Understanding the Asian crisis: Systemic risk as coordination failure

David Marshall

Introduction and summary

The vast array of financial regulations in the United States and other developed economies is justified largely as a way of protecting the public from the dangers of *systemic risk* or *systemic crisis* in financial markets. Even the title of a recent General Accounting Office report on financial derivative regulation ("Financial derivatives: Actions needed to protect the financial system" [GAO, 1994]) and images in the popular press (a close-up of a snake with jaws wide open on the cover of *Fortune* magazine) appeal to our fear of systemic risk. While many different (and often mutually contradictory) characterizations of systemic risk have been proposed, it is somewhat disturbing that we lack a consensus as to what, precisely, systemic risk is.

In its most general usage, the term systemic crisis describes a shock to the financial system that impairs crucial functions of the system, such as asset valuation, credit allocation, and payments.[1] This characterization, however, is not too helpful. What sort of mechanism can result in this sort of impairment? Economists do not agree. Proposed answers include: irrational piling-on of debt;[2] moral hazard induced by mispriced government-provided deposit insurance;[3] complex relationships among counterparties;[4] an unwillingness of dealers to trade;[5] a failure of the central bank to provide liquidity as needed;[6] unpredictable adverse shocks that come from outside the financial sector;[7] and bank runs.[8] Regardless of which characterization one prefers, a satisfactory theory of systemic risk requires a fully articulated, internally consistent economic model. A rigorous economic model may leave uncertainty as to how the model maps to reality, but there can be no uncertainty about what is meant within the context of the model itself.

While economists may disagree as to the causes and nature of systemic risk, there have been specific events in history that are generally recognized as examples of systemic crisis. The most recent such event is the Asian crisis that began in mid-1997. This crisis displays certain textbook characteristics generally associated with systemic crisis: It appeared to originate in financial markets; it displayed *contagion*, with problems in one country seeming to induce crises in other countries; there was clear evidence of confidence loss by investors; there were substantial real costs in economic output; and the crisis clearly called for a policy response. One aspect of the Asian crisis that is more difficult to explain using standard theories is that it seemed to emerge almost spontaneously. Although, with hindsight, one can point to conditions that may have made some East Asian economies vulnerable to economic disturbance, the crisis was not forecasted by knowledgeable observers, nor was it triggered by any shocks commensurate with the scale of the upheaval.

I argue in this article that the *standard neoclassical model* commonly used in economic analysis is poorly suited to explain the sort of crisis in which a small impulse induces a large change in economic performance. Rather, the Asian crisis is best explained as an example of *coordination failure*. Suppose the economic performance of a country (or a firm, industry, or financial market) depends on large numbers of investors being willing to provide funds. If it is generally believed that *other* investors will withhold funds, it is rational for any *given* investor to refrain from investing. Thus, these beliefs become self-fulfilling. This

David Marshall is an economic advisor and senior financial economist at the Federal Reserve Bank of Chicago. *The author acknowledges helpful comments and suggestions from Doug Evanoff, Mike Kouparitsas, and Jim Moser.*

represents a coordination failure because *everyone* would be better off if *all* investors provided funds to the affected country. Unfortunately, there is no way to coordinate investor actions in this way.

In this article, I formalize this notion of coordination failure in a simple static model. My model implies that, as in the Diamond–Dybvig (1983) model of bank runs, some credible insurance mechanism is necessary to avoid costly coordination failure. The U.S. savings and loan crisis shows the danger of government-provided insurance without adequate regulatory oversight. However, it is not obvious how to establish such oversight in an international context. The recent events in Asia suggest that the world community eventually must face this issue.

The Asian financial crisis

Since a detailed chronology and discussion of the events surrounding the Asian crisis can be found elsewhere,[9] I summarize only the essential facts in this article. The years preceding the crisis were a period of exceptional growth in the East Asian economies. In the months preceding the crisis, no real macroeconomic distortions were observed. The economies of the five crisis countries (Thailand, Malaysia, Indonesia, South Korea, and the Philippines) were characterized by low inflation (less than 10 percent), budgets generally in surplus, and declining government foreign debt (as a fraction of gross domestic product [GDP]). For example, Indonesia experienced 10.4 percent export growth in the year preceding the crisis, its government budget was in surplus each of the previous four years, and its current account deficit was only 3 percent to 5 percent of GDP. During the 1990s, these governments engaged in responsible credit creation and monetary expansion. Unemployment rates were low and did not provide an incentive for governments to engage in currency depreciation or monetary expansion as a short-term stimulus.

