Financial Stability Paper No. 6 – August 2009

A risk-based methodology for payment systems oversight

Ben Norman, Peter Brierley, Peter Gibbard, Andrew Mason and Andrew Meldrum

BANK OF ENGLAND

Financial Stability Paper No. 6 – August 2009

A risk-based methodology for payment systems oversight

Ben Norman,(1) Peter Brierley, Peter Gibbard, Andrew Mason and Andrew Meldrum

We are very grateful for helpful suggestions and comments from colleagues at the Bank of England, including Ian Bond, Paul Chilcott, Stephen Collins, Stephen Denby, Jack Garrett-Jones, Cathy Hayes, Steve Miller, Julian Oliver, Vicky Saporta, Chris Shadforth, Gabriel Sterne and Matt Willison.

Financial Stability, Bank of England, Threadneedle Street, London, EC2R 8AH

(1) Corresponding author: ben.norman@bankofengland.co.uk.

ISSN 1754–4262

Contents

A risk-based methodology for payment systems oversight

Ben Norman, Peter Brierley, Peter Gibbard, Andrew Mason and Andrew Meldrum

The Bank of England has developed a risk-based methodology to support its oversight of payment systems. The methodology provides more precise estimates of risks in payment systems than previously available. Because it is consistent and systematic in its application, the methodology assists the Bank in focusing its attention and resources — the intensity of oversight — where the level of risk is estimated to be greatest. This article provides an overview of the framework.

Members of the Payment Systems Oversight team in the Bank of England's Financial Stability Directorate developed a risk-based methodology in 2005 to assist its oversight of payment systems in the United Kingdom. This article describes how the methodology has been designed and applied to date. It also provides an overview of the broader framework in place for assessment and monitoring of payment systems risks. Two key aspects of the risk-based methodology are elaborated in an annex. The article finishes by describing in general terms how outputs from the risk-based methodology can be used by the Bank's Oversight team and management to support targeted risk-reducing actions.

The context for a risk-based methodology for oversight

The Bank's oversight responsibilities were first formalised in 1997 in the Tripartite Memorandum of Understanding (MoU) with HM Treasury and the Financial Services Authority (FSA). The MoU, which was revised and updated in March 2006, assigns to the Bank a general responsibility for oversight of UK payment systems. In practice, oversight resources are focused on those payment systems which have been judged to pose the greatest risk to financial stability.(1) As part of the Banking Act 2009, the Bank's role in oversight is being put on a statutory basis, with HM Treasury responsible for recognising those payment systems that the Bank will formally oversee.(2) Individual payment systems themselves remain responsible for the identification, assessment and, crucially, mitigation of risks — a responsibility which the Bank's oversight is not intended to dilute. **Figure 1** provides a schematic overview of the Bank's oversight framework.

The oversight risk methodology described in this article provides the basis for a more precise approach to assessing risks in payment systems — resulting in more cardinally ranked risk estimates — than was previously the case.(3) Risks in payment systems can never be precisely quantified, and the implementation of the methodology described here does not represent the adoption of a mechanical or 'model-driven' form of oversight. This methodology is, though, an important adjunct to the Bank's qualitative assessment of risks, which includes assessment against a number of principles for payment systems that incorporate the international 'Core Principles'.(4)

Among other things, the methodology described in this paper helps facilitate consistency of oversight, because it can be applied systematically across different payment systems. This is particularly useful in a UK context, where a number of wholesale and retail payment systems co-exist. A consistent approach to oversight helps to ensure that the Bank focuses its risk mitigation actions on those systems where risks have been assessed to be significant. It also enables the identification, across systems, of the types of risk on which the Bank should focus its oversight resources. And it allows for risks that would

(1) A fuller explanation of the Bank's role in the oversight of UK payment systems is given in Bank of England (2005), *Payment Systems Oversight Report 2004*, chapters 1 and 2 in particular. See also Haldane, A G and Latter, E (2005), 'The role of central banks in payment systems oversight', *Bank of England Quarterly Bulletin*, Spring. A further, forthcoming publication, Manning, M, Nier, E and Schanz, J (eds) (expected 2009), *The Economics of Large-value Payments and Settlement*, contains references to a wide range of literature relevant to the research and policy work that underpins the Bank's oversight function.
(2) Office of Public Sector Information (2009), *Banking Act*, Part 5.
(3) Several Bank of England *Payment Systems Oversight Reports* have given an overview of the Oversight Risk Framework. This article is the first, more detailed exposition of it.
(4) Bank for International Settlements, Committee on Payment and Settlement Systems (2001), *Core Principles for Systemically Important Payment Systems*.

Figure 1 Schematic overview of the oversight risk framework

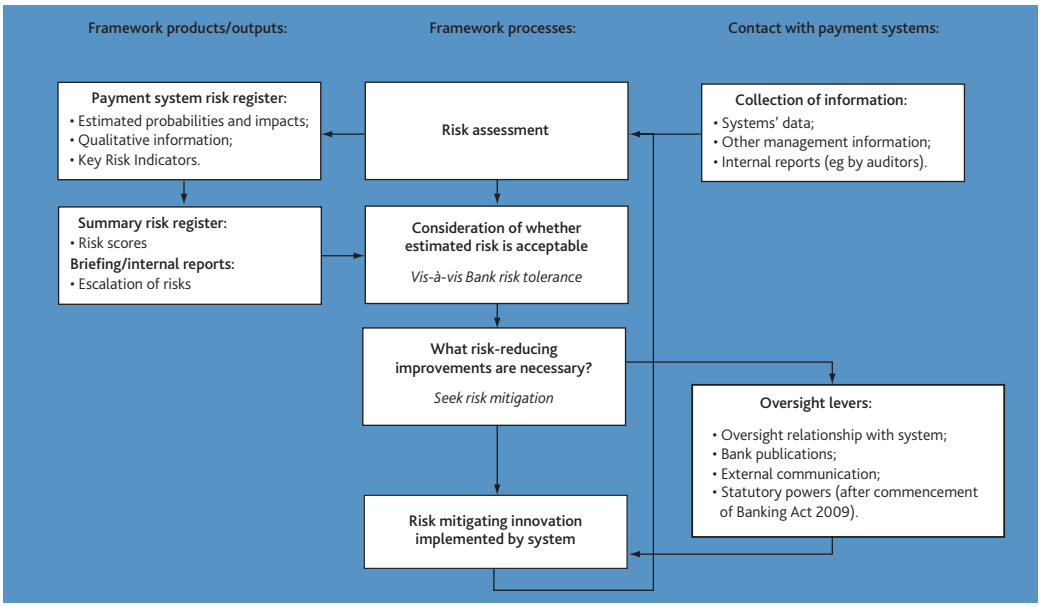

crystallise simultaneously across several systems to be reflected in the risk assessments.

This is by no means the first formal risk framework to be developed by a financial regulator. In the United States, 'CAMELS' ratings are used by a number of financial regulators (notably the Federal Reserve Board, Federal Deposit Insurance Corporation and Office of the Comptroller of the Currency) to provide a means of drawing together a substantial evidence base to support their supervisory work.(1) In the United Kingdom, the FSA has implemented its ARROW II framework to allow its staff to perform a detailed risk assessment, in order to identify the main risks to achieving its regulatory objectives.(2) Neither CAMELS nor ARROW II is specifically focused on assessing risks in payment systems.

The hitherto most developed framework specifically for assessing risk in payment systems and associated financial infrastructure is perhaps that developed by Citigroup's Payment Systems Risk Management function.(3) However, Citigroup's framework was not developed for regulatory purposes.

A key difference between these risk-based frameworks and the Bank of England's oversight risk methodology is in its approach to deriving estimates of risk. While the Bank of England's methodology still embodies a number of assumptions (in order to make it operational), it goes beyond a simple ordinal scale (eg 1 to 5) or arbitrary weightings (eg multiples of 5% or 10%) that tend to characterise the way that probabilities and impacts are assessed in these other frameworks. By design, the Bank's methodology provides more precise estimates of both the probability (expressed as a 'one in *x* years' event) and the impact (expressed in monetary units) of risks in payment systems.(4) Even so, the assumptions made in deriving these risk estimates are such that the methodology is still best regarded as providing a (more refined) ordinal assessment of the relative risks.

Risks captured in the Bank of England's oversight methodology

The Bank's risk-based methodology is organised around the objective of the Bank's oversight — to assess and, if necessary, promote the mitigation of those risks within UK payment systems that could have adverse effects on the financial sector and the wider economy. In essence, the Bank seeks through its oversight to reduce systemic risks that could arise from and be propagated by payment systems. At the same time, the Bank recognises that designing and operating a payment system to minimise systemic risks would be counterproductive if the system thereby became so inefficient or impractical to use that payment traffic migrated to less safe alternatives.

Since the global financial crisis started in Summer 2007, the financial infrastructure has remained resilient in the face of both significant credit events and operational challenges (examples of the latter include processing record volumes of trades whose associated payments have needed to be cleared and settled). Operational problems that temporarily prevented a member of a payment system from making

(1) Federal Reserve Bank of San Francisco (1999), 'Using CAMELS Ratings to Monitor Bank Conditions', *Economic Letter 99–19*, June.
(2) Financial Services Authority (2006), *The FSA's risk assessment framework*.
(3) Details about the Citigroup risk framework are proprietary. A flavour of the issues that form the basis of the Citigroup framework can be gleaned from the New York Payments Risk Committee's (2007) Report on *Financial Market Infrastructure Risk*.
(4) The monetary impact measure encompasses estimated losses both to financial institutions and to end-users of payment systems.

Table A Examples of events that have crystallised in recent years in UK payment systems

Risk type/detailed risk category	Event	Source
Settlement risk: settlement member insolvency/illiquidity	'On 15 September 2008, Lehman Brothers…was placed in administration. …Lehman Brothers' default occurred after some intraday funding via the self-collateralising repo mechanism had been undertaken by [CREST] settlement banks. This demonstrated the importance of settlement banks ensuring adequate liquidity management planning for a client default.'	*Payment Systems Oversight Report 2008* (page 12)
Settlement risk: settlement member operational problems	'One particular [CHAPS] member had an outage that lasted most of the day on 3 January 2008, owing to an extremely rare software failure. … Communication between members meant that they were able to stop or delay sending payments to the stricken bank, so that it did not become a liquidity sink… .'	*Payment Systems Oversight Report 2008* (page 10)
Operational risk: disasters/terrorist attack	'Following the London bombings on 7 July 2005, LCH.Clearnet Ltd was required to evacuate its head office and operate from its secondary office site.'	*Payment Systems Oversight Report 2005* (page 29)
Operational risk: systems/network failures	'On 29 August [2006]…a software bug affecting communication between the SSE [Single Settlement Engine] and the CREST system resulted in a three hour outage. As a result, CHAPS processing had to be extended, sterling deadlines were pushed back to around 19:15, and major banks were only able to close their systems and process client accounts after 20:00.'	*Payment Systems Oversight Report 2006* (page 16)
Operational risk: systems/network failures	'…on 12 February [2007]…connectivity problems…affected the RTGS [Real Time Gross Settlement] infrastructure, preventing CHAPS members from submitting settlement instructions to RTGS via SWIFT for around six hours. This was caused by localised problems affecting the software supporting RTGS… .'	*Payment Systems Oversight Report 2007* (page 10)
Operational risk: systems/network failures	'…on 20 and 21 August 2008…several members [of the Faster Payments Service, FPS] started to have problems accessing the central infrastructure due to a problem with a security certificate authentication server maintained by BT [British Telecom]. The initial fix exacerbated the problem, which was resolved on 21 August 2008. The LINK system experienced similar problems, as it shares the secure communications network with FPS.'	*Payment Systems Oversight Report 2008* (page 23)
Operational risk: utilities failure (and systems/network failures)	'…there was a double failure of the firewalls surrounding the RTGS processor on 7 July 2008. When the main firewall at the secondary site…was unable to start due to a power failure, the backup firewall should have taken over, but it was unable to do so. …RTGS was unavailable in total for over 200 minutes … [and] necessitate[d] extensions in two other currencies' payment systems… .'	*Payment Systems Oversight Report 2008* (page 11 and page 30)
Operational risk: external threat to networks/theft	'…there was an incident in September 2008 where some Bacs components were stolen from a BT exchange. This caused delays to Bacs processing… .'	*Payment Systems Oversight Report 2008* (page 21)

