The UK approach to controlling risk in large-value payment systems

At SWIFT's annual SIBOS Conference, in Geneva in September 1993, a number of speakers described the steps being taken in their countries to manage, and where possible to eliminate, the settlement risks in their large-value payment systems. Brian Quinn, an Executive Director of the Bank of England, described the UK approach to this issue, which differs in some important respects from the approach adopted in other countries.

Introduction

Changing the way payment systems work is not an easy task nor one to be undertaken lightly. That is not necessarily because of the complexity of the systems themselves: in fact, although they often represent a very substantial investment in terms of both infrastructure and human resources, payment systems can be relatively straightforward compared with some of the other systems banks have to cope with. It is more because payment systems are of absolutely crucial importance both to banks and to their customers. Payment systems are, if you like, the heart of a bank, keeping the life blood of money moving through its various businesses; and they are no less vital to the efficiency and well-being of the economy as a whole. Any change to them has therefore to be undertaken most carefully. Any disruption or deterioration in performance has to be avoided, and also the inadvertent creation of any new risks.

This point is particularly relevant when discussing risk management policies, because the policies many countries are considering or are already putting in place frequently involve major change. This is certainly true in Britain where the banks and the Bank of England have agreed that CHAPS —our principal large-value payment system, which at the moment is an end-of-day net settlement system—should be converted to real-time gross settlement. Getting that change right is vital. Of course, circumstances vary from country to country. The payment systems themselves, the banking systems in which they operate, the relations between the central bank and the commercial banks may all differ; and so the specific objectives and details of the risk management policy may vary case by case. Nevertheless, the broad thrust is everywhere very much in the same direction, and thus it may be of some interest to describe why and how in Britain we have gone about making the change in case it provides any lessons—positive or negative—for others. I will also try to say something about the possible consequences of the changes in CHAPS for the customers of the system—a topic that may be of particular interest to those of you who, via SWIFT and your UK correspondent banks, are indirect users of CHAPS.

CHAPS is an important payment system. On an average day it handles payments worth a total of some £90 billion, while on a peak day the value can approach twice this. Much of its traffic comes from the settlement of large financial market deals—for example, interbank lending or the sterling leg of forex transactions. But the remainder comes from a wide range of other urgent transactions throughout the economy—all of which are important to the people making and receiving them. The conversion of CHAPS to real-time gross settlement will have to take place without disrupting this existing activity; the new system will also have to be able to cope, at a later stage, with a large number of additional payments arising out of securities transactions, since in due course we aim to forge a truly robust link between real-time securities transactions and real-time payment facilities.

Because of the need to get it right first time, the transition from net to gross settlement will not be a particularly rapid process. By the time it is complete, at the end of 1995, six years will have passed since the then Governor of the Bank of England started the ball rolling with a speech in which he raised the issue of risk in the wholesale sterling payment systems. But the change will not come as a big bang at the end of that six years—some of the elements are already being put in place, and we will be moving cautiously but steadily towards the objective. The banks have described this process as an evolutionary one, and in essence it contains, I think, four key steps: assessment; containment; agreeing the long-term objective; and implementing that objective. It is these steps that I would now like to describe in turn.

Assessment

First, assessment—perhaps the most important step. Awareness of the risks in national payment systems is still relatively recent. In 1989 it was not clear that there was a problem in the United Kingdom's systems, let alone exactly what it was, or how serious. It was important, therefore, that *everyone* involved in the process should first find out more about the subject.

I should stress here that our risk-management project is very much a joint exercise between the banks and the Bank of England. As a central bank, the Bank of England has a natural interest in the operation of the large-value sterling payment systems—not just because of the economic importance of their smooth functioning for the conduct of

public policy and the efficient functioning of the economy, but also because of the risk to the banks, and indeed to the Bank of England, that exists in the systems. But in the United Kingdom those systems are owned and operated by the banks—not by the Bank of England. Close co-operation is therefore particularly important. Whatever their nature and whether owned by the banks or by the central bank, one must of course never forget that payment *systems* are not an end in themselves but, rather, a means to enable banks to carry out their own payments business, and that of their customers, efficiently and effectively. Change to the systems can affect that business and so change has to be a joint process in which all the parties are comfortable with the objective and with the means of achieving it.