However, there were some worrisome signs of imbalance that, in retrospect, made these economies vulnerable to crisis. First, the ratio of short-term debt to short-term assets was high and growing. In particular, the ratio of short-term debt to foreign exchange reserves in Thailand, Korea, and Indonesia had exceeded 1.0 since 1994. Radelet and Sachs (1998b) show that this ratio is positively correlated with systemic crises, but does not inevitably lead to crisis. In addition, the crisis was preceded by a large increase in foreign bank lending. In particular, from yearend 1995 to yearend 1996 foreign bank lending to the five crisis countries increased by 24 percent, with an additional 10 percent (annualized) increase during the first half of 1997. Finally, anecdotal evidence suggests that much of this expanded credit was invested in real estate markets, rather than increasing productive capacity.

Looking back, it is possible to discern developments in 1996 and 1997 that may have been unfavorable to these Asian economies. For example, Corsetti, Pesenti, and Roubini (1998) point to a fall in the demand for semiconductors in 1996 and expectations of contractionary monetary policy in Japan and the U.S. in spring 1997. However, one can always use hindsight to point to less-than-perfect business conditions before a crisis. Without a doubt there were also unfavorable developments during the years of unprecedented growth. Why did they precipitate a crisis this time?

The evidence seems clear that the crisis took the world completely by surprise. Bond spreads between emerging market debt (including that of the Asian five) and U.S. Treasury securities actually *fell* between mid-1995 and mid-1997.[10] Similarly, syndicated loan spreads were generally lower in early 1997 than they had been in 1996. Moody's and S&P's ratings of sovereign debt gave no indication of an impending crisis. Long-term sovereign debt ratings for the five crisis countries remained unchanged throughout 1996 and the first half of 1997, right up to the onset of the crisis. The only exception was an *upgrade* for the Philippines.[11]

The first events leading to the crisis were failures of certain large Korean chaebols (industrial conglomerates). In January 1997, Hanbo Steel declared bankruptcy. This was particularly significant because it was the first bankruptcy of a leading Korean conglomerate in a decade. This was followed in March 1997 with the failure of another Korean conglomerate, Sammi Steel. Around the same time, Kia, Korea's third largest auto maker, had difficulty rolling over its debt. (In July, it asked for emergency loans.)

Meanwhile, in Thailand, nonbank finance companies suffered under the weight of bad real estate loans. It was widely thought that the Thai government would aid these institutions. In February, Samprasong Land, a finance company, missed payment on its foreign debt. On March 10, the Thai government appeared to state that it would buy up to $3.9 billion in bad property debt from various finance companies. However, it reneged. The collapse of Thailand's largest finance company, Finance One, on May 23 provided explicit evidence that the Thai government would not bail out investors in these finance companies.

After repeated declarations that the baht would not be devalued, on July 2 the Bank of Thailand announced that it would allow the baht to float. The baht fell by about 15–20 percent. Immediately, the Philippine peso came under attack, followed a few

days later by the Malaysian ringgit. The crisis was under way. The Indonesian rupiah did not come under severe pressure until mid-August. (By late January 1998, it had lost 80 percent of its value against the dollar.) By November, currency weakness had spread to the Korean won. When the crisis reached its peak, foreign exchange rates came close to a "free fall." On several trading days in November and December, the Korean won plunged the limit of its trading range in the first few minutes of trading. In the first four trading days of 1998, the Indonesian rupiah hit four consecutive record lows, for a total loss of 27 percent against the dollar. On January 8, it lost another 26 percent in a single day.

The crisis soon spread to the stock markets. During four days in mid-October the Hong Kong stock market lost nearly 25 percent of its value. In early November, stock markets in Taiwan and South Korea plunged. Equity markets elsewhere also posted sharp losses. For example, on May 19 (in apparent response to the political upheaval in Indonesia), Russian stocks plunged nearly 12 percent, with equity markets posting smaller declines in Brazil (6.4 percent decline), Argentina (4.7 percent), and Mexico (3 percent).

Clear evidence of a collapse in investor confidence can be seen in the dramatic reversal of capital flows. In 1996, the capital inflow to the five Asian crisis economies was $93 billion. In 1997, the figure was –$12.1 billion, with estimates for 1998 of –$9.4 billion.[12] The reversal from 1996 to 1997 represented 11 percent of the combined GDP of these five countries. Korean firms had great difficulty raising short-term capital. Part of the collapse of the Korean stock market was attributed to selling of shares by cash-short institutions in an effort to raise operating cash. Indeed, at a meeting of finance ministry officials in November, U.S. Deputy Treasury Secretary Lawrence Summers stressed that this was not a traditional balance of payments crisis but fundamentally a crisis of confidence.

Asset and foreign exchange prices seemed to stabilize in Asia by February 1998. However, the economic impact of the crisis continued. South Korea entered a recession, with a first quarter 1998 decline in output of 3.8 percent/annum. (The previous year's growth rate was 4.3 percent.). Korean consumer price inflation was 9 percent in March 1998 (year over year), compared with 4.5 percent in the previous year. Producer prices recorded a more dramatic rise: Korean producer price inflation was 17.5 percent for the year ending March 1998 (versus 3.8 percent in the previous year). The Korean jobless rate in March was 4.7 percent, compared with 2.5 percent a year earlier. A record high 1.2 million people were unemployed.

The real impact of the crisis in Indonesia was even worse. By mid-May, the consensus forecast was for a 1998 economic contraction of 7–8 percent. (Indonesia's GDP grew at a rate of 5 percent in the previous year.) The official data showed an April 1997–April 1998 inflation rate of 44.9 percent. The official unemployment estimate was 10 percent, although many analysts regard this as an underestimate. Nor were other countries spared. First quarter data for Hong Kong reported a 2 percent GDP contraction, the first quarterly contraction in Hong Kong in 13 years. The consensus forecast for GDP growth in Thailand was –6.0 percent, compared with a 2.5 percent expansion in 1997. The full impact of the crisis has yet to be assessed. There remains a substantial short-term debt overhang that must be renegotiated.

Characteristics of systemic risk

The Asian crisis exemplifies certain characteristics that are incorporated in most definitions of systemic risk.

A systemic crisis originates in, or is substantially magnified by, financial markets. More precisely, systemic risk must originate in the *process of financing*. That is, the capital needed by a firm is provided by investors outside the firm. According to this property, systemic risk would not be present if all firms were purely financed internally. This characteristic is clearly present in the Asian crisis. Prior to the crisis, the macroeconomic foundations were strong, with no large shocks from outside the financial system.

A systemic crisis involves contagion. Problems in one country (institution, firm) cause (or appear to cause) insolvency, distress, or bankruptcy in other, *otherwise healthy*, countries (institutions, firms). For example, the initial business failures in the Asian crisis were in Korea and Thailand, but financial markets were roiled almost immediately in Malaysia, Philippines, Indonesia, and Hong Kong, and eventually in Latin America and Russia.

A systemic crisis involves a loss of confidence by investors. Usually this means that investors or financial institutions cut back the amount of liquidity they are willing to provide. Furthermore, the withdrawal of liquidity is *not* due to any objective deterioration in the quality of the borrower firms but is more a function of investor sentiment. This problem was apparent in the rapid reversal of capital flows to East Asia.

A systemic crisis involves substantial real costs in economic output and/or economic efficiency. For example, a stock market crash is not necessarily a systemic crisis. If the crash simply redistributes wealth from speculators on the losing side (those who sold

when prices were low) to the speculators on the winning side (those who bought low) without affecting real economic activity, it would not be a systemic crisis. A crisis must hurt Main Street, not just Wall Street. The Asian crisis is clearly having a huge impact on the well-being of ordinary citizens.

A systemic crisis calls for a policy response. In particular, a systemic crisis must result in a suboptimal economic performance that, at least in principle, could be improved by government action. According to economic theory, government action can be justified if there is an *externality*, whereby individuals do not fully internalize the costs or benefits of their actions. For example, some economists associate systemic crisis with a wave of bankruptcies (see, for example, Feldstein, 1991). If a bankruptcy only affects the individuals who have a contractual relationship with the bankrupt firm (such as the shareholders, creditors, and employees), there is no justification for government action. Presumably, the terms of the original contracts between the firm and these individuals should have taken the possibility of bankruptcy into consideration. However, if a wave of bankruptcies affects individuals with no contractual relationship with the bankrupt firms (say, because fearful investors withdraw funds from healthy firms), then these bankruptcies have economy-wide costs that would not have been considered in the original contracts. In such cases, there may be scope for government action. More generally, if a particular characterization of systemic risk does not imply lost economic efficiency, and therefore the possibility of corrective policy action, it would not be particularly useful for regulatory purposes.

There is one additional characteristic that the Asian crisis has in common with many other systemic crises—*the lack of a clear triggering event*. There were obvious structural problems in the East Asian crisis economies, including a high reliance on external capital, a poorly developed banking system (with an even less well developed regulatory apparatus), cronyism between the financial sector and high-ranking government officials, and a lack of transparency in the workings of financial institutions. However, all of these problems were well known before the crisis, and foreign investors were perfectly willing to provide capital when these economies were performing very well. What happened to shatter their confidence so dramatically and so quickly? If there was a precipitating event, the economic response seems wholly disproportionate to the size of the triggering shock.

Neither the 1929 nor the 1987 stock market crash had an obvious precipitating event (in contrast to the sharp market declines in fall of 1973 and 1980, which were caused by OPEC's announcement of its oil embargo and the fall from power of the Shah of Iