FEDERAL RESERVE BANK *of* NEW YORK *Serving the Second District and the Nation*

OPSummit 98
Paris, France

February 24, 1998

Ernest T. Patrikis
First Vice President
Federal Reserve Bank of New York

It is a great pleasure to be here today to discuss operational risk in payment systems with many of the leading experts in the financial services industry. In preparing my remarks for today, I began by asking myself some very basic questions: (1) what exactly do we mean by the term "operational risk," (2) can an institution effectively measure and manage this risk, and (3) should banks develop further the concept of reserves and/or capital allocations for such risk. I am not sure I can provide you, up front, with easy answers to any of these questions. But by discussing operational risk first at a somewhat more abstract level, we might gain some insights in how to answer these questions later on in this presentation. Therefore, let me begin first with some general discussion of operational risk and, with this background in mind, come back later to the task of answering these three difficult questions with some precision and, I hope you will conclude, with some "creative insights."

Many of us probably use as a working definition of operational risk in our day-to-day, payments-related activities something along these lines: an unexpected problem with our computers that prevents us from making or receiving payments in a normal, routine way. A problem that prevents us from doing so could come about because of a glitch in the system, a human error by an employee, or because of an "act of God" such as a severe storm that causes a major power failure. To deal with these possibilities, we continuously test software with our vendors and customers as well as maintain fully functional back up sites on alternative power grids. We also plan for full recovery from many difficult situations including worst-case scenarios. Stringent security procedures are in place, and we aggressively monitor our networks and systems to protect them from insider and outsider attacks. And we encourage our employees, through rather detailed sets of procedures and controls, to alert top management when operational problems occur, rather than trying to correct them in isolation.

These exercises, procedures, and controls become more and more important as we rely on our payment systems to process greater volumes with less intervention points, as payments flow from the initial payor to the final receiver over compatible, linked message and payment networks -- what is now commonly called straight through processing. Increasingly, these additional linkages mean that the successful operation of one bank's system depends on the successful operation of other systems outside of that bank's control.

But what exactly is "at risk" if an institution encounters an operational problem? The answer, of course, depends on the nature of the problem and may not consist only of the loss of monetary value. It may also involve the reputation of the bank in the marketplace. And therefore, the ability to compete in the future is at risk if customer confidence in a bank's ability to make and receive payments, promptly and accurately, is lost.

Thinking now in somewhat broader terms, the consequences of an operational problem could: (1) be limited to the firm experiencing the problem, (2) spread to other participants in the payment network, or (3) endanger the operation of an entire payment system. Let me give you some examples of these three cases. In some of our large-value payment systems, incoming payments during the day, actual or anticipated, become the liquidity for outgoing payments. Recently, we had a case in the United States in which a bank's operational problem resulted in a situation where the bank could receive incoming payments but not make outgoing payments. As a result, this bank absorbed a large part of the available liquidity in the system, and at the end of the day some other banks, without any operational problems of their own, suddenly found themselves without adequate liquidity to close the day without overdrafts. In other words, this liquidity was anticipated and scheduled to be there, but it simply did not materialize. Hence, I always like to point out with some emphasis that one bank's operational risk could well turn out to be another's liquidity risk.

But it does not always work out that way. Several years ago a large money- center bank experienced an operational problem that permitted it to receive but not send securities against payment. It incurred an enormous overdraft in its funds account during the course of the day that eventually was corrected through intervention by the Federal Reserve Bank of New York that consisted of the largest discount window loan ever extended, secured by those securities the bank could pledge and repledge and all pledgeable domestic assets of the bank. In this case, that bank's operational risk turned out to be its own liquidity problem. The cost to the bank was substantial -- the cost of the discount window loan, as well as the cost of borrowing funds later during the reserve calculation period to maintain the required level of reserves even with the large shortfall on the day of the operational problem. In addition, the purchasers of the securities which did not receive them got a free ride for that day -- they would earn that day's interest on the securities but did not have to pay for the securities on that day. Somehow, this outcome seems a little more fair to me than the previous case when an operational failure of one bank caused liquidity problems for other banks, but please believe me when I say it was the type of situation that we do not want to see reoccur.

