, short- and long-term liabilities receiving the same credit for the benefit of liquidity.

1.1.3 Sizing and attributing the costs of liquidity cushions

For a large majority of the banks surveyed, liquidity cushions were derived from stress assumptions stemming mainly from idiosyncratic funding scenarios, revolving around a single bank's sudden inability to raise funds. Having little or no regard to systemic funding scenarios, most cushions were too small to withstand prolonged or deep market disruptions.

In addition, cushions comprised liquid assets that were themselves funded short-term. This meant that the cost of carrying the liquidity cushion was quite small, but the real value of the cushion in addressing sudden (contingent) liquidity risks was also minimal. This costing and funding arrangement provided insufficient incentive for banks to attribute true costs back to business units on an expected or predicted usage basis but, rather, to opt for the simpler but incorrect method of averaging the cost across all assets.

1.2 Regulatory developments

The Basel Committee on Banking Supervision (BCBS) has been central to regulatory developments in liquidity, first publishing *Sound Practices for Managing Liquidity in Banking Organisations* in February 2000. Following this, in 2006, the BCBS established the Working Group on Liquidity (WGL) to "serve as a forum for information exchange on national approaches to liquidity supervision and regulation". The group's initial mandate was to review and evaluate liquidity supervision practices, and banks' approaches to liquidity risk management, with respect to the sound practices already established.

This work was the first to highlight the basic yet fundamental elements that were missing from bank liquidity management. These findings formed the basis of the report *Liquidity Risk: Management and Supervisory Challenges* (February 2008) and sparked a review of the February 2000 sound practices. An updated version of these practices, articulating 17 principles, was released as *Principles on Sound Liquidity Risk Management and Supervision* (September 2008). Since 2008 the BCBS has released *Principles for Sound Stress Testing Practices and Supervision* (May 2009) and more recently, *Basel III: International Framework for Liquidity Risk Measurement, Standards and Monitoring*, aimed at improving the resilience of the financial system (December 2010).[5] As part of this, two global standards for liquidity risk were developed. First, a Liquidity Coverage Ratio (LCR) to ensure banks have sufficient high quality liquid assets to survive an idiosyncratic shock and, second, a Net Stable Funding Ratio (NSFR) to encourage banks to fund their business activities using more stable sources of funding.

1.3 The need for more guidance on LTP

The scale and extent of liquidity reform is large. But, given the weaknesses in bank liquidity risk management approaches unveiled by the recent crisis, it is not surprising that certain principles require further guidance. This is particularly the case for LTP.

Extant guidance is broad but merely encourages banks to include liquidity risk in their internal pricing mechanisms, without providing specific help.[6] For example, Principle 4 of the

---

[5] These papers can be accessed via http://www.bis.org/list/bcbs/sac_1/index.htm.

[6] A complete list of principles and/or recommendations provided by various regulatory and non-regulatory bodies is included in Appendix 1.

BCBS *Principles on Sound Liquidity Risk Management and Supervision* states that: "a bank should incorporate liquidity costs, benefits and risks in the internal pricing, performance measurement and new product approval process for all significant business activities (both on- and off-balance sheet), thereby aligning the risk-taking incentives of individual business units with the liquidity risk exposures their activities create for the bank as a whole" (p 3).

The Committee of European Banking Supervisors (CEBS) (now the European Banking Authority) has also highlighted the importance of banks having "an effective allocation mechanism for liquidity costs, benefits and risks". Recommendations provided in the *Second Part of CEBS's Technical Advice to the European Commission (EC) on Liquidity Risk Management* [7], and Point 14 in Annex V of the amendments to the Capital Requirements Directive (Directive 2009/111/EC of the European Parliament and of the Council of 16 September 2009)[8], led to the development of CEBS's, *Guidelines on Liquidity Cost Benefit Allocation* (October 2010).[9]

Efforts by the BCBS and others to improve LTP are acknowledged, but a lack of detailed supporting guidance has left some supervisors and banks asking: "what exactly constitutes better practice?" This paper seeks to assist with this question. It focuses on some of the more widespread poor

---

[7] See Recommendation 2 in *Second Part of CEBS's Technical Advice to the European Commission on Liquidity Risk Management*, September 2008, which is available at http://www.eba.europa.eu/getdoc/bcadd664-d06b-42bb-b6d5-67c8ff48d11d/20081809CEBS_2008_147_(Advice-on-liquidity_2nd-par.aspx.

[8] Available at http://eurlex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:302:0097:0119:EN:PDF.

[9] *Guidelines on Liquidity Cost Benefit Allocation*, CEBS, October 2010, can be accessed via http://www.eba.europa.eu/cebs/media/Publications/Standards%20and%20Guidelines/2010/Liquidity%20cost%20benefit%20allocation/Guidelines.pdf.

LTP practices that were identified via the international survey with the aim of drawing out better practices. This more detailed guidance is intended to support and supplement the principles already promulgated, and assist supervisors and banks in achieving better LTP practices.

2. Governing LTP

Broadly speaking, all policies, processes and practices require governing. This is normally achieved through a combination of external control factors, such as regulation and competition, and internal control factors, such as board oversight and risk management.[10] Because external control factors affect institutions in much the same way, governance is differentiated largely by the internal control factors that are employed.

How well an institution is governed can bear heavily on whether group-wide objectives are met. While institutions with strong internal controls are more likely to achieve their goals, institutions with weak internal controls are more prone to the problems of moral hazard and adverse selection.[11] This is likely to weigh on performance.

---

[10] The Committee of Sponsoring Organisations of the Treadway Commission (COSO) defines internal control as "a process, effected by an entity's board of directors, management and other personnel, designed to provide "reasonable assurance" regarding the achievement of objectives in the effectiveness and efficiency of operations, reliability of financial reporting, and compliance with applicable laws and regulations" (http://www.coso.org/resources.htm).

[11] Moral hazard occurs when a party, insulated from risk, behaves differently than they would if they were fully exposed to the risk (www.wikipedia.com). For example, a bank might be more inclined to engage in risky behaviour, knowing that it will be bailed out if the risks turn bad. Adverse selection, on the other hand, is a process whereby bad

Many of the poor LTP practices that were identified across banks that participated in the survey were the direct result of weak internal controls. Some of these are discussed in more detail below.

2.1 Management of the LTP process

Broadly speaking, there were severe deficiencies in how the LTP process was managed.

2.1.1 LTP policies

Few banks in the survey had an effective LTP policy. As a result, LTP was not defined, nor were there any principles and/or rules in place to assist businesses understand how LTP should operate. Having no LTP policy is clearly poor practice, given what we now know about liquidity risk. Banks have traditionally relied on internal transfer pricing to manage interest rate risk in the banking book, and to assess and monitor the performance of products and business units, but with no or only minimal adjustments for liquidity costs, benefits and risks. For most banks in the survey, the internal pricing of liquidity risk is a relatively new concept, brought to light by the recent breakdown in wholesale funding markets, and the consequent increase in funding costs. It will take time for banks to establish adequate LTP policies and procedures, but this is a necessary first step towards better LTP practice.

---

results occur because of information asymmetries between buyers and sellers. For example, a used car salesman might sell a car, which he knows has mechanical problems, to a buyer that is less informed.

2.1.2 Internal funding structure – centralised vs decentralised

There is substantial debate surrounding the optimal internal structure of banks – is it better to have a centralised funding centre, whereby wholesale funding is restricted to a group or subsidiary treasury or, alternatively, decentralised funding centres, whereby certain business units are able to raise funding themselves from their own sources to cover their own liquidity needs? There are reasonable academic and economic arguments that provide support for both approaches. However, the survey identified that banks with decentralised funding centres, particularly those with large prime brokerage business activities, were more susceptible to poor LTP practices. For example, some business units that were able to raise wholesale funds from external sources then sold the funds to treasury and in some cases to other business units, at a higher rate. This resulted in a "risk-free" profit to the business unit at the cost of more and possibly badly managed risk for the bank as a whole.

It is clear, however, that decentralised funding structures were not the sole cause of internal arbitrage. Poor oversight and inadequate risk controls also played a role. Jointly, these factors limited the ability of treasury and business units to know what price other business units had paid for funds from external sources and thus provided a basis for arbitrage. In addition to this, most of the banks with decentralised funding structures employed inconsistent LTP regimes and relied on manual off-line processes to update funding costs.

2.1.3 Trading book funding policies and identifying funding requirements

Probably the worst LTP practices identified in the survey were in relation to trading and investment banking activities. A combination of poorly designed trading book policies, inadequate risk controls and limits, as well as a lack of oversight were to blame. For example, some banks that took part in the survey lacked trading book funding policies and

procedures, which allowed for over-aggressive trading behaviour and the accumulation of illiquid assets in search of revenues, not risk-adjusted profit. Most of the banks included in the survey did have trading book funding policies, but nearly all of these policies assumed that assets were only held short-term (ie for 180 days or less). One problem with this approach is that, irrespective of whether assets are likely to be held for more than the 180-day threshold, long-term funding charges only apply when assets roll from the trading book to the banking book. This provided little incentive for banks to develop risk controls and limits to adequately measure, monitor and assess the liquidity risk in traded assets, and was evident through the build-up of positions that were highly illiquid.

Many of the larger banks included in the survey, particularly those with substantial trading businesses, lacked a line of sight to individual business balance sheets, and thus could not identify the funding requirements of individual trading desks. As a result, trading and investment banking activities were funded based on the total net funding requirement across all related business units. This method essentially provides a line of credit to the trading book, and gives no regard to the liquidity risk embedded in business activities. This approach is therefore considered to be poor practice. On a separate but related issue, banks with large trading businesses that participated in the survey also applied insufficient haircuts to many of the traded assets they held. These banks clearly underestimated the likelihood of a market disruption, and the extent to which market liquidity could evaporate. The severe drop in market prices led to calls on margin positions and placed severe pressure on banks' abilities to meet funding requirements. Part of the reason this occurred was because no one had previously thought of the need to price the liquidity costs of potential margin calls.

2.1.4 Oversight

Ineffective oversight of the LTP process contributed to many of the problems that were identified at banks that took part in the survey. For example, the accrual of long-term illiquid assets and short-term volatile liabilities created a large and poorly understood mismatch between the maturities of assets and liabilities, and therefore exposed banks to greater structural liquidity risk. Probably the most striking example highlighting the implications of poor oversight was how some of the banks' LTP processes enabled them to accumulate significant amounts of highly rated, yet highly illiquid, tranches of collateralised debt obligations (CDOs) in their respective trading accounts. These portfolios were assumed to be safely funded with much shorter-term liabilities, typically in the order of overnight to 90-day funds.

2.1.5 Towards better LTP practice

In one form or another, all of the banks included in the survey are enhancing the way LTP is managed. A large portion of the banks surveyed, for example, are creating LTP policies for the first time to outline the purpose of LTP and, to provide some principles and/or rules to ensure business units understand the reasoning behind charges relating to the use of liquidity.

For the small proportion of banks in the survey that were operating with decentralised funding centres, they are all moving towards having wholesale funding managed centrally by a treasury function. In part, this is to restrict arbitrage between business units and treasury, and between business units themselves.

The survey also identified a small number of banks that are developing trading book policies and procedures for the first time. To complement this change, these banks are also developing risk controls and limits for trading activities to properly measure, monitor and assess the liquidity risk embedded in products and business units.

Most of the banks included in the survey, however, were found to be updating existing policies. The most notable enhancement includes the application of higher funding charges to trading positions that are more likely to become "stale" (ie positions that have a higher probability of rolling from the trading book to the banking book). Banks in this category are also enhancing existing risk controls and limits to better manage liquidity risk exposures.

The follow-on effects from these enhancements are improving risk-adjusted profit measures, and this is prompting business units to consider the cost of liquidity as part of their decision to book certain assets. Haircuts on traded assets are also being widened to account for more severe and prolonged market disruptions, and to ensure that assumptions surrounding the amount of liquidity that can be generated during a crisis are appropriately conservative.

Banks that were included in the survey whose trading book exposures are small relative to their main business activities, are attempting to curb over-trading behaviour by imposing higher funding charges on net funding requirements when certain funding limits are breached. Other banks in the survey, whose trading book exposures are large relative to their main business activities, are devoting more attention to understanding the funding requirements of individual trading desks, and are looking to apply charges on a more granular basis.

Across all banks in the survey, there is an emphasis on improving oversight. Management at all levels, treasury functions, as well as independent risk and financial control personnel are becoming more engaged in the LTP process. In addition, meetings to discuss changes in funding costs are being held more regularly, for example, monthly instead of quarterly or semi-annually, as they were prior to the financial crisis.

In one form or another, all of the banks included in the survey are enhancing the way their LTP process is managed. One positive stemming from these enhancements is that related

parties involved in the management of LTP are being forced to better understand the LTP process. Broadly, banks are encouraged to continue with similar changes.

2.2 Liquidity Management Information Systems (LMIS)

LMIS are widely used by management as a primary source of measuring and monitoring the performance of businesses. LMIS provide information that assists management in liquidity strategic decision-making. In this regard, LMIS play a pivotal role in helping management achieve group-wide goals. Weak LMIS could easily distort the information for decision-making and prevent the bank from achieving its objectives.

One application of LMIS is to support internal pricing mechanisms. In relation to LTP, LMIS enable the costs, benefits and risks of liquidity to be attributed to appropriate business activities. Many of the pre-2009 LMIS employed by banks that were included in the survey were too basic, and this limited the effectiveness and efficiency of the LTP process. In some cases, for example, the basic and rigid nature of LMIS meant that certain business activities failed to receive a charge for the cost of liquidity or, conversely, a credit for the benefit of liquidity. Another weakness in many of the LMIS that were employed by banks in the survey was that they prevented the costs, benefits and risks of liquidity from being attributed at a sufficiently granular level.

The SSG (2009) also report similar findings. “Many firms acknowledged shortcomings in their LMIS infrastructure and in their ability to produce useful reports during the crisis, recognizing that better-quality and more timely liquidity reporting was essential to effective management of liquidity and funding issues during a crisis” (p 15). In light of these shortcomings, the SSG recommended that banks improve their LMIS.

One implication of the weakness in LMIS and the poor LTP practices that resulted is that businesses reported performance (and employees claimed bonuses) on a basis

that might not have reflected their actual performance. Essentially, this would limit management's ability to monitor performance, accurately distinguish good performing businesses from those that were not performing so well, and make reliable decisions pertaining to their objectives.

A large proportion of the banks included in the survey are in the process of upgrading LMIS after their short-comings were unmasked by the GFC. From a supervisory perspective, it is essential that this upgrading continues. As outlined above, LMIS are an essential part of the decision-making process so it is vital that the information they provide is accurate and reliable. Upgrading LMIS in a large bank is a costly and long-term process. But the benefits of appropriately charging business activities for the cost, benefits and risk of liquidity, and at a sufficiently granular level, will far outweigh the costs and limitations of the basic LMIS that were previously employed. LMIS that are sufficiently advanced to achieve these outcomes will promote better LTP practice.

2.3 Remuneration practices

If designed well, incentive pay can have enormous benefits. It encourages behaviour that is consistent with the culture of an institution, and assists management in achieving group-wide objectives. On the other hand, poorly designed remuneration can promote perverse behaviours such as excessive risk-taking, which could severely impact the performance of an institution.

In 2009, the Financial Stability Board (FSB) reported that poor remuneration practices were one of the factors that contributed to the GFC. "High short-term profits led to generous bonus payments to employees without adequate regard to the longer-term risks they imposed on their firms.

These perverse incentives amplified the excessive risk-taking that severely threatened the global financial system and left firms with fewer resources to absorb losses as risks materialised" (p 1).[12]

The SSG also identified poor remuneration practices as one of the factors that contributed to the funding and liquidity problems witnessed at some banks during the recent crisis. Following their survey of firms, the SSG reported that remuneration was largely insensitive to the risks taken to generate income, and to costs associated with long-term funding commitments that were required to hold illiquid assets (p 24).

Similarly, many of the banks that participated in the survey on LTP failed to adequately account for the costs, benefits and risks of liquidity in the pricing and performance assessment of various products and business units. As a result, profit measures used as a basis for determining remuneration were often distorted. Profit pools, for example, which are generally used to determine short-term incentives, or bonuses for employees, were derived from a simple percentage of accrued revenues, without any regard to the cost of liquidity (or capital). This placed more emphasis on maximising revenues rather than risk-adjusted earnings.

Another reason the costs, benefits and risk of liquidity were poorly allocated through the LTP process for most of the banks in the survey, was because of the way remuneration was structured, particularly for those employees responsible for oversight. For many staff in these areas, remuneration was designed such that it largely depended on the performance of front-line businesses they were responsible for overseeing. Thus, including the actual costs for liquidity would have

[12] *Principles for Sound Compensation Practices,* FSB, (April 2009) is available at http://www.financialstabilityboard.org/list/fsb_publications/tid_123/index.htm.

impacted negatively upon business unit performance, which inevitably would have reduced personal remuneration and benefits for employees. Clearly, this would have also impacted the independence of their role.

Recognising these weaknesses, many of the banks that took part in the survey are developing their respective LTP processes to ensure that profit and performance measures include the relevant costs for liquidity (and capital, although this is a separate issue). Under the new regime, assets will receive a charge for the cost of liquidity consistent with the positions that are funded. In addition to this, many of the larger banks in the survey are moving towards re-designing remuneration for persons in risk control positions consistent with Principle 3 of the FSB's Principles for Sound Compensation Practices. Principle 3 states that, "staff engaged in financial and risk control should be compensated in a manner that is independent of the business areas they oversee and commensurate with their key role in the firm" (p 2). It is envisaged that this will also promote a more appropriate attribution of liquidity costs to business activities and restore independence in these vital roles.

3. LTP in practice: managing on-balance sheet funding liquidity risk

3.1 Why banks need LTP

In their daily operations, banks make money by funding long-term loans (assets) with short-term deposits (liabilities), a process that is commonly referred to as maturity transformation. As pointed out by the BCBS (September 2008), this "makes banks inherently vulnerable to liquidity risk, both of an institution-specific nature and that which affects markets as a whole" (p 3). But provided banks use LTP to account for the costs, benefits and risks of liquidity in product pricing, new product approval processes and profit and

performance assessments, they should not be discouraged from engaging in maturity transformation. Banks with poor LTP practices are more likely to accrue larger amounts of long-term illiquid assets, contingent commitments and shorter-dated volatile liabilities, substantially increasing their vulnerability to funding shortfalls.

3.2 An example of what can go wrong with poor LTP

In October 2009, the SSG revealed that firms that encountered the most severe funding and liquidity problems through the financial crisis were those that relied excessively on short-term financing of longer-term illiquid assets. That is, those that engaged most in maturity transformation. The SSG highlighted that one of the drivers behind the development of these business models was poor LTP practices, which failed to penalise businesses for the liquidity risk embedded in the assets that were booked, and which also allowed banks to build up significant amounts of contingent liquidity risk in off-balance sheet exposures. These banks made large apparent profits before the GFC, but failed to recognise that these profits were based upon what proved to be extraordinarily fragile liquidity arrangements.

The international survey identified many poor LTP practices, which reflected weaknesses in the LTP methods/approaches that were used to manage funding liquidity risk. These are discussed in more detail below.

3.3 "Zero" cost of funds approach – liquidity as a "free" good

Probably the most striking example of poor practice identified in the survey was that some banks failed to account for the costs, benefits and risks of liquidity in all or some aspects of their business activities. These banks came to view funding liquidity as essentially free, and funding liquidity risk as essentially zero. As a result, there was simply no charge

attributed to some assets for the cost of using funding liquidity, and conversely no credit attributed to some liabilities for the benefit of providing funding liquidity. This was undoubtedly the worst practice identified in the survey. Figure 2 below provides a graphical representation of what this would look like in practice. Note that the rate charged to users of funds in this instance would have been derived from the swap curve only. If we assume that interest rate risk is properly accounted for using the swap curve, then a zero spread above the swap curve implies a zero charge for the cost of funding liquidity.

Figure 2

Zero cost of funds approach to LTP

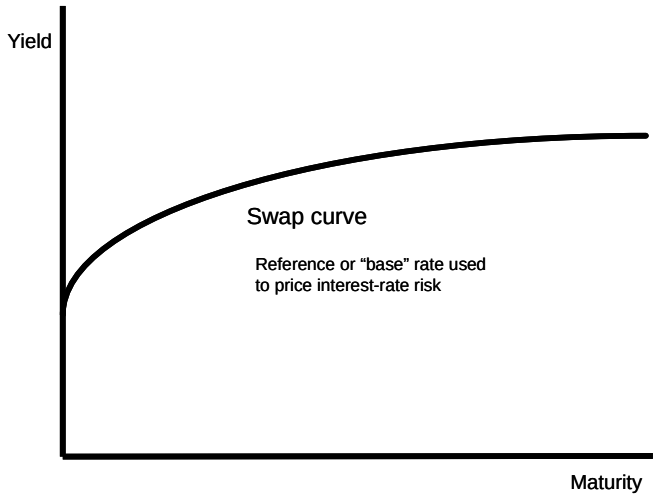

A zero charge for the cost of liquidity and, conversely, a zero credit for the benefit of liquidity exacerbated maturity transformation to the largest degree possible. This approach resulted in the hoarding of long-term highly illiquid assets, and very few long-term stable liabilities to meet funding demands as they became due.

3.3.1 Why did some banks choose this approach?

Ideal funding conditions in the years preceding the crisis could provide one explanation of why some banks viewed liquidity as a free good, and funding liquidity risk as essentially zero. Figure 3 below shows how the spread between one-year LIBOR and the one-year swap rate changed during the period June 2005 to October 2010.

Figure 3

1-year LIBOR/Swap spread (Currency = USD)

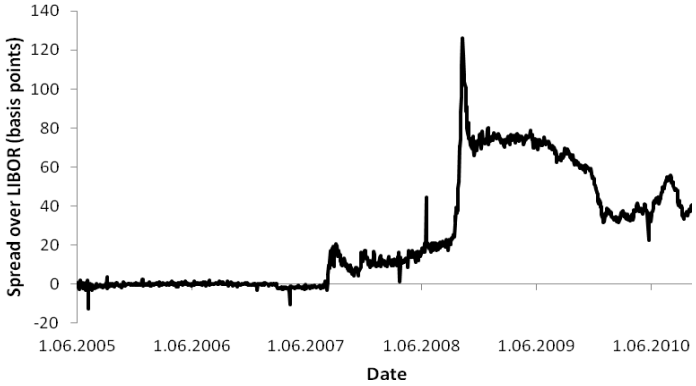

Source: Bloomberg.

In June 2005, at the peak of robust share market growth, the spread was only 0.5 basis points (bps). With funding conditions so easy, it is likely that banks viewed spreads as pure credit risk adjustments and neglected (ignored) funding liquidity risk altogether. If banks believed funding would always be available and at permanently cheap rates, this simply could have masked the need to charge assets for the cost of liquidity, and conversely, credit liabilities for the benefit of liquidity.

3.4 Pooled "average" cost of funds approach to LTP

Some banks recognised the need to charge users and credit providers of funding liquidity and employed a pooled approach to LTP, where an average rate was calculated based on the interest expense (cost of funds) across all existing funding sources. For example, if deposits were a bank's only source of funding the average rate would be based on the total interest expense for all deposits divided by average total deposits, adjusted for floats and reserve requirements. This approach is much better than the zero cost of funds approach, but because there is only one "average" rate calculated, all assets irrespective of their maturity are charged the same rate for their use of funds (cost of liquidity), as depicted in Figure 4 below.

Figure 4

Single average for the cost and benefit of funds

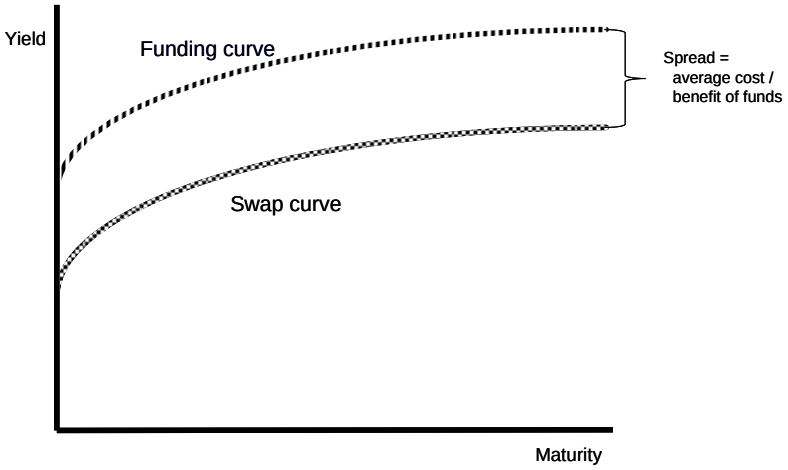

To illustrate how charges and credits for the use and benefit of funds would be allocated under an average cost of funds approach, consider the following example. If the average rate across all funding sources was 10 bps, all loans would receive

a charge of $1,000 on a principal amount of $1 million, irrespective of their maturity. Assuming this rate was also used to reward fund providers, then all deposits would receive a credit of $1,000 on a principal amount of $1 million, irrespective of their maturity. This can be seen in Table 1.

Table 1

Costs and benefits of funds under an average cost approach

Term in years	1	2	3	4	5
Loan/deposit principal	$1 million	$1 million	$1 million	$1 million	$1 million
Average cost of funds (bps)	10	10	10	10	10
Charge for use of funds	$1,000	$1,000	$1,000	$1,000	$1,000
Credit for benefit of funds	$1,000	$1,000	$1,000	$1,000	$1,000

3.4.1 Problems with the pooled "average" cost of funds approach

Despite its simplicity there are two major weaknesses with this approach. First, it ignores the heightened liquidity risk embedded in longer-term assets. Charging one "average" rate for the use of funds inherently assumes that all assets, irrespective of their maturity, pose the same liquidity risk. Moreover, if this "average" rate is also used to credit fund providers, then an incentive to write loans will be met with a direct disincentive to gather deposits. For example, decreasing the rate charged to fund users from 10 bps to five bps will encourage loan generation, but at the same time, this will provide less incentive for business units to raise deposits.

Having separate "average" rates for the costs and benefits of funds is a better approach. This is depicted in Figure 5 below.

Figure 5

Separate averages for the cost and benefit of funds

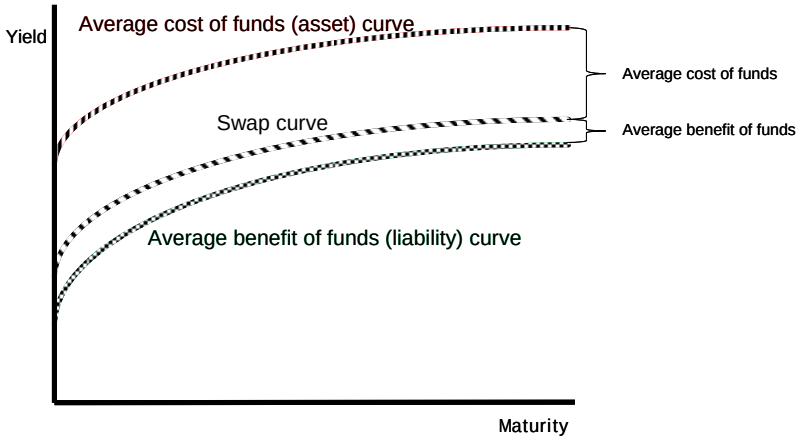

To illustrate the effect of having separate average rates for the cost and benefit of funds, consider the following example. If the average cost of funds is 10 bps, as in the example presented above, all loans would be charged $1,000 on a principal of $1 million, irrespective of their maturity. Further, if the average benefit of funds is four bps, all deposits would be credited $400 on a principal of $1 million, irrespective of their maturity. Under this approach, lowering the average cost of funds from 10 bps to five bps will encourage loan generation. However, because of the separate rate for the average benefit of funds, this change will not directly discourage business units from raising deposits. This information is presented in Table 2.

Table 2

Costs and benefits of funds under a separate average cost approach

Term in years	1	2	3	4	5
Loan/deposit principal	$1 million	$1 million	$1 million	$1 million	$1 million
Average cost of funds (bps)	10	10	10	10	10
Average benefit of funds	4	4	4	4	4
Charge for use of funds	$1,000	$1,000	$1,000	$1,000	$1,000
Credit for benefit of funds	$400	$400	$400	$400	$400

However, in much the same way as one "average" rate for fund users ignores the heightened liquidity risk in longer-term assets, having one "average" rate for fund providers ignores the increased benefits of liquidity in longer-term liabilities. That is, all deposits irrespective of their maturity receive the same credit for the benefit of funds, as can be seen in Table 2.

Second, using an average cost of funds reflects historical rates and prices, but does not appropriately reflect the actual market cost of funds. If five-year funding was to increase by 20 bps, for example, the respective change in the average (cost of funds) rate would be much less. Changes in the actual market cost of funds would need to be sustained for a period of time for the effect to be fully integrated into the average cost of funds. Because the average cost of funds lags changes in the actual market cost of funds, it does not appropriately reflect market perceptions of risk for new business entering a bank's books.

3.4.2 Implications of pooled "average" cost of funds approach

Promotes maturity transformation

One implication of employing a pooled "average" cost of funds approach to LTP is that it promotes unhealthy as well as healthy maturity transformation. Business units will be unduly encouraged to write long-term assets because they do not receive higher charges for their use of funds over a longer period. Conversely, business units will be discouraged from raising long-term liabilities because there is no premium credited to liabilities that provide funding for longer periods of time. The net effect of this is a larger mismatch between the maturities of assets and liabilities on banks' balance sheets, which inherently exposes them to greater structural liquidity risk. This point is supported by the SSG (2009), which claims that "borrowers had taken advantage of the opportunity the market afforded to obtain short-term (often overnight) financing for assets that should more appropriately have been funded with long-term, stable funding" (p 2).

Moreover, some institutions ignored maturity mismatch liquidity risk by not appropriately match-funding originated transactions in their funds transfer pricing (FTP) systems on a cash-flow basis. When combined with average costs of existing funds across all funding sources to banks' balance sheets, the lack of reference to term cash-flow matched funding entailed the cross-subsidisation of longer-dated liquidity risk at the expense of shorter-dated risk. Such subsidisation skewed business incentives and behaviours to the detriment of bank soundness.

Other factors, such as remuneration and information asymmetries, naturally encourage long-term asset generation, but under an average cost of funds approach the incentive is exacerbated. For example, if remuneration is based on performance, which is measured via net interest income, businesses will ordinarily be encouraged to write long-term loans because they generate more interest income, with less effort, over several years. Where an average cost of funds approach is employed, this incentive becomes even more

attractive for business units because assets that require funding for longer periods of time are not charged more for the cost of liquidity. In regard to information asymmetries, business unit managers are likely to know more about their businesses' activities than treasury. Hence, if business unit managers believe treasury is under-charging for the use of long-term funds, it will naturally encourage them to write long-term assets. But since all funds are charged the same rate for the use of funds under an average cost of funds approach, where information symmetries exist, this incentive will be magnified. A similar but opposite effect will exist for liabilities.

Distorts profit assessment

Another implication of the pooled average cost of funds approach to LTP is that it distorts profit assessment. As outlined above, the average cost of funds lags changes in banks' actual market cost of funds, especially in volatile markets. Banks employing this approach found that their pricing methodologies resulted in the mispricing of and accumulation of assets on significantly distorted risk-adjusted terms. This made it difficult to identify poor performing products and business units on a risk-adjusted basis.

There are several reasons why some of the banks included in the survey might have chosen to adapt a pooled average cost approach to LTP. First, averaging funding costs across all assets is much simpler than having to charge individual assets, products or transactions based on their contractual or behavioural (expected) maturities. Second, the simplicity of the average cost of funds approach makes it easier for business units to understand the LTP process and therefore provides more incentive for them to comply. Third, under this approach, the LTP process could be managed efficiently using basic LMIS. Fourth, the average cost of funds is less susceptible to intermediate changes in banks' actual market cost of funding, thereby reducing net interest income volatility across businesses. This is advantageous because it limits the subjective decision-making of business unit managers and provides central management with more control over group-wide objectives.

3.5 Matched-maturity marginal cost of funds approach to LTP

A matched-maturity marginal cost of funds approach to LTP is current best practice for assets and liabilities on the balance sheet. From banks' actual market cost of funding, this approach calculates the portion of the cost that is attributable to liquidity. It seeks to achieve this by converting fixed-rate borrowing costs to floating-rate borrowing costs through an internal swap transaction and observing the spread over the reference rate, which is depicted from the swap curve. This spread is usually referred to as a term liquidity premium and is the rate that charges assets for the use of funds, and credits liabilities for the benefit of funds. This is presented graphically in Figure 6.

Figure 6

Matched-maturity marginal cost of funds approach to LTP

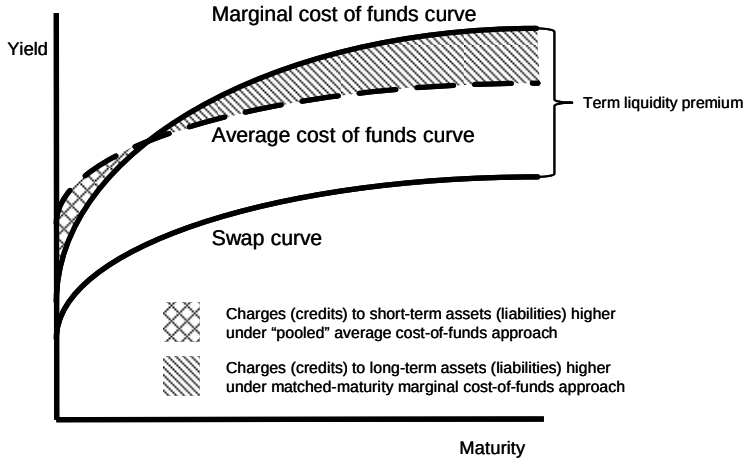

To explain this process more fully, banks incur fixed-rate costs when issuing unsecured wholesale term debt. Using these

costs alone it is difficult to strip out the portion that is attributable to liquidity. But swapping fixed rate costs to floating rates provides a solution. The process generally involves stripping structured debt issuances into embedded derivatives and floating rate cash instruments, which are pegged to a reference rate.[13] The spread above the reference rate is the rate that values the internal swap transaction at par. This is the term liquidity premium. It reflects both idiosyncratic credit risks and market access premiums and is considered to be a much better measure of the cost of liquidity than an average cost of funds.

Reference rates are generally depicted from a swap curve, which is constructed from a combination of LIBOR or Euribor rates for funding up to one year, and interest rate swaps for funding above one year. This curve reflects a term structure of interbank lending rates. Although credit risk is somewhat mitigated by the fact that principal amounts are not exchanged between respective parties in a swap agreement, swap curves are still considered to provide better estimates of "base" reference rates for the purpose of teasing out liquidity than, say, government curves. This is because swap curves more closely reflect the risks to which banks are exposed when borrowing and lending money in the interbank market. Swap curves also capture changes in general market conditions.

3.5.1 How are rates for users and providers of funds determined?

Under the matched-maturity marginal cost of funds approach, rates charged for the use of funds and, conversely, rates credited for the benefit of funds are based on the term liquidity premiums corresponding to the maturity of the transaction, or in the case of amortising or indeterminate-maturity products,

[13] This process is described in detail in Matz and Neu (2007).

blended term liquidity premiums consistent with their known or estimated cash-flow profiles.

Even though the matched-maturity marginal cost of funds approach to LTP is considered to be better practice, some of the more advanced banks surveyed that had employed this method failed to actively update term liquidity premiums. As a result, assets were mispriced and risk-adjusted profit assessments were distorted, especially as market volatility increased in the early stages of the GFC.

3.6 Examples of pricing funding liquidity risk

To illustrate how the matched-maturity marginal cost of funds approach should be applied in practice, and to compare it to the average cost of funds approach, assume the following term liquidity premiums and average cost of funds were recorded by a bank at a point in time prior to the crisis (pre-GFC), and more recently (current).

Table 3

Pre-GFC and current
Term liquidity premiums and average cost of funds

In basis points

Term in years	1	2	3	4	5
Panel A: pre-GFC					
Term liquidity premium	1	2	3	6	10
Average cost of funds	2	2	2	2	2
Panel B: current					
Term liquidity premium	5	10	18	28	40
Average cost of funds	8	8	8	8	8

Some examples of how LTP should apply to various transactions are presented below.

3.6.1 Non-amortising bullet loans

As the name implies, non-amortising bullet loans provide no repayments (cash flows) throughout the life of the loan. Since all principal and interest is repaid at maturity, a funding commitment is required for the entire life (term) of the loan. Hence, using a matched-maturity marginal cost of funds approach, a one-year non-amortising bullet loan should have received a charge of one bp (Panel A), if originated pre-crisis, and five bps (Panel B) if originated more recently. For simplicity, if the principal of the loan was $1 million, this should have translated to charges of $100 and $500, respectively to the business unit(s) writing the loans. In much the same way, a five-year non-amortising bullet loan should have received a charge of 10 bps (Panel A) if originated pre-crisis, and 40 bps (Panel B) if originated more recently. Assuming the same loan principal of $1 million, this should have translated to charges of $1,000 and $4,000, respectively, to the business unit(s) writing the loans.

In contrast, had an average cost of funds approach been applied, both the one- and five-year non-amortising loans would have been charged two bps (Panel A) if originated pre-crisis, and eight bps (Panel B) if originated more recently. Table 4 below presents the differences in the charges for the uses of funding for each of the non-amortising bullet loans in this example.

Table 4 shows that non-amortising bullet loans with a term of one-year would have received a higher charge for the use of funding if banks applied an average cost of funds approach rather than a matched-maturity marginal cost of funds. However, for all other maturities, the opposite is true. Using an average cost of funds approach in the pre-crisis period, a five-year loan would have been undercharged eight bps ($800 on a loan of $1 million). If the same loan was originated more recently it would have been undercharged 32 bps ($3,200 per

$1 million). This example highlights one of the major weaknesses of the average cost of funds method, viz., its inability to immediately reflect changes in the actual market cost of funds. For banks in the survey employing this approach, it would have encouraged business units to write long-term loans at the expense of short-term deposits.

Table 4

Differences in funding charges

Basis points

Term in years	1	2	3	4	5
Panel A: pre-GFC					
Term liquidity premium	1	2	3	6	10
Average cost of funds	2	2	2	2	2
Difference	–1	0	1	4	8
Panel B: current					
Term liquidity premium	5	10	18	28	40
Average cost of funds	8	8	8	8	8
Difference	–3	2	10	20	32

3.6.2 Amortising loans

Unlike non-amortising bullet loans, amortising loans do provide repayments (cash flows) throughout the life of the loan. Since a portion of principal (and interest) is repaid prior to maturity, a funding commitment is not normally required for the entire life (term) of the loan. This is because, at some point between origination and maturity, the loan becomes self-funding.

Consider the simplest loan type in this category, a five-year linearly amortising bullet loan with a principal amount of

$1 million. If you think of this as five separate annual loans, each of $200,000, using a matched-maturity marginal cost of funds approach, this loan (assuming it was originated pre-crisis) should have received a charge of:

$$\frac{1(1)+2(2)+3(3)+4(6)+5(10)}{1+2+3+4+5}=5.9\,\text{bps}.$$

This is a tenor-weighted (blended) term liquidity premium, derived from what is commonly referred to as the tranching approach[14]. Following this same approach, if the loan was originated recently it should have received a charge of:

$$\frac{1(5)+2(10)+3(18)+4(28)+5(40)}{1+2+3+4+5}=26.1\text{bps}.$$

In both cases, the charge for the use of funds indicates that a funding commitment is required for somewhere between three and four years and not the entire term of the loan, which was five years.

If an average cost of funds approach had been employed, the loan originated pre-crisis would have received a charge of two bps (Panel A). This would have resulted in an undercharge of 3.9 bps (5.9 – 2). If the loan had been originated more recently, it would have received a charge of eight bps, which would have resulted in an undercharge of 18.1 bps (26.1 – 8). Although the differences in the funding charges are not as severe as in the non-amortising bullet loan example above, it still highlights the weakness of the average cost of funds approach in reflecting changes in the actual market cost of

[14] Another method used by banks to calculate the charge for the use of funds is the internal rate of return (IRR) approach. This involves calculating an IRR using the rates depicted from the swap curve, and an IRR using the rates depicted from the marginal cost of funds curve. The difference between the resulting rates is the rate used to charge business units for the use of funds. For more detail, and an example of this approach, see Matz and Neu (2007).

funding. Once again, this would have encouraged long-term loan (asset) generation.

Not all amortising loans provide known cash flows for the entire life of the loan. Take standard variable- (adjustable-) rate mortgages, for example. Often their contractual maturity will be 25 or 30 years at origination, but their actual maturity will vary depending on factors such as repayment frequency and repayment amount.

The uncertainty surrounding future cash flows makes it more difficult to calculate an appropriate charge for the commitment of funds required to service these types of loans. For example, a simple tenor-weighted (blended) term liquidity premium cannot be derived because of the unknown timing of future cash flows.

Consider a standard $500,000 variable rate mortgage, with a contractual term of 25 years. Attributing a 25-year term liquidity premium essentially overcharges the loan for the cost of funding liquidity and could discourage asset growth. A better approach is to bundle mortgages into monthly vintages, based on their origination date, and model the repayment history (decay) over time as depicted in Figure 7 below.

Figure 7

Mortgage run-off for 3 different vintages

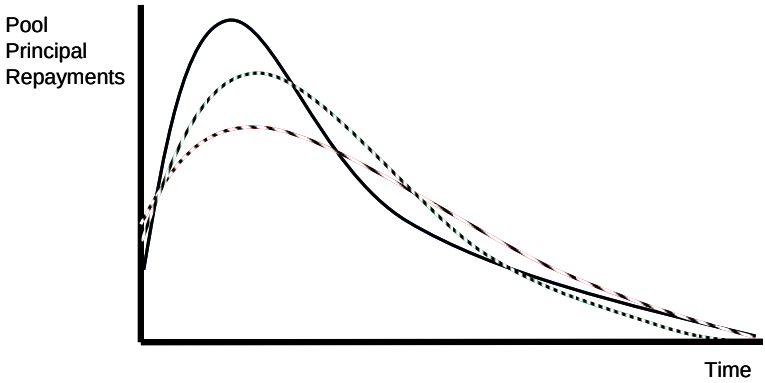

If mortgages tend to behave similarly, as highlighted in the figure above, irrespective of the vintage to which they belong, then a single charge for funding liquidity can be attributed to the entire portfolio, instead of to each individual transaction. This charge is based on the behavioural maturity of the portfolio, which is often calculated by banks using the weighted-average life (WAL) method.

$$WAL = \sum_{i=1}^{n} \frac{p_i}{P} t_i \qquad (1)$$

where P_i = principal amount in distribution *i*, P = amount of loan, and t_i = time (in years) of payment *i*.

WAL can be interpreted as the weighted-average time it takes to recoup \$1 of principal (ie the time it takes for the loan to start paying for itself).[15]

As an example, suppose a large bank writes around \$2 billion of mortgage loans on average, per month, and upon examining the decay of its loans finds the behavioural maturity (WAL) of the mortgage portfolio to be approximately four years. If a matched-maturity marginal cost of funds approach is employed, then all mortgage loans should receive a charge, at point of origination, based on the four-year term liquidity premium. Using the figures from Table 3, this would be six bps (Panel A) or 28 bps (Panel B) depending on when the loan was originated.

Across the entire portfolio, this would translate into dollar charges of \$1.2 million or \$5.6 million, respectively. In contrast, if an average cost of funds approach is employed,

[15] The WAL is not the time it takes to repay 50 per cent of the loan. That would be a median calculation. The WAL is an average. Only in the special case of when the interest rate on the loan is zero, will 50 per cent be repaid at the WAL. As the interest rate increases from zero, less than 50 per cent of the loan will be paid at the WAL. This is because most of the initial repayments comprise interest and not principal.

mortgage loans should receive a charge of two bps (Panel A) if originated pre-crisis, and eight bps (Panel B) if originated more recently. Collectively, for all mortgage loans, this would translate into dollar charges of $400,000 or $1.6 million, respectively.

This example further demonstrates how the average cost of funds lags changes in banks actual market cost of funds and, at the same time, highlights how costly this could be, especially when products are priced at the portfolio level and comprise a large portion of bank assets.

3.6.3 Deposits

Because deposits are a source of funding for banks, business units responsible for raising retail, and in some cases wholesale, deposits should be credited for the benefit of liquidity they provide.

Deposits should be categorised as “sticky” or “hot/volatile” and credited based on their likelihood of withdrawal. As a general rule, sticky money, such as term deposits, are less likely to be withdrawn and should therefore receive larger credits than hot/volatile money, such as demand deposits, savings and transaction accounts, which are more likely to be withdrawn at any time.

Using a matched-maturity marginal cost of funds approach, term deposits should receive a credit based on their maturity. For example, using the same figures as presented in Table 3, a one-year term deposit should have been credited one bp if originated pre-crisis and five bps if originated more recently. Similarly, a five-year term deposit should have received a credit of 10 bps if originated pre-crisis and 40 bps if originated more recently.

Had an average cost of funds approach been employed, all term deposits would have received a credit of two bps if originated pre-crisis and eight bps if originated more recently, irrespective of their maturities. In the pre-crisis period, this would have resulted in over compensating the one-year term

deposit by one bp (2 – 1), and by three bps (8 – 5) if originated more recently. The five-year term deposit on the other hand would have been under compensated by eight bps (2 – 10) if originated pre-crisis and a staggering 32 bps (8 – 40) if originated more recently. As above, this example highlights the limitations of the average cost of funds approach. For banks employing this approach, it would have encouraged business units to raise short-term deposits rather than long-term, more stable sources of funding. Collectively, with the finding from above, this would have led to more structural liquidity risk on the balance sheet.

Hot/volatile sources of deposits are often referred to as indeterminate maturity products, given the uncertainty surrounding their cash flows. However, despite being categorised as hot or volatile, these types of deposits sometimes provide stable sources of funding. Demand deposits, for example, can be withdrawn at any time without notice. But, if all similar accounts were to be pooled and the behaviour of the cash flows modelled over time, there would be a proportion that is rarely withdrawn (stable or core part) and a proportion that is more often withdrawn (hot or volatile part). Making this distinction is important, because if a bank were to simply apply a matched-maturity marginal cost of funding approach, all demand deposits would only receive a credit based on the overnight term liquidity premium. Given this is likely to be very close to zero, which translates to a cheap funding source for banks, business units would be discouraged from raising demand deposits. A better approach would be to assign larger credits to core parts of funding, based on the modelled behavioural maturity, and smaller credits to hot/volatile parts of funding.

Banks employing an average cost of funds approach would have no incentive to make the distinction between core and volatile parts of funding since, under this approach, the same credit for the benefit of funding is applied to all deposits, irrespective of their maturity.

3.7 Summary

Failing to price liquidity is unacceptably poor LTP practice for a bank, and supervisors should not tolerate this failing.

The average cost approach to LTP is simple, but has two major defects. First, it neglects the varying maturity of assets and liabilities by applying a single charge for the use and benefit of funds and, second, it lags changes in banks' actual market cost of funding. These defects essentially promote maturity transformation, which inherently exposes banks to more structural (mismatch/funding) liquidity risk.

Overall, a matched-maturity marginal cost of funds approach promotes better LTP practice. It is more complex than the pooled average cost of funds approach, but it has some significant advantages. First, it recognises that the costs and benefits of liquidity are related to the maturities of assets and liabilities, and therefore allows higher rates to be assigned to products that use or provide liquidity for longer periods of time. Second, it recognises the importance of having changes in market conditions incorporated quickly and efficiently into the rate used to charge and credit users and providers of funds, and therefore relies on the actual market cost of funds. Banks should be encouraged to move towards this approach, if they are not already doing so.

4. LTP in practice: managing contingent liquidity risk

For many on-balance sheet items, calculating the charge for using, or the credit for providing, funding liquidity is quite straightforward. However, the same cannot be said about contingent commitments such as lines of credit, collateral postings for derivatives and other financial contracts, and liquidity facilities to name a few. In these cases, the best approach is to impose a scenario model, determine a reasonable low probability worst-case outcome and charge at

the most granular level the transaction, product, or business unit for the costs of covering this outcome.

Banks carry a liquidity cushion, a "buffer" of highly liquid assets or, alternatively, stand-by liquidity to help them survive periods of unexpected funding outflows. A graphical illustration of this is depicted in Figure 8.

Figure 8

Unexpected funding outflows and the need for a liquidity cushion

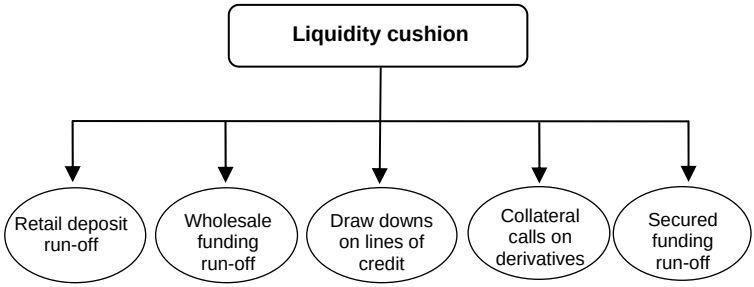

In December 2010, the BCBS published two global standards for liquidity risk. First, a Liquidity Coverage Ratio (LCR) to ensure banks have sufficient high quality liquid assets to meet their daily net cumulative cash outflows during an idiosyncratic shock, for a period of one calendar month. Second, a net stable funding ratio (NSFR) aimed at reducing banks structural liquidity risk by encouraging the use of longer-term funding of assets and other business activities.[16] The move to make banks more self-sufficient and stable over a longer period is in

[16] For more details, see *Basel III: International Framework for Liquidity Risk Measurement, Standards and Monitoring*, BCBS, (December 2010).

part to reduce the burden of central banks having to act as the lender-of-last-resort, and the potential implications of moral hazard as a result of these actions.

4.1 Liquidity cushions: a principle of liquidity risk management

Liquidity cushions are considered a fundamental principle for the management of liquidity risk. This is clearly outlined in Principle 1 of the BCBS *Principles for Sound Liquidity Risk Management and Supervision* (September 2008), and also reinforced by Principle 12, which states that "a bank should maintain a cushion of unencumbered, high quality liquid assets to be held as insurance against a range of liquidity stress scenarios, including those that involve the loss or impairment of unsecured and typically available secured funding sources. There should be no legal, regulatory or operational impediment to using these assets to obtain funding" (p 4).

4.2 Extant guidance focuses on size, composition and marketability

Extant guidance provided in association with liquidity cushions focuses mainly on size, composition and marketability of the assets contained within the cushion. To ensure banks' liquidity cushions are adequately sized, the BCBS recommends they be aligned with stress-testing outcomes that consider both idiosyncratic and systemic scenarios, plus a combination of the two. The level of stress assumed in the tests should reflect a bank's overall risk tolerance. To assess their risk tolerance, banks should consider factors such as structural liquidity risk (ie the liquidity gap between the weighted average maturities of assets and liabilities) and the complexities of both on- and off-balance sheet business activities, which affect the frequency and irregularities of cash flows.

According to the BCBS, liquidity cushions should comprise "a core of the most reliably liquid assets, such as cash and high quality government bonds or similar instruments, to guard against the most severe stress scenarios" (p 30). Banks should also consider the marketability of these assets. Although this is likely to vary in relation to the stress scenario and survival period (ie assets will generally remain more marketable throughout less severe market disruptions), there are some generic characteristics that tend to improve asset liquidity. For example, assets that are more transparent are generally also easier to value, and the certainty surrounding this will inherently improve marketability. In addition, assets that are central bank-eligible and/or have good market depth will generally be more marketable. A bank's reputation, credit rating and active participation in certain markets will also impact asset marketability (p 30).

4.3 Problems with banks liquidity cushions unveiled by the GFC

The recent crisis exposed some fundamental problems with banks' liquidity cushions. First, for the banks that participated in the survey, very few used the results of stress-testing to determine the size of their liquidity cushion. For the few banks that did consider the results of stress-tests, the size of their liquidity cushion was based on outcomes stemming from idiosyncratic funding scenarios only. Having little or no regard for prolonged market-wide disruptions meant that cushions were inadequately sized to protect the banks from larger-scale unexpected (contingent) outflows. On a separate but related issue, one of the flaws with many of the banks' stress-testing processes was that the parameters used were too narrow, and were based purely on historical data. This meant that events that had not previously occurred were neglected.

Second, most of the banks that were surveyed had liquidity cushions comprised of assets that were thought to be highly liquid, but were found to be highly illiquid and highly correlated. There were instances where some assets held as

stand-by liquidity were not unencumbered, meaning that the bank did not have legal claim over the asset or that the asset was not entirely free from debt.

Third, nearly all of the banks included in the survey funded their liquidity cushions short-term (e.g. overnight), consistent with the perception that funding could be easily accessed and any market disruption would only be short-lived. While this minimised negative carry costs, it also provided banks little incentive to attribute the relevant costs back to the businesses that created the need to carry additional liquidity. When assets in the cushion could not be sold to generate funding, it became apparent that the real cost of carrying stand-by liquidity was much greater than what the banks had assumed.

4.4 LTP and liquidity cushions – both principles, both treated separately

To date, there has been limited guidance about how to attribute the cost of carrying liquidity cushions, but this paper offers some ideas. Even though it is common practice for banks to attribute this cost via LTP, no link between LTP and liquidity cushions is established in extant material. In fact, LTP and liquidity cushions are very much treated as mutually exclusive principles for sound liquidity risk management.

4.5 Poor attribution of cost of carrying a liquidity cushion

Carrying a "buffer" of highly liquid assets is costly for banks because the cost of funding assets comprising the cushion generally outweighs the return they generate.[17] As such,

[17] The survey identified that this is the typical method banks use to calculate the cost of carry.

banks often seek to minimise the size of their liquidity cushion so that the negative carry does not drag on profits.

Most of the banks included in the survey consider the cost of carrying additional liquidity a cost of doing business, rather than an opportunity cost. In this regard, the cost of carry should not be borne by central management (ie treasury). It should instead be attributed back to businesses via the LTP process. While this appears common practice amongst banks participating in the survey, generally through incorporating a liquidity premium in the LTP process, most simply averaged the cost across all assets without giving specific attention to those businesses and products that generate the need to carry additional liquidity.

As a simple example, if it costs a bank $30 million to carry a buffer of liquid assets and total assets of the bank equate to $300 billion, then under the approach described above, all assets would receive a charge of one bp through the LTP process. This is depicted in Figure 9. Note that the one bp charge is in addition to the term liquidity premium, which is charged to assets based on their commitment of funds. This is highlighted in equation 2.

Figure 9

Recouping the cost of carrying a liquidity cushion via LTP

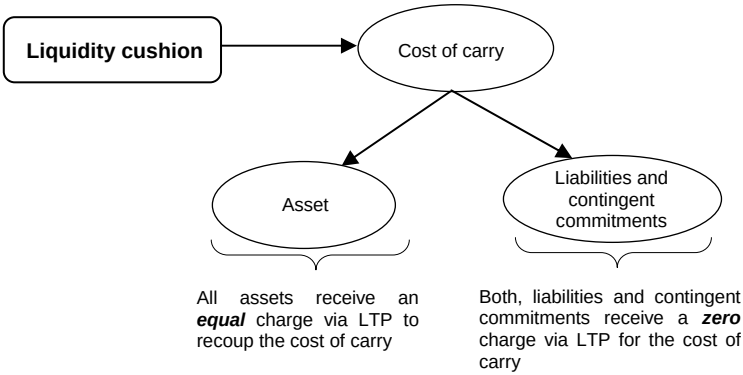

One explanation of why this approach might have been adopted by most of the banks surveyed is because they underestimated the "actual" cost of carrying additional liquidity. As mentioned above, assets in the cushion most often incurred short-term (overnight) funding charges on the premise that funding could be easily accessed and that any market disruption would be short-lived. This minimised the negative cost of carry, making it easy for banks to recoup the cost by simply adding a small spread (liquidity premium) to the funding cost that was charged to assets, as depicted in Equation 2.

$$\text{FTP} = \text{base rate} + \text{term liquidity premium} + \text{liquidity premium} \quad (2)$$

where FTP = funds transfer price, base rate = rate depicted from the swap curve corresponding to the asset's contractual/behavioural maturity or repricing term, whichever is less, term liquidity premium = spread between the swap curve and the bank's marginal cost of funds curve based on the contractual/behavioural maturity of the asset, and liquidity premium = cost of carrying liquidity cushion averaged over total assets of the bank.

Generally speaking, banks would not have believed this would create problems despite the fact that only assets were being charged because the spread charged to recoup the cost of carrying the liquidity cushion was so small. But there are some serious implications associated with this poor practice. First, it inherently assumes that all assets expose the bank to the same unexpected (contingent) liquidity risk. Second, it completely neglects the contingent liquidity risk embedded in liabilities, for example, deposit run-off during stress environments, and off-balance sheet activities such as draw-downs on lines of credit. Third, it makes no attempt to charge businesses based on their predicted liquidity usage during stress environments. Fourth, the attribution of charges is not granular enough to discourage businesses from writing or buying products that pose more contingent liquidity risk than others.

Although some banks that were surveyed did attempt to charge the negative cost of carry back to businesses on a predicted usage basis, the attribution was generally at a very high level. One problem with not having a granular charge is that it encourages businesses to deal in products that are not being charged for the contingent liquidity risk they actually present. The implication of distorting behaviour by not charging products for the risks they present was discussed in Section 2.

4.6 Towards better management of contingent liquidity risk

In one form or another, all banks that participated in the survey are enhancing the way they manage contingent liquidity risk. Many are incorporating a wider variety of scenarios as part of their stress-testing processes to account for different types of market disruptions that might occur. These are largely in conjunction with BCBS recommendations and include idiosyncratic and systemic funding shocks and a combination of the two. In another step forward, senior management are becoming more engaged with stress-testing results and using them as a basis for deriving the size of the liquidity cushion. The composition of assets in liquidity cushions is broadly improving, once again in line with BCBS recommendations. For example, many of the larger banks are now holding a larger proportion of cash and government securities than previously. This is most likely due to the development of the LCR.

Probably the most substantial enhancement that is occurring is the application of higher funding costs to liquid assets. Before the GFC, banks believed funding could be accessed almost immediately and always. But the recent market turmoil has demonstrated that funding markets can remain disrupted for a significant period of time. As such, banks are applying higher funding charges to assets held as part of the liquidity cushion on the premise that it could take longer than expected to generate liquidity when needed. The charges applied

depend on banks' assumptions surrounding the length and severity of potential market disruptions. If, for example, a bank assumes funding markets for a particular asset could remain stressed for two years, then the cost of holding that asset as additional liquidity should be based on the two-year term liquidity premium.

The move to apply higher funding costs to liquid assets is considered significant for banks because carrying a more costly liquidity cushion creates more profit drag. As described above, banks previously recouped this cost by charging all assets equally, a small liquidity premium via LTP. Following the same approach now, however, is likely to cause conjecture amongst business units and distort business unit behaviour given the larger costs. As such, this method is no longer feasible for banks. A better approach is to examine the contingent liquidity risk embedded in various business activities and to attribute charges based on their predicted, or expected, use of funding liquidity. Higher contingent liquidity charges should be applied to business activities that pose more threat to large and unexpected funding outflows. This process is depicted in Figure 10.

Figure 10

Toward better management of contingent liquidity risk

1. Identify contingent commitments that are likely to create unexpected funding demands.

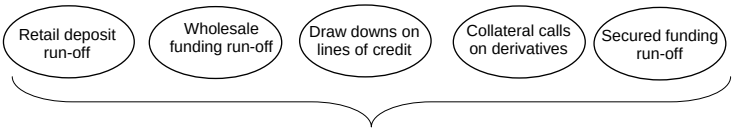

2. Perform stress tests under various scenarios to approximate the funding that might be required.

3. Net approximations from above against inflows generated, for example, through the sale of marketable securities to derive the size of the liquidity cushion.

Liquidity cushion

4. Calculate the cost of carry as the cost of funding liquid assets minus the return they generate. Ensure that appropriate haircuts and unsecured term funding charges have been applied to assets.

Cost of carry

5. Recoup cost of carry by charging a liquidity premium, at the most granular level, to the business unit, product or transaction that creates the need for the bank to carry such liquid assets.

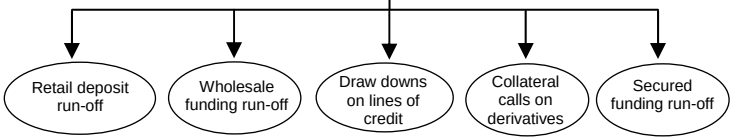

Some examples of pricing contingent liquidity risk are presented in the following section.

4.7 Example of pricing contingent liquidity risk

The uncertainty surrounding future cash-flow demands stemming from contingent commitments makes it particularly difficult for banks to assess and price contingent liquidity risk. This is one of the reasons why it was neglected prior to the GFC. Some of the products that received little attention but then warranted significant funding included: credit card loans and investments, trading positions and derivatives, revolving lines of credit, and liquidity lines.

The first step towards better management of contingent liquidity risk is not to address the question of how much should be charged but, rather, for banks to understand that all contingent commitments need to be charged. Once this is clear, then methods for pricing contingent liquidity risk can be refined and improved.

At the most basic level of what is considered to be better practice, all banks should be charging contingent commitments based on their likelihood of drawdown. For example, suppose a line of credit with a limit of $10 million has $4 million already drawn. The rate charged for contingent liquidity risk should be derived as:

$$\frac{\text{limit - drawn amount}}{\text{limit}} \times \text{likelihood of drawdown} \times \text{cost of funding liquidity cushion}$$

The likelihood of drawdown (sometimes referred to as a drawdown factor) should be assessed using behavioural modelling and should depend on factors such as customer drawdown history, credit rating of the customer, and other factors the bank deems important in making this prediction. In the example above, assume there is a 60 per cent chance the customer will draw on the remaining credit and that the cost of term funding assets in the liquidity cushion is 18 bps (depicted from the three-year term liquidity premium in Table 3). The

rate charged for the cost of contingent liquidity risk should be equal to:

$$\frac{(\$10\text{m} - \$4\text{m})}{\$10\text{m}} \times 0.6 \times 0.0018 = 0.000648\% \text{ or } 6.48 \text{ bps}$$

Multiplying this by the limit of $10 million on the line of credit yields a dollar charge of $6,480.

As explained earlier, prior to the crisis, banks applied short-term and often overnight funding charges to assets comprising their liquidity cushions on the belief that funding was abundant and permanently cheap. If the overnight funding rate was 0.5 bp and banks had applied the same approach as in this example, the dollar charge for the contingent liquidity risk would have been $1,800. However, banks did not follow this approach. Instead they averaged the cost of funding their liquidity cushions across all of their assets. By doing this it is likely that the lines of credit such as the one in this example did not even receive a charge for the cost of contingent liquidity risk. This would have encouraged business units to grant lines of credit and other contingent commitments.

A similar approach to that in the example above can be applied to other types of contingent commitments. For example, credit card accounts will generally have a proportion of the limit that is undrawn. With some behavioural analysis the likelihood of drawdown can be estimated and contingent funding liquidity costs can be attributed accordingly. More advanced banks might assess the behaviour of individual customers and assign a weighted probability of drawdown. All banks, however, are encouraged to at least examine and attribute charges based on the behaviour of a portfolio of like-contingent commitments.

5. Conclusion

The international survey identified some badly deficient LTP practices. The worst practices are summarised below.

- Management of the LTP process was poor. There were a lack of LTP policies, decentralised funding structures accompanied by weak internal risk controls and limits, inconsistent LTP regimes, off-line and manually adjusted LTP processes, a lack of trading book funding policies, as well as poor independent oversight by risk and financial control personnel.
- LMIS were often not advanced enough to incorporate the costs, benefits and risks of liquidity at a sufficiently granular level.
- Profit pools, which were used to determine bonuses to employees, were derived from unadjusted revenues, without any regard to the risks (liquidity and capital) taken to generate such profits.
- Probably the most striking example of poor practice was that some banks applied a zero charge for the cost of funding liquidity based on the premise that liquidity was a free good.
- Other banks applied a single pooled approach to LTP, whereby one average rate was used to charge users of funds and to credit providers of funds.
- Liquidity cushions were not linked to stress-testing outcomes, and scenario analyses were not severe enough to account for prolonged market-wide disruptions.
- Charges applied to fund liquid assets were often based on short-term rates, reflecting the belief that funding was abundant and would remain permanently cheap.
- To recoup the cost of carrying a liquidity cushion, most banks simply charged all assets an equal and small liquidity premium, which failed to account for the varying amounts of contingent liquidity risk embedded in different business activities.

Collectively, these poor LTP practices encouraged long-term illiquid asset creation and discouraged long-term stable liability creation, with obvious consequences.

By and large, banks have realised that many of their LTP practices were insufficient. As a result, banks are now working towards enhancing their LTP processes to ensure their business activities adequately account for the costs, benefits and risks of liquidity. To assist banks and supervisors throughout this process, below is a compilation of what is considered to be better LTP practices.

Governing LTP

- Banks should have an LTP policy that defines LTP, states the purpose of LTP and provides some principles and/or rules to ensure LTP achieves its intended purpose. The LTP policy should apply to all business units that are material users or providers of funding liquidity.
- LTP should be managed centrally, within group treasury or a subsidiary treasury, and applied consistently across the group. Wholesale funding should be confined to this function.
- Trading book funding policies should exist and, where appropriate, funding should be provided at the most granular level e.g., at the trading desk level as opposed to the trading book level. This requires treasury to have a full line of sight to individual business balance sheets. Banks are also encouraged to examine individual positions and apply higher funding charges to those that are more likely to become stale, or that present significantly greater amounts of funding liquidity risk. Finally, banks should have in place limits and adequate controls to curb over-trading behaviour.
- Oversight of the LTP process should be provided by independent risk and financial control personnel. Senior

management should also be involved in the LTP process. Meetings should be held regularly, and include various stakeholders such as ALCO, senior management and treasury functions to discuss changes in funding costs.

- To assist with performance assessments and decision-making, LMIS should be advanced enough to attribute, at a sufficiently granular level, all of the relevant costs, benefits and risks of liquidity to the appropriate business activities.
- Remuneration practices should be more sensitive to the risks taken to generate profits. For example, profit pools, which are used to derive short-term incentives such as bonuses, should be adjusted for the cost of liquidity through the LTP process.

Using LTP to manage on-balance sheet funding liquidity risk

- When applying LTP in practice, banks should by now have come to the realisation that liquidity is not a free good. Hence, employing a zero cost of funds approach to LTP is extremely poor practice and should not be tolerated by supervisors.
- Banks should instead be moving towards incorporating a matched-maturity marginal cost of funding approach to LTP. This approach is superior to any other established. First, it recognises the need to charge more for the cost of liquidity for assets that require funding for longer periods of time. Conversely, it recognises the need to credit more for the benefit of liquidity for liabilities that provide funding for longer periods of time. Second, the rate charged for the use and for the benefit of funds is based on banks' actual market costs of funds. This rate incorporates both idiosyncratic credit risk adjustments and market access premiums.

- Charges for the use of liquidity and credits for the benefit of liquidity should not be borne by central management (ie treasury), but should instead be attributed, at a sufficiently granular level, to the business activities using or providing liquidity.

Managing contingent liquidity risk

- The size of liquidity cushions should be derived from stress-testing outcomes and scenario analyses that at a minimum account for idiosyncratic and systemic scenarios, including prolonged market disruptions, and a combination of the two.
- Assets held as part of the liquidity cushion should be of the highest form of liquidity.
- Funding charges applied to assets in the liquidity cushion should not be based on short-term overnight rates but should instead be derived from longer-term rates to account for the possibility of longer than expected market disruptions.
- The cost of carrying a liquidity cushion is a cost of doing business and should thus be recouped from the business activities that require the bank to carry such liquid assets. The charge attributed to business activities should not be equal and based on the assumption that the contingent liquidity risk is the same for all types of business. Instead, individual business activities should be charged based on their expected, or predicted, use of contingent liquidity.

To conclude, one size does not fit all. Banks vary in size and complexity, and the LTP process employed should reflect these factors. While banks should, at least, consider all better practices promoted through this paper, only those that are appropriate and will most likely improve their own LTP process should be adopted.

Appendix:
LTP principles and recommendations

Working Group	Principles/Recommendations
Basel Committee on Banking Supervision (BCBS) Principle 4 of "*Principles for Sound Liquidity Risk Management and Supervision*" (September, 2008).[1]	"A bank should incorporate liquidity costs, benefits and risks in the internal pricing, performance measurement and new product approval process for all significant business activities (both on- and off-balance sheet), thereby aligning the risk-taking incentives of individual business lines with the liquidity risk exposures their activities create for the bank as a whole" (p 3).
European Commission (EC) Point 14 in Annex V of *"Directive 2009/111/EC of the European Parliament and of the Council of 16 September"*.[2]	"Robust strategies, policies, processes and systems shall exist for the identification, measurement, management and monitoring of liquidity risk over an appropriate set of time horizons, including intra-day, so as to ensure that credit institutions maintain adequate levels of liquidity cushions. Those strategies, policies, processes and systems shall be tailored to business lines, currencies and entities and shall include adequate allocation mechanisms of liquidity costs, benefits and risks" (L302/116).

[1] Available at http://www.bis.org/publ/bcbs144.pdf. [2] Available at http://eurlex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2009:302:0097:0119:EN:PDF.

Working Group	Principles/Recommendations
Committee for European Banking Supervisors (CEBS) Recommendation 2 in *"Second Part of CEBS' Technical Advice to the European Commission on Liquidity Risk Management"* (September, 2008).[3]	"Institutions should have in place an adequate internal mechanism – supported where appropriate by a transfer pricing mechanism – which provides appropriate incentives regarding the contribution to liquidity risk of the different business activities. This mechanism should incorporate all costs of liquidity (from short to long-term, including contingent risk)" (p 8).
The Institute for International Finance (IIF) Recommendation III.4 of *"Final Report of the IIF Committee on Market Best Practices: Principles of Conduct and Best Practice Recommendations"* (July, 2008).[4]	"Firms should ensure that they have in place effective internal transfer pricing policies to reflect implied or incurred actual or potential costs related to reasonably anticipated liquidity demands from both on- and off-balance sheet business. Transfer pricing should take closely into account the liquidity of relevant underlying assets; the structure of underlying liabilities, and any legal or reasonably anticipated reputational contingent liquidity risk exposures. Transfer pricing should be designed to ensure that lines of business within the firm that create liquidity exposures are proportionately charged for the cost to the firm of maintaining corresponding prudent liquidity positions" (p 56).

[3] Available at http://www.eba.europa.eu/getdoc/bcadd664-d06b-42bb-b6d5-67c8ff48d11d/20081809CEBS_2008_147_(Advice-on-liquidity_2nd-par.aspx. [4] Available at http://www.iif.com/press/press+releases+2008/press+75.php.

Working Group	Principles/Recommendations
Counterparty Risk Management Policy Group III (CRMPGIII) Recommendation IV-17 of "*Containing Systemic Risk: The Road to Reform*" (August, 2008).[5]	"The Policy Group recommends that all large integrated financial intermediaries incorporate appropriate pricing-based incentives for the full spectrum of their funding activities. This includes a funds transfer pricing policy that assigns the cost of funding to businesses that use funding and credits the benefits of funding to businesses that provide it. This must encompass both on- and off-balance sheet activities (for example, contingent funding), as well as potential funding needs related to actions that might be taken to preserve the institution's reputation. The funds transfer pricing process should be informed by stress testing efforts that identify potential vulnerabilities and assign the related costs to the businesses that create them. The methodology should provide direct economic incentives factoring in the related liquidity value of assets and behavioral patterns of liabilities. The costs and benefits identified should be assigned to specific businesses and, under all circumstances, used in evaluating the businesses' performance" (p 30).

[5] Available at http://www.crmpolicygroup.org/.