The first step, therefore, was for the members of CHAPS—including the Bank of England—to carry out a study into the nature and scale of the risks in the system. The results were not comforting. As a net settlement system, the CHAPS banks exchange payment instructions during the day but interbank settlement (the movement of the funds themselves across accounts at the Bank of England) does not take place until the end of the day. Because the system is designed to provide a speedy service for urgent payments, banks receiving a payment instruction are reluctant to wait until after settlement has taken place before giving value to their customers. Rather they are under pressure to give immediate (or at least same-day) value, and thus until the time of interbank settlement they are exposed to the risk that the sending bank will not provide the funds. In practice, it turned out that these intra-day exposures could be very large —so large, in fact, that if one of the larger banks in the system had failed, payment system losses alone could have wiped out the entire capital base of several other CHAPS banks and critically weakened several more. Moreover, there was no mechanism within the system that enabled banks to control this risk. Everyone therefore agreed that something needed to be done.

Containment

This analysis led fairly rapidly to the second step—namely an immediate plan for containing and as far as possible reducing the risks. In other words, short-term measures to put in place while the longer-term objective was being agreed. This decision was taken in the winter of 1990/91 for implementation starting in April 1992—more than a year ago. It involved retaining end-of-day net settlement in CHAPS, but making three major improvements.

The first of these concerned some important tidying-up of the *agreements* underpinning the system—agreements which previously were only partially codified. This may sound dull, but actually it was of great importance. One of the problems with payment system risk, especially in net settlement systems, is that participants sometimes misunderstand the nature of the risks they bear—they make assumptions about the risks which perhaps cannot be justified. Thus perhaps the most significant but least heralded change in CHAPS was an explicit agreement among the members that the intra-day exposures they bear within the system are the bilateral net amounts, rather than the multilateral net amounts as had sometimes previously been assumed. That at least made it clear where all the parties stood.

And that agreement helped to reinforce the second improvement, namely the introduction in the system of a mechanism to allow banks to set *limits* on their bilateral exposures to each other. This gave them for the first time a means of restraining, up to a point, the scale of the risks they bore. The bilateral limits constituted a sea change in the system because, also for the first time, banks were obliged to manage their payment flows in order to keep within those limits. Previously banks had been able to enter payments into the system in any order without having to worry about the consequences in terms of the size of any intra-day deficit they ran. Having to manage their flows was a fundamental step towards managing their risks. Because of the possibly disruptive consequences of this change, the limits were phased in gradually over a period of some eight months, giving the banks the opportunity to devise a *modus vivendi* between the two conflicting requirements inherent in payment system limits—namely, the prudential need to keep the limits low to reduce risk, and the operational need to keep them high to prevent payments being blocked or unduly delayed. The limits have now been 'live'—that is, fully under the control of the bank bearing the risk—since last December; and substantial rescheduling of payments instructions during the business day has been apparent since the middle of last year.

The third improvement has been taking place since March this year. This is the introduction of so-called net sender limits or system-wide caps as, in effect, a back-up to the bilateral limits. Whereas a bilateral limit constrains a bank's debit position *vis-à-vis* another particular bank in the system, the net sender limit constrains its aggregate debit position *vis-à-vis* all the other banks in the system collectively. I will return later to the usefulness of these net sender limits which, like the bilaterals, are being phased in gradually.

Agreeing the longer-term objective

If assessment was the most fundamental step, and containment the most urgent, then perhaps the third step—agreeing the longer-term objective—could be described as the most difficult.

Broadly speaking the major issue was choosing between what have been called risk reduction and risk elimination strategies. *Risk reduction* meant, in essence, retaining the short-term measures for the longer-term, albeit with some modifications. CHAPS would have remained an end-of-day net settlement system and although risks in the system would have been substantially reduced, they would still have been present, in possibly significant amounts. *Risk elimination* meant designing a real-time gross settlement system whereby a bank would not be able to send a payment instruction to another bank in the system unless it had

available the necessary central bank funds to settle that instruction.

Many commercial bankers instinctively felt that risk elimination was too extreme a strategy. After all, banks are used to managing interbank exposures in the foreign exchange and money markets, so why not in payment systems as well? Surely it was excessive or even perverse to say that the best way to manage something was to remove it entirely?

That argument was understandable—but, I believe, it was also fundamentally flawed. Settlement risk in the payment systems is *not* the same as other kinds of interbank risk, because it is difficult if not impossible to control it adequately. The only safe way to control it is therefore to remove it.

But why is it so difficult to control? As I noted earlier, in net settlement systems like CHAPS the credit (and hence the risk) is provided automatically by the bank receiving the payment instruction, who does not get paid by the sending bank until settlement takes place, typically at the end of the day. Such a payment system is therefore carrying out two functions—it is carrying out its primary task of processing the payments from Bank A to Bank B, but at the same time it requires Bank B to give credit to Bank A as a by-product of that process. And it is too much to ask one mechanism to perform both functions. It means that the exposures are primarily determined *not* by the direct credit judgments made by the receiving bank about the sending bank but rather by the pattern of payments between the two banks concerned. If on a particular day customers of Bank A send more payments than usual to customers of Bank B, then, other things being equal, Bank B will have a higher exposure to Bank A. If Bank B tries to contain that exposure by setting a limit it is likely to delay or block the payments which its customers are expecting. That is the choice between operational and prudential needs that I mentioned earlier when talking about the short-term measures; and, given the competitive pressures facing banks today, who could be confident that the choice will always be exercised in the direction of prudence.