Finally, if the operator of the payment system, the central bank in the case of many real-time gross-settlement systems, has an operations problem, it becomes the problem of everyone else in the system. Payments simply do not flow at all among the participants, and the system "freezes in place" at the time of the problem, leaving some participants in potentially very difficult situations. Needless to say and I should touch wood when I say it, central banks make every effort to ensure this situation does not occur in their real-time gross-settlement systems -- too much is at stake.

In any case, the point I am trying to make here at the beginning of my remarks is that it is very difficult to predict in advance the exact nature of the problems that will occur if an operational failure in processing payments takes place. It could affect a bank on the send or receive side only, or on both sides at the same time. The problem could occur late or early in the day, affecting the amount of time a bank would have to fix the problem or attempt to route transactions through alternative systems. Or, the computers could be apparently working just fine on the receive and send side all day long, but somehow a security breach could occur and false transactions initiated or data records contaminated. In which case, the system might need to be shut down until the security breach could be repaired.

Operational risk can come from many sources, and currently it is coming form one source that I sincerely hope is very much on the mind of each and every person in this room. Of course, I am talking about the year 2000 challenge. If we are not adequately prepared for the turn of the century, our payment systems could be adversely affected in ways that we cannot fully anticipate at this time. And I believe that banks, other firms, and clearance and settlement systems in the payments and finance industry need to look at this problem at least at two levels. The first level is making sure our own individual systems are tested and ready to process after the turn of the century. The second level, because we are all connected through payment networks, power grids, and telecommunications networks, consists of contingency planning in the event our own systems are able to process, but some of our counterparties or utility suppliers are less fortunate. In other words, "their problems" could suddenly become "our problems" on January 3, 2000 through what I will call "network effects" of various types. What precisely should we be prepared to do under these circumstances in terms of contingency planning? Would a bank, for example, that is ready for the year 2000 want a credit exposure in terms of a large volume of incoming payments from a bank or corporate customer that is not prepared, or from a bank or

corporate customer that is prepared in terms of its own internal systems, but there is serious doubt about whether the utility providing power to that bank will be ready? Are these credit risks that need to be managed in advance, or to put a slightly different spin on this question, will banks that are fully ready for the year 2000 operationally also need to be prepared to route transactions for their major customers around banks that are not ready? Finally, should a centralized information clearing house on the preparedness of various organizations be established so that informed credit and other decisions can be made well in advance? I believe these are questions we all -- central banks, commercial banks, utilities, and other participants in the payments system -- need to ponder.

At the Federal Reserve Bank of New York, both as participants in Fedwire and as the home of the Wholesale Payments Product Office of the Federal Reserve Banks which operates of one of the world's largest interbank payment systems, we are taking every step necessary to ensure that both our funds and delivery-versus-payment, securities systems are ready for the century date change, including making sure that our customers have made the necessary adjustments to their systems to be able to successfully inter-face with our systems without any year 2000 problems. End-to-end testing between customer end points will begin later this year.

As a bank supervisor, we are undertaking a coordinated effort with other supervisors to closely monitor the efforts of commercial banks to get ready for the year 2000 in all aspects of their operations, covering awareness, assessment, renovation, testing, and implementation. Target dates have been established for banks to complete each of these phases, with banks not in compliance being subject to supervisory actions.

At the next level, given our membership in the banking and financial services community more generally, we are working with various industry groups and observing closely the efforts of firms, both on and off Wall Street, to prepare for this event that is now less than two years away. To facilitate this effort, we are hosting quarterly meetings with representatives from the banking industry, insurance companies, mutual funds, securities firms, payment system operators, and others to share knowledge and make sure we understand this rather unique problem from every potential angle.

Of course, many angles to this problem can be identified. Banks buy various computing and processing services from several vendors and, in turn, sell payment products to their our customers. Outsourcing of the data processing for various products is also quite common place. As a result, there are indeed "many angles to consider" in preparing our payment systems for the turn of the century, not all of them under our direct control, but all of them worthy of our attention and careful scrutiny.