payments might, at a time when market participants were particularly nervous, have been misinterpreted as a signal of liquidity and/or solvency problems. Similarly, if the technical capacity of the payment system itself were exhausted, preventing further payments from being settled on high-volume days, exposures would be likely to build up until the problem was fixed, at just the time when market participants are most concerned to contain their exposures. So a thorough assessment of risks is important in informing risk mitigation priorities. In this way, the Bank's oversight contributes to a more stable financial system, by reducing risks in payment systems, which might trigger, or amplify the impact of, a financial crisis.

Channels through which payment systems risks are propagated

To operationalise the assessment of risks in UK payment systems, the methodology specifies two channels through which payment systems risks may have an adverse impact upon the financial sector and the wider economy:

- contagion, whereby the financial or operational difficulties of one member of the payment system are transmitted through it to one (or more) other member(s) of the system; and

- disruption to transactions, whereby the financial or operational difficulties of the operator(s) of payment systems have so-called 'real economy' effects, by delaying, or even preventing, payments being made by financial institutions, businesses and/or consumers, or by requiring such payments to be made via materially less efficient/more risky methods than payers would freely choose.

The annex to this article explains in more detail how risks via these two channels are assessed.

Risk types, events and the risk register

To organise the landscape of risks in payment systems into a format which lends itself to consistent assessment and monitoring, the Bank's methodology assigns risks to three distinct risk types:

- 'Settlement risk', which is the risk that a participant in a payment system cannot or does not meet its financial obligations when they fall due, or that another institution facilitating settlement of those obligations — eg the settlement agent — becomes insolvent or suffers an operational outage such that settlement is impeded.(1) For example, a participant in a payment system may fail to meet its financial obligation as required by the rules of the system, but then meet that obligation at some later date, giving rise to the need for liquidity; or the participant may become insolvent and be unable to meet its obligation at all, creating an outright loss.

(1) For many payment systems, the central bank is the settlement agent, in which case insolvency is not relevant to the assessment process for this aspect of settlement risk.

- 'Business risk', which is the risk that a payment system or one of its components — eg an infrastructure provider[1] — is no longer financially viable and so is not able to continue to operate and enters administration, which may disrupt or terminate its capacity as a business to process payments.

- 'Operational risk', which is the risk that a payment system operator or an infrastructure provider to the system is operationally unable to process or settle payments as intended. For instance, participants or users might face losses or material inconvenience from a failure of a payment system's software, hardware or internal processes; from internal fraud; or from external events (eg a major power outage).

The scope of these three main risk types (and the more detailed categories of risks within each) have been defined to ensure, as far as possible, that they are all-encompassing and yet do not overlap.

Within each risk type various (more granular) trigger events can give rise to settlement, business or operational risks in a payment system. It is on the basis of its analysis of such detailed events that the Bank seeks to assess the probabilities and impacts of the various risks. **Table A** gives a (non exhaustive) list of events that have occurred in UK payment systems in recent years, which inform the Bank's assessment of settlement and operational risks.

Risk assessment and monitoring

Risk assessments are carried out against the same list of detailed risks for each payment system. A typical risk register is shown in **Figure 2**. In essence, the assessment process both reviews observed events and analyses unobserved vulnerabilities, and attaches estimated probabilities to each risk and its associated estimated impacts. These estimates represent the main components of the methodology.

So how are probabilities and impacts actually estimated? To start with, the overseer seeks to understand *how* risks can crystallise in each system and the nature of any impact they might have. This requires the overseer first to develop a sound knowledge of the key processes which support the day-to-day operation of the payment system in question. This assessment also takes into account mitigants (such as legal agreements, or operational control procedures) that are in place to control the risks.

The next step is to determine through which channel(s) — contagion and/or disruption to transactions — each particular risk may affect the financial sector and the wider economy. Certain aspects of settlement risk have the potential to cause both contagion and disruption to transactions; whereas business and operational risks cause impacts only via disruption to transactions. More specifically, for example, failure of a payment system's IT hardware would be judged by the overseer to have an impact only through disruption to transactions, whereas member default might be judged also to affect other members of the payment system via contagion (unless other members are fully protected, for example by collateralising the exposures).

Greater precision in the Bank's risk assessment methodology is then achieved by using the available data to estimate *how likely* the risk might be and *how much* of an impact it might have. In theory, the estimation process should begin with the conditional loss distribution of a risk — that is, the distribution of possible impacts and associated probabilities which are conditional on the event occurring. Put another way, there is a probability of some event occurring (eg hardware failure) and, conditional on that occurrence, there is a probability that the event will have a certain impact (eg a probability of *p1* that hardware failure will result in closure of the payment system for a period of time *t1*, a probability of *p2* that the closure lasts a period of time *t2*, etc). For many events, it is reasonable to assume that the mode of the conditional loss distribution is relatively small; but there is likely to be a long tail of larger impacts which could occur, usually with decreasing probabilities. **Figure 3** shows a stylised conditional loss distribution of this form.

In practice, as described in the annex, there are sufficient data to estimate a conditional loss distribution only for some aspects of settlement risk. For business and operational risks, the methodology instead makes the simplifying assumption that there are just two different outcomes if such a risk crystallises; a 'typical' (modal) and an 'extreme' (tail) scenario. Over time, the estimation method for each of the risk types may improve further, as more data become available.

The initial qualitative assessment becomes particularly important when considering operational risks. The vulnerability of a payment system to operational risks, and the effectiveness of associated controls, are summarised in a simple ordinal score. This assessment is then used to support the probability and impact estimates of various operational risks. In this case, the qualitative assessment helps to compensate for the rarity of material operational events, and the consequent lack of useful data on which to base the risk assessment.[2]

(1) Examples of infrastructure providers to payment systems include: SWIFT, which provides secure messaging services to financial institutions and market infrastructures globally; and VocaLink Ltd, which provides the infrastructure to some of the United Kingdom's retail payment systems.

(2) On operational risk estimates, see also De Fontnouvelle, P, Jordan, J and Rosengren, E (2006), 'Implications of alternative operational risk modelling techniques', in Carey, M and Stulz, R (eds) (2006), *The Risks of Financial Institutions*.

Figure 2 A stylised payment system risk register

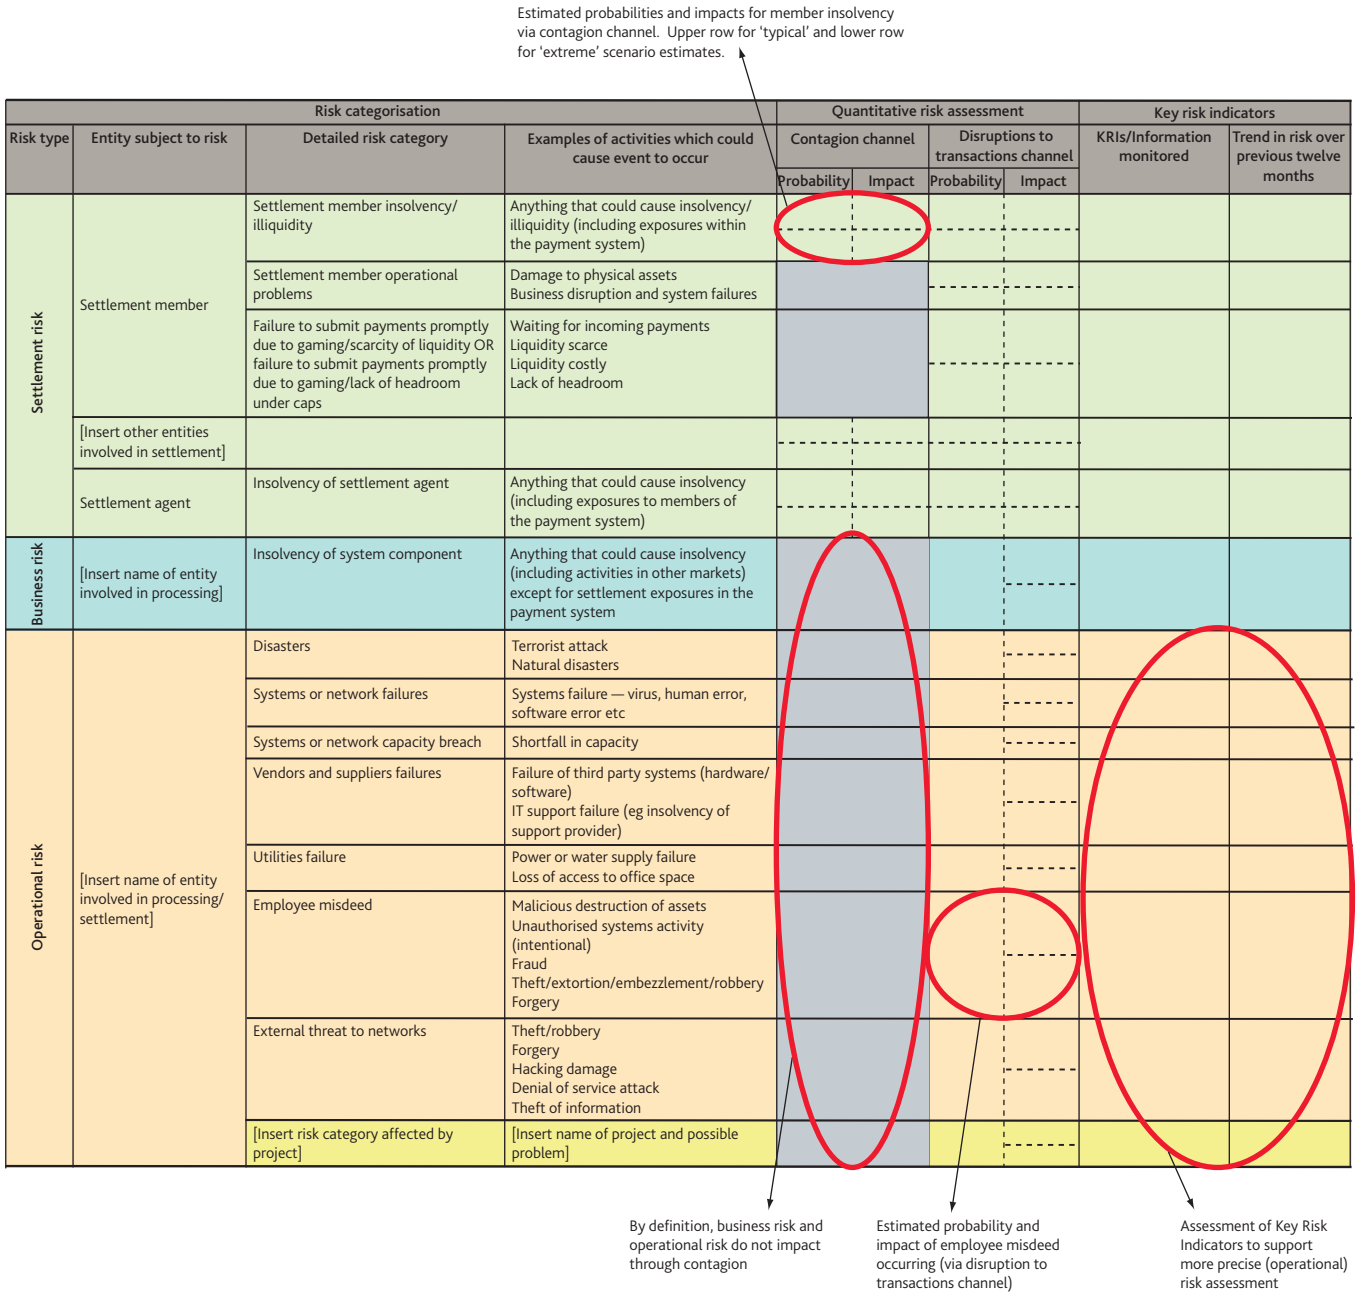

Risk categorisation				Quantitative risk assessment				Key risk indicators	
Risk type	Entity subject to risk	Detailed risk category	Examples of activities which could cause event to occur	Contagion channel		Disruptions to transactions channel		KRIs/Information monitored	Trend in risk over previous twelve months
				Probability	Impact	Probability	Impact		
Settlement risk	Settlement member	Settlement member insolvency/ illiquidity	Anything that could cause insolvency/ illiquidity (including exposures within the payment system)						
		Settlement member operational problems	Damage to physical assets Business disruption and system failures						
		Failure to submit payments promptly due to gaming/scarcity of liquidity OR failure to submit payments promptly due to gaming/lack of headroom under caps	Waiting for incoming payments Liquidity scarce Liquidity costly Lack of headroom						
	[Insert other entities involved in settlement]								
	Settlement agent	Insolvency of settlement agent	Anything that could cause insolvency (including exposures to members of the payment system)						
Business risk	[Insert name of entity involved in processing]	Insolvency of system component	Anything that could cause insolvency (including activities in other markets) except for settlement exposures in the payment system						
Operational risk	[Insert name of entity involved in processing/ settlement]	Disasters	Terrorist attack Natural disasters						
		Systems or network failures	Systems failure — virus, human error, software error etc						
		Systems or network capacity breach	Shortfall in capacity						
		Vendors and suppliers failures	Failure of third party systems (hardware/ software) IT support failure (eg insolvency of support provider)						
		Utilities failure	Power or water supply failure Loss of access to office space						
		Employee misdeed	Malicious destruction of assets Unauthorised systems activity (intentional) Fraud Theft/extortion/embezzlement/robbery Forgery						
		External threat to networks	Theft/robbery Forgery Hacking damage Denial of service attack Theft of information						
		[Insert risk category affected by project]	[Insert name of project and possible problem]						

Figure 3 A stylised example of a conditional loss distribution

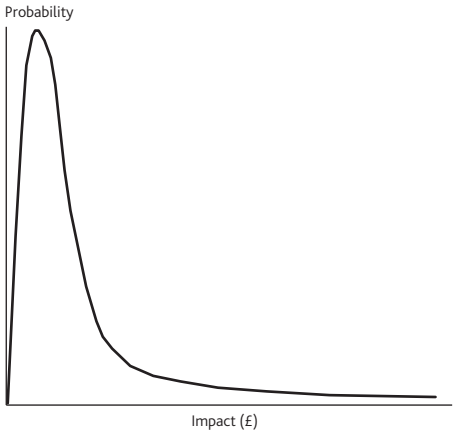

The 'typical' and 'extreme' scenarios are chosen by the individual overseer, although a risk framework co-ordinator compares different overseers' assumptions in order to maintain a broad consistency of approach. The overseer will typically document a number of assumptions about how a particular scenario might play out, while also making judgements about the robustness of relevant controls and capabilities the payment system has in place. The overseer also draws further on their qualitative understanding of the payment system, to determine how a particular risk would affect the members of the payment system and the real economy. For example, a typical scenario for an operational risk such as hardware failure might be specified as resulting in the temporary intraday closure of the payment system (eg a few hours), whereas the extreme scenario might be a closure

Figure 4 A stylised summary risk register

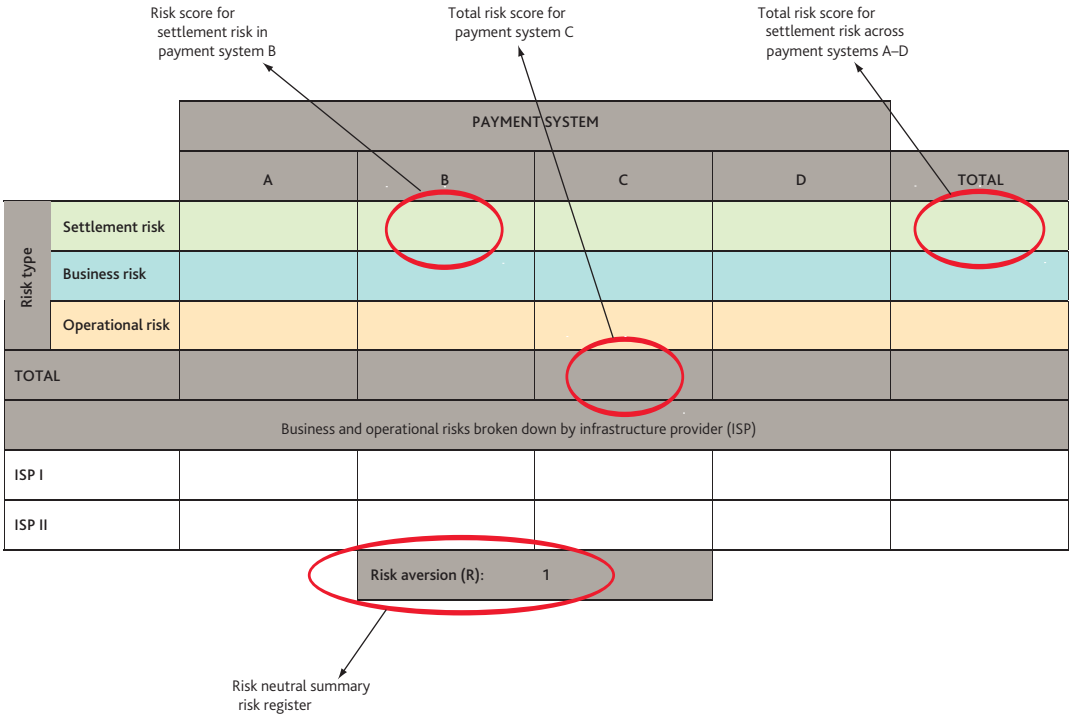

lasting a number of days, as a result of (for instance) the failure of recovery processes put in place by the system. Both scenarios could cause disruption to transactions, though the latter would be much more disruptive (yet less likely to occur) than the former.

Probabilities estimated for each risk are recorded in the risk register as frequency ranges. These frequency ranges are spread across a spectrum from 'more frequent than annually' through to 'less frequent than once in every 200 years'.(1) Impacts are estimated in terms of a monetary cost for each scenario ('typical' and 'extreme') attached to each risk, and are also presented as ranges (£1 million–£20 million, £21 million–£50 million, and so on).

Since estimated probabilities and impacts, and supporting qualitative information, represent an assessment of risks only at a particular time, the Bank is careful to monitor existing risks and analyse new ones that may need to be reflected in future assessments. Such monitoring is facilitated by consideration of key risk indicators (KRIs). These can be high frequency data series which give an indication of trends in probabilities and/or impacts, based on observed events (eg operational performance statistics showing how many incidents of different severity have been observed in a particular system during a particular time period). They can also take the form of qualitative information which highlights changes in a payment system's vulnerability to a specific risk and the quality of associated controls (eg indications from a system's audit reports that risk controls have improved or deteriorated).

Monitoring such KRIs is an important part of the continual work of an overseer, and identification of a material change in a KRI can trigger a reassessment of the risk in question. The risks identified and included within each payment system risk register are formally reviewed by the Bank on an annual basis, with quarterly updates to reflect any significant developments.

Interpreting the outputs

The methodology of risk assessment described above is used to populate the Bank's risk register for each payment system. The estimates contained in individual risk registers are aggregated into management information which is used to help identify where mitigating action should be sought. Through this process, the risk estimates also help the Bank to determine how to allocate its oversight resources.

The principal piece of management information is the summary risk register (**Figure 4**). This shows aggregated risk estimates for each of settlement, business and operational risk in each of the overseen payment systems, based on the probabilities and impacts estimated for the detailed risk categories within individual payment system risk registers. The row totals of the summary risk register represent the estimated overall risk for each risk type, whereas the column totals represent the estimated overall risk for each payment

(1) For risks with extremely low estimated probabilities, there is a limit to the meaningfulness of the more precise approach to risk assessment. Where such risks have a very high estimated impact, the Bank would, where appropriate, still seek improvement of controls in order to mitigate the risk — but it would do so on a largely qualitative analysis of the risks.

system. The summary risk register also breaks down the latter to show the contribution of different infrastructure providers to the business and operational risk estimates. If it transpired that risks were concentrated in a particular infrastructure provider, this would show up in the summary risk register and could — where necessary — prompt the overseer to press the payment system(s) concerned to deal with the risk issues with that infrastructure provider.

The aggregation methodology used to derive the estimates populating the summary risk register starts off with the simplifying assumption that risks occur independently across all payment systems. For example, it assumes that a network failure in one system does not also result in the same failure in another system. By calculating the products of the impact and probability estimates for each of the risks in the detailed risk register, and summing them as appropriate, a 'neutral' (or 'actuarially fair') set of risk estimates is calculated for the summary risk register.

A number of risks in different payment systems, however, are unlikely to be completely independent of one another. They may in fact be partially correlated (eg through overlapping membership of payment systems); or they may be perfectly correlated (eg through payment systems sharing common infrastructure); or their impacts may be inseparable (eg if different systems were sharing the same collateral pool). In any of these cases, ignoring such interlinkages would result in an underestimate of aggregate risks in payment systems. So the risk framework co-ordinator brings together overseers of those systems where risks are not independent of each other, and ensures that risk estimates in the detailed risk registers are consistently applied. For instance, for perfectly correlated risks, the overseers of the systems in question assign the same probability estimate to the risk in question.

Even this might not be sufficient to reflect fully the interdependencies across different payment systems. For example, an overseer assessing the impact of a particular risk might, in isolation, make assumptions about the likely availability of another system as an (imperfect) substitute, should the risk crystallise in the first system. If, because of interdependencies, the other system is not in fact viable as a substitute (or to a lesser degree than assumed), then the impact of this risk would — if not mitigated — be greater than that estimated. Estimating aggregate risks (across systems) in the context of such interdependencies is complex. So the aggregation method in the oversight risk framework proxies for such interdependencies by weighting the individually assessed impact estimates with an index. The index is calibrated, such that a value of one replicates the summary risk register estimates in a risk 'neutral' setting. By setting this so-called 'risk aversion index' greater than one, the estimates in the summary risk register increase the emphasis that is placed on higher impact risks — which, among other things, will increase the risk estimates where interdependencies exist.

Using all these risk estimates, both in detail and in summary form, and with both risk neutral and risk averse settings, the Bank's framework provides a rich set of management information which allows the Bank to monitor the risk mitigation efforts of system operators in a structured fashion, and provides the basis for the Bank's oversight dialogue with individual payment systems. Any such efforts which prove effective are captured in revised estimates of probabilities and impacts within the payment system's risk register. All other things being equal, this results in lower risk estimates in the summary risk register and, ultimately, potentially some reallocation of oversight resources.

Conclusions

The Bank of England has developed a risk-based methodology for its oversight of payment systems in the United Kingdom. The methodology enables risks in payment systems to be assessed on a more precise basis. Consistency of application means the methodology helps the Bank to judge the relative intensity of its oversight from system to system in a risk-based fashion.

Incorporating the new methodology into the Bank's oversight of payment systems has generally served to confirm the Bank's oversight priorities. In particular, after the focus on settlement risk issues during the 1990s (when RTGS was introduced), the initial results from the framework have underlined the importance of also dealing with operational risks in payment systems.

Since it was first developed in 2005, the methodology has been continuously evolving (and is likely to do so in the future) as refinements are made. The authors would welcome feedback from academics and practitioners on the technical approach set out in this paper.

Annex: Estimating risks in more detail

Notwithstanding that the methodology is still evolving, this annex provides further detail concerning the estimation of probabilities and impacts for settlement, business and operational risks. It does so by explaining some key aspects of the approach taken to estimating the risks via each of the two channels that are considered in the oversight risk framework: the disruption to transactions and contagion channels. An integral element of estimating the impact arising from any event that causes a disruption to transactions is the calculation of the estimated 'cost of a one-day outage' (CODO) for each payment system. This is presented first in the annex. Aspects of settlement risk also have the potential to cause losses through the contagion channel. This annex then sets out how, in broad terms, the Bank's overseers go about calculating such settlement risk estimates.

Estimating impacts via the disruption to transactions channel

Anything that causes payments not to be made in the way normally intended has the potential to cause inconvenience or financial losses to members of payment systems and end-users. Such disruption to transactions can arise from a crystallisation of any of the risks in the risk register. The Bank's approach to estimating the losses caused by this channel entails estimating what the losses would be if any given payment system were operationally unavailable for a period of 24 hours: the CODO.

When a risk crystallises and disrupts transactions, payments are classified as being affected in one of three ways. They are:

- delayed — until the outage affecting the system through which the payments are (to be) processed has been resolved; or
- denied — such that the intended payments never actually take place; or
- substituted — whereby the payments are made in a (reasonably) timely fashion, but only by being processed using a different (and potentially more expensive or risky) method or system.

Different cost assumptions can be attached to these three classifications based on an analysis of the individual payment type. For example, any payee that receives a payment with a delay may (pending receipt) not have the funds available in order to make further payments itself, which may in turn mean that its trading opportunities are lost (or inferior to what they would have been); or a retail firm that has its means of receiving payments denied may be unable to make a sale that day and will forgo the profits it might have made; or a bank that decides to re-route payments via a different, functioning payment system may incur greater staff costs, especially if the alternative procedures are unfamiliar to the staff concerned because they are not often invoked; and so on. **Table B** provides a fuller description of the different types of cost that are considered in building the CODO estimates.

The CODO approach requires certain parameters to be assumed about each of these costs. For instance, the search costs following a denied payment could be approximated by an assumption that, on average, a payee would spend *x* hours arranging an alternative transaction at a cost of *x* times the average hourly wage. Or some of the re-routing costs for a payment that is made via a substitute system could be approximated by an assumption that, on average, a bank would need to pay additional overtime to a number of its back office staff, incurring a larger wage bill. Individually, the costs to payees/banks could be expected to be relatively small in absolute terms — although they could still be non-trivial (for instance, to businesses that operate on tight margins). However, given the volumes of payments that are normally settled by the main payment systems in the United Kingdom, the aggregate costs of delay, denial or substitution are potentially significant.

Table B Costs arising from disruption to transactions

Effect	Consequent cost types	Examples
Delayed	Opportunity cost to the payee	Interest the payee could have earned on the funds had the payment not been delayed.
	Liquidity constraint faced by the payee	If the payee were relying on (delayed) incoming funds to make outgoing payments.
	A decline in the value of the underlying transaction	If the goods against which payment is being made can lose their value relatively quickly (eg if they are perishable).
	Additional risk	Credit risk, if one of the counterparties to the payment fails.
Denied	Lost profit	If the payment would have realised a profit to the payee, eg from selling goods/services to the payer, which are then not purchased.
	Search costs	Arising from a counterparty needing to start afresh in order to achieve the effect of the denied transaction.
	Additional risk	If the utility of the transaction denied was to hedge an exposure in a financial market.
Substitution	Net cost of using the substitute method or system compared to the intended system / additional risk	Additional fees and/or increased financial risks arising from using alternative method/system.
	Cost of switching to (and back from) the substitute system / additional risk	Having to use non-standard procedures, which takes more staff time and potentially increases operational risks.

Broadly speaking, the following steps are taken in order to derive a CODO estimate for each payment system:

- establish recent typical volumes and values of the payments handled by the system, disaggregated by the various types of transaction it processes;
- choose a particular 24 hour period for the outage scenario (eg for a worst-case estimate, the 24 hour period that is thought to cause maximum disruption);
- decide what proportions of each of the transaction types will be affected by this outage;
- consider whether these transactions (by each type) will be delayed, denied and/or substituted;
- determine which types of costs affect which transaction;
- quantify these costs using agreed parameters and charges; and
- sum the different costs that are assumed across the proportion of the different payment types affected.

Many of the risks and outage scenarios considered as part of the CODO approach have never materialised in reality. Therefore a significant degree of judgement is used to estimate the effects. The calculations are not an exact reflection of what would happen in reality. Nevertheless, the CODO approach gives a more refined estimate of the potential impacts of a number of risks in payment systems than previously available.

One reason for standardising the CODO estimate for each payment system on the basis of a 24 hour outage is that this helps to make the estimation approach consistent across different payment systems. But this 24 hour metric is not a straight-jacket to estimation in the oversight risk framework. The CODO estimates can be scaled down by the overseer, if it is determined that the typical or extreme scenarios captured in the oversight risk framework would be expected to last for less than 24 hours. On the other hand, it is not inconceivable that, especially in an extreme scenario, the payment system could be out of operation for longer than 24 hours.(1) If the overseer considers this to be a plausible (if extreme) scenario, then there is discretion to scale up the CODO estimate.

The scalar for adjusting the CODO estimate up or down could be linear (for simplicity), or it could be based on an exponential relationship between the duration of an outage and its impact. Typically, for very short outages of, say, a few minutes, the costs are negligible, since the delay in making payments a few minutes later than intended is usually of no consequence to the payment system participants and their end-users. However, in particular in payment systems that settle in real-time, an outage of a few hours can result in more substantial costs. And an outage that is not resolved before the end of the operational day arguably represents a step-change in impact, since it precludes any payments not yet made from being settled that day. (Often the contractual requirement is for a payment to settle at some point on a particular business day.) On the other hand, the impact of an outage lasting several days could be less than the linear multiple of the CODO estimate — on the basis that, if payment system participants are anticipating a longer such outage, then they are likely to find alternatives (for making the payments) which might not be available within 24 hours.

Estimating the risk of contagion

Every business day, the Bank settles payments across the settlement accounts it operates for members of the CHAPS payment system. For those payments settled on a gross basis in real-time ('real-time gross settlement' (RTGS)), the credit risk that payment systems can give rise to is mitigated. For those payments that are settled on a deferred basis, say once a day, with obligations between members of the payment system added up and offset such that only the net position of each member is settled ('deferred net settlement' (DNS)), interbank credit risk can exist from the moment the payment instruction is irrevocably processed to the moment the Bank settles the net obligations with finality.(2)

When settling the net obligations of the members of a DNS system, the Bank obtains each day a dataset that allows it to infer what the shortfall would be if (all other things being equal) a member in a net debit position were to default that day. By collecting these data over a long period, the Bank can build up a detailed distribution of the net debit positions that typically occur in each DNS system. The Bank uses this distribution to infer the probabilities associated with different losses if a member of the system should default.

This is the conditional loss distribution — ie the distribution of losses arising, conditional upon one member defaulting. In order to complete its calculations of settlement risk estimates, the Bank also needs estimates of the likelihood of any one member of the payment system defaulting.

This approach to estimating settlement risk is not without difficulties. Among other things:

(1) For instance, Lacker, J M (2003), 'Payment System Disruptions and the Federal Reserve Following September 11, 2001', reports how events such as the terrorist attacks on New York in September 2001 forced the closure over several days of key parts of the US financial infrastructure (notably the stock market), although the main payment system, Fedwire, continued to be operational throughout the disruption.

(2) Specifically, interbank credit risk crystallises in DNS systems if a bank making payments fails, and either the receiving customer can (and does) make use of the funds prior to interbank settlement, or the surviving banks guarantee the settlement.

- It is based on the assumption that the net debit positions observed during the normal course of payments business is what would be observed if a member of that payment system actually defaulted. In practice, if market rumours regarding its viability should begin to circulate prior to its demise, a defaulting bank may be subject to a run, and this would be reflected in larger net debit positions in the payment system than is normally the case. Similarly, in the case of a quick-burn default (eg arising from fraudulent activity), it is likely that the cause of the default would result in larger obligations in the payment system than would be the case in ordinary circumstances.

- As for estimating each payment system member's probability of default in the first place, whether these are derived from Credit Ratings Agency estimates or by some other means, it is in practice extremely difficult to produce an accurate estimate of the probability of such an event, given its relative infrequency. Furthermore, while it is difficult enough to produce such an estimate in steady-state, the true probability of default may alter quickly as market conditions change.

- In addition, it is likely (based on the guidance set out in the 'Core Principles') that the DNS payment system has rules in place, which oblige its members to post collateral in advance that would cover the default of at least the single largest participant in the system. This would allow settlement to occur as intended. The overseer can take this mitigation into account when estimating the impact of settlement risk. But the surviving members of the payment system would, in due course, presumably need to replenish the collateral fund (net of any recoveries by the defaulting bank's administrators). This then is a deferred cost of settlement risk. But it is contingent on various assumptions — which, if all taken on board, would complicate the calculations significantly.

Notwithstanding these (and other) shortcomings, the methodology seems to be a step forward in estimating settlement risk.

References

Bank for International Settlements (2001), *Core Principles for Systemically Important Payment Systems*, Committee on Payment and Settlement Systems.

Bank of England (2005), *Payment Systems Oversight Report 2004*, Issue No. 1.

Bank of England (2006), *Payment Systems Oversight Report 2005*, Issue No. 2.

Bank of England (2007), *Payment Systems Oversight Report 2006*, Issue No. 3.

Bank of England (2008), *Payment Systems Oversight Report 2007*, Issue No. 4.

Bank of England (2009), *Payment Systems Oversight Report 2008*, Issue No. 5.

Federal Reserve Bank of San Francisco (1999), 'Using CAMELS Ratings to Monitor Bank Conditions', *Economic Letter 99–19*, June.

Financial Services Authority (2006), *The FSA's risk-assessment framework*.

De Fontnouvelle, P, Jordan, J and Rosengren, E (2006), 'Implications of alternative operational risk modelling techniques', in Carey, M and Stulz, R (eds) (2006), *The Risks of Financial Institutions*, NBER/University of Chicago Press.

Haldane, A G and Latter, E (2005), 'The role of central banks in payment systems oversight', *Bank of England Quarterly Bulletin*, Spring, pages 66–71.

Lacker, J M (2003), 'Payment System Disruptions and the Federal Reserve Following September 11, 2001', *Federal Reserve Bank of Richmond Working Paper Series no. 03–16*.

Manning, M, Nier, E and Schanz, J (eds) (forthcoming: expected 2009), *The Economics of Large-value Payments and Settlement*, Oxford University Press.

New York Payments Risk Committee (2007), *Financial Market Infrastructure Risk*.

Office of Public Sector Information (2009), *Banking Act*.