Of course there *are* measures that banks can take to try to manage these risks—indeed, the short-term measures in CHAPS do exactly this, and have been very successful in their way. But they can only go so far. There is a minimum level of credit which has to be provided (and thus risk to be borne) if the business of the national economy is to flow smoothly, as it must. Even if that level of risk were to be acceptable now, it is unlikely to remain so in the future as the system handles more traffic. Over the years since CHAPS started, in 1984, the market for sterling same-day payments has grown considerably—the real value of payments made (that is, after discounting the effects of inflation) has increased by around 200%. In the future, moreover, payments related to securities transactions are likely to be put through the system as part of the process of introducing real-time delivery versus payment; this alone is likely to add a further 50% to the value of the total payments the system has to handle.

European developments may also increase the scale and complexity of the problem. The single market is likely to bring with it increased participation in domestic systems by banks from other member states, thereby compounding the difficulty of ensuring that the legal basis of a net settlement system is robust: netting is a legally complex issue under the national laws of most, if not all, member states and for a net settlement system with multinational participation to be safe, the law of every country involved has to be adequate. That is far from clear at present.

We felt very strongly, therefore, that any large-value system that was going to be robust enough to serve the sterling payments market's needs both now and in the future needed to be based on real-time gross settlement, where the inherently uncontrollable and legally difficult form of interbank credit that exists in net settlement systems was removed entirely—and that indeed is the solution that the banks have decided to adopt. Of course, intra-day credit of some kind will still be necessary in such a system to keep it liquid; in the United Kingdom, it has been agreed that this credit will be provided by the Bank of England, collateralised in some form. In that way the banks will get the credit they need, and the Bank of England will be protected against the resulting risk by the assets provided in exchange. Liquidity will be maintained, without the uncontrollable risk.

Implementing the long-term objective

The decision in principle to convert CHAPS into a system that settles in real-time across accounts at the Bank of England was taken by the banks last autumn and, following an initial feasibility study, reaffirmed in March this year. Final specification of the project still has to be approved but we are now effectively into the fourth and final stage, that of implementing the decision by the end of 1995.

That is a somewhat daunting task at first sight. But the evolutionary approach has made it significantly more manageable because the ground has already been well prepared. Comparing CHAPS as it was before the risk-management programme started to the essential features of the future RTGS system, there are three main changes involved: changing the way messages are routed through the system; constraining the availability of credit; and collateralising that credit.

Payment systems depend crucially on means of communication. As part of the step-by-step way in which change has been implemented, the banks have decided not to build a new communications system from scratch, but rather to adapt the arrangements that are already in place. In the existing CHAPS system, payment messages are routed directly from the sending bank to the receiving bank: unusually, even for a net settlement system, CHAPS has no central clearing house or computer. The major *system*

change needed for RTGS is thus a change in the routing, involving a new element of centralisation, so that messages are sent first to the Bank of England for settlement.

Detailed designs are still being considered, but the way in which this is likely to be done will, I think, be unique to the United Kingdom. When it comes to payment routing, there is a kind of alphabet of RTGS systems. The classic design, if I can call it that, which is typified by, say, Fedwire in the United States or the Swiss system SIC, is V-shaped: payment messages go from the sending bank to the central bank (or to Telekurs on behalf of the Swiss National Bank), and then, after settlement, on to the receiving bank. In contrast you will already have heard today about the Premium service offered by SWIFT which is to form the basis of, for example, the TBF system being introduced by the Banque de France. That system is often described as T-shaped—one message goes from the sending to the receiving bank, whilst a copy gets routed to the central bank. In contrast to the V and the T, the UK arrangements are likely to be L-shaped. The core details of a payment message—those that are needed to enable settlement to take place—will first go from the sending bank to the Bank of England; once settlement has taken place, confirmation will be passed back to the sending bank, who will then send the entire payment message directly to the receiving bank. This L-shape has no great merit in its own right, but in system terms it is convenient because it fits in well with the existing distributed architecture of CHAPS. Adapting the existing communications system in this way will, we hope, mean less disruptive change—and less expense—than starting from scratch.

The second of the three main changes needed is a change in management. Whereas in CHAPS as it used to operate there was no constraint on the intra-day debit positions that banks could incur, under RTGS banks will need to manage their payment flows in order to keep within the credit they obtain from the Bank of England. As I noted earlier on, this is a major change—requiring banks to implement new internal systems (to enable them to reorder their payment flows) and new management skills. Again, however, the impact of the change has been minimised by the evolutionary approach. Because of the short-term containment measures already in place, banks are *already* living with bilateral credit limits. And with system-wide 'net sender' limits now being phased in, the approximation to a regime of real-time gross settlement will, in terms of credit availability, soon become even closer. One of the most difficult changes needed for RTGS is thus already well under way.

The third element of change necessary is the provision of collateral to the Bank of England. This is perhaps the most controversial area: there is bound to be some pain involved in moving from an arrangement where unlimited, free intra-day credit is available (albeit at significant risk, ultimately to the central bank) to one where the credit is still free but obtained by providing collateral. It is vital from the point of view of the smooth functioning of the system that adequate liquidity *is* available. But if the central bank is to supply that liquidity, I think it is also appropriate that collateral should be provided to protect the central bank from the resulting risk: it is not the business of the central bank to underwrite bank risk.

Of course that leaves the banks concerned about whether they have enough collateral available to obtain the credit they need and, if not, what the marginal cost of obtaining the balance will be. I understand those concerns. Indeed an important part of the gradual process of change is discussion between the Bank of England and the CHAPS banks how these concerns can be minimised. Thus, for example, we have drawn up a list of the assets we would be willing to accept as collateral and have indicated that we would be prepared to consider a fairly wide range provided a good case for their inclusion can be made. We have agreed, moreover, that the banks can make intra-day use of those assets they already hold on their balance sheets as part of their liquidity management. And, as far as the system design allows, the banks will be able to move collateral to and from the Bank as they need it, during the day. These discussions are still continuing but are already, I hope, helping to lead the banks to the conclusion that the provision of collateral may not be too expensive.

But in the end it has to be the responsibility of each bank to ensure that it has enough liquidity to handle the volume of payments it is putting through the system. Put in another way, it is up to each bank to decide how much of its resources it is prepared to devote to its own payment business, and to that of its customers. That is a commercial decision for each bank. The responsibility of the central bank should, I think, be confined to ensuring the system *as a whole* has enough liquidity to meet the needs of the economy. Again, however, the evolutionary approach should help. By the time CHAPS is converted to RTGS at the end of 1995, CHAPS banks will have had over two years of living with the net sender caps I mentioned a few moments ago, and thus they will have had ample opportunity to discover how much liquidity (and hence collateral) they do actually need to process their payments smoothly. They will therefore have time to make whatever adjustments they deem necessary.

Consequences for customers

So much for the four main steps in our evolutionary approach. Before I finish, I would like to say a little about the consequences of the changes in CHAPS for its customers—and I suspect that that will include many of you in the audience, for most banks at one time or another are indirect users of CHAPS via the links you have with your correspondent banks in London.

A key objective of the exercise is, of course, to minimise the consequences of the change for customers. At the moment the fourteen existing CHAPS banks are able to use the system to offer their customers a very fast, efficient and reliable service for sterling same-day payments, and the transition to RTGS has to be managed to ensure that they can continue to

do so. But it would be foolish to pretend that the change can be costless.

The cost arises from the change from a world in which intra-day credit is free and available in unlimited amounts to one where it is rationed and available only in return for collateral. It is a change from a world where the costs, in the form of risk to the system and to the central bank, are implicit and unmanageable to one where, in the form of collateralised credit, they are explicit and controllable. But to keep the amount of collateral they provide for the purpose to a minimum, CHAPS banks will, as I have already mentioned, need to manage their payment flows more actively, and since these flows are to a great extent determined by their customers' payments, the banks will in turn be looking to their larger customers in particular to engage in similar management. Moreover, to the extent that they do, in spite of this management, have to provide additional collateral, CHAPS banks may need to look to their customers to bear a share of the resulting cost.

Overall, however, I am convinced that the result will be very much worthwhile. The system will be significantly more robust because of the removal of its interbank risk; and in due course the way will be open to provide customers with a range of new services including true DVP in the securities and foreign exchange markets. That will be a major gain.

Conclusion

In conclusion, let me say that in implementing our risk-management programme we have learnt much from the experiences of other countries—including the United States, because they have been travelling along this road the longest, and also Switzerland and France, who have recent experience of the conversion to RTGS. I hope that, in turn, others will be able to benefit at least a little from Britain's step-by-step approach to the problem. That approach will not, unfortunately, make the conversion from net to gross settlement entirely painless, but it should do much to avoid many of the headaches that can be caused by a big bang.