And finally, the Federal Reserve Bank of New York has been active at the international level as well in addressing the year-2000 problem. The G-10 central banks will be making information available on the preparation of their respective interbank payment systems for the year 2000 at the web site of the Bank for International Settlements. In addition, William J. McDonough, President of the Federal Reserve Bank of New York and Chairman of the Committee on Payment and Settlement Systems of the G-10 central bank governors, has written to central banks throughout the world encouraging them to promote this effort to establish that single BIS Internet site where information on the year-2000 readiness of payment and settlement systems worldwide would be available to participants.

By now, you can clearly see that all of my time here today could be spent discussing the turn-of-the-century, operational problems for payment systems. With the assumption that I have at least impressed you with the seriousness of this problem, as well as the intensity of our efforts to cope with it, let me return now to the more general topic of operational risk. In recent years, banks have taken a new approach to assessing the risks they take. Banks have constructed models to calculate the "value at risk" in many of their business lines. The simulations from the models, in turn, allow them to assess whether capital is adequate given the nature of the business risks they undertake. Banks can even "stress test" to see what worst case scenarios might look like. Developing models to help us get at the value at risk from an operational failure is conceptually far more difficult than is the case for market or even credit risk, although I understand some banks are beginning to work on models for this risk factor as well.

I view this work as a positive development because central banks, as many of you already know, have been actively concerned about operational, systemic, and other risks in payment systems for a number of years and have undertaken several joint efforts to reduce risk under the sponsorship of the Committee on Payment and Settlement Systems. You can review these efforts -- which cover a fairly broad range of topics -- by visiting the BIS's web site on the Internet.

And while you are "surfing the net" to take a look at these documents, it might also be worthwhile to reflect on what types of operational problems could develop as retail, electronic payment systems are developed for the Internet. Our large-value systems are usually closed networks, but the Internet is an open network with perhaps a broader set of potential "operational problems." Currently, technology companies are busy creating electronic versions of currency, checks, and credit-card payment systems. Electronic bill presentment and payment is being actively promoted, and many believe this is the application that will give PC banking on the Internet the necessary momentum to make a significant dent in our paper check volumes.

Our banks, of course, will be right in the middle of all this, using this technology to provide payment services to their retail customers. The failure, however, of these retail systems to operate as promised will at the very least put the banks at substantial reputational risk, and I would add, potentially at risk of onerous legislation and regulation because individual consumers are also voters who might complain quite loudly to their elected officials. As a result, as banks open up their computer systems to access through an open network, they must be very careful not to create the potential for new operational problems, whether these problems are the more conventional glitches or the result of hackers attempting to practice their trade for fun and profit.

In addition, as banks rely increasingly on technology companies to provide them with the means to implement electronic retail payment systems, I cannot help but wonder if banks are fully aware of the operational risks they might face. This new technology is highly complex, involving sophisticated encryption and other high-tech features. But once large sums of money are at stake, we can only assume the sophistication of hackers will increase as well. Can banks afford to take technology companies at their word when these companies claim the banks' operations can be adequately protected? Or is there a new type of operational risk developing here as well?

I am beginning to believe that Internet banking and payments could well prove, once we get past the challenge of the year 2000, to be the banking system's next big challenge at managing operational risk. I am concerned about these potential problems because: (1) our experience in managing these open networks is still quite limited, (2) a severe shortage appears to be developing of people with the highly technical skills necessary to administer, operate and support these complex networks, and (3) organized crime, hostile foreign governments, and others with substantial resources will have an interest in creating operational problems, or in making us believe our Internet payment systems are operating just fine, as they remove large sums of money and/or information.

Therefore, I believe that security for payments on the Internet as well as for our large-value, payment systems will continue to be an operational challenge in the years ahead. In the United States, we have a number of small, high-tech companies consisting of engineers, cryptologists, and other scientists who work full time using powerful computers to test other firms' electronic security systems. We have retained security consultants to perform these types of tests. I assume that you also have done so. What I have learned is that we will never feel comfortable with the level of security. It is something that we must all strive to improve continuously.

With this background discussion of operational risk, including the security aspects, in mind, let us now go back to the questions I raised at the beginning of my remarks, that is: (1) how should we define operational risk, (2) can we measure and manage this risk, and (3) should banks put aside reserves and/or hold capital for this risk? For me, these are three very difficult questions, but I doubt you would find my remarks all that interesting today if I did not take make an attempt to provide some answers. Nonetheless, before I attempt some answers please allow me to emphasize that: (1) the answers are my own personal views, (2) my thinking on this subject is still very preliminary in nature and likely to

change over time as more information becomes available, and (3) I would also be interested in your views on these issues because I am still very much in the information gathering and processing mode with respect to these issues.

Regarding the definition of operational risk, I believe we can break it down into internal and external risk, along the lines of the examples I gave you earlier. So if we limit ourselves to payments processing, our definition of operational risk becomes: the risk of not being able to process payments without incurring other difficulties because of either internal or external technical problems.

For measurement of internal operational risk, I suspect we are limited to looking at the historical performance of our payment systems and see what happened in the past when these systems did not operate properly. This historical record, if it has been maintained for a number of years, might give us some idea of the potential for operational risk from internal factors. I use the phrase "might give us some idea" because our payment systems are not static over time; we are constantly improving and upgrading them, and if a problem occurs, we always try to "implement a fix" that will insure the problem never occurs again. On the other hand, if we spend all our time looking in the past at the historical record of our operational problems, we might never see a problem, like the turn-of-the-century problem, coming at us without any precedent, from the future. And as I mentioned earlier, the operational risk associated with Internet banking could be another example of "no comparable experience from the past."

To attempt a measure of the potential for operational risk from external sources, we would need performance data from other participants as well as from payment system operators. It is not likely these data from external sources would be readily shared, and most likely the only information we would have as individual institutions would be for those instances in which we were directly affected, if in fact we ever bothered to record this information in a systematic way. And once again, external operational risk, like that created by the year 2000 problem, would not be captured in a historical data base.

If measurement is likely to be difficult, is management of this risk impossible? The management of this risk is indeed possible; what is not possible as a result of the incomplete, static, backward-looking data is a precise cost-benefit analysis. Operational risk can be managed by: (1) extensive internal and external audits, (2) review by external technology and security companies, (3) work with industry groups to identify "best practices" for every dimension of our operations, (4) cooperation with central banks and other regulators to insure our payment systems are not potential sources of external operational problems, and (5) careful attention to all this information by top management and directors who give the entire effort a truly high priority.

One of the more difficult issues in addressing operational risk is that the probability of occurrence of an event can be quite low while the potential loss from such an event very high. Clearly, large amounts of money and human resources can be spent in reducing operational risk to extremely low levels. But because the effort, by its very nature, is largely "preventative maintenance," designed to insure against outcomes we cannot precisely measure or predict, we will never be sure the "cost is worth the benefit." I tend to think of these high-level, management problems, for which cost-benefit analysis is not possible, as the ones where we really earn our annual salaries. Problems that can be neatly reduced to a cost-benefit ratio are the "easy ones" that can be delegated on down the line.

In any case, if measurement of operational risk is difficult, but management possible, what can we say about the question of reserves and/or capital for operational risk? The banking industry is only beginning to focus on this question, along with the measurement and management issues we just discussed. My current, and still quite preliminary, thinking regarding reserves for operational risk is that, to the extent a bank can identify a predictable, stable component of operational risk that lends itself to coverage by insurance or reserves, it might not be a bad idea to do so. But what about events, such as the year 2000 or the compromise of a new Internet payment system, where the probability of losses might be fairly low, but the potential for losses, if something happens, totally unpredictable at this time? Would an allocation of capital be appropriate at this time? If so, how much? I must admit I have just done something that I truly dislike when other people do it to me. I have answered my own question, not with one, but three additional questions. But for a reason. Reserves and/or capital for operational risk would need to be studied a great deal more before anyone could take a firm stand on this last issue.

In conclusion, I hope that you will find my remarks, occurring near the opening of this program, as provocative ones that will help initiate a lively discussion of the issues. The concept of operational risk is a slippery one at best. It is not always clear who will bear the consequences of an operational problem and whether operational risk can be neatly separated from liquidity and the other risks banks face. Nor is operational risk likely to be static, easily measurable, or always predictable in nature. It will evolve over time as we expect more from -- and expand the linkages of -- our large-value payment systems, at the same time we move our retail payments to electronics on open networks. Nonetheless, even with all these conceptual and practical caveats, operational risk remains potentially far too important a risk to ignore, even by the most senior of management. The year 2000 reminds us of this each and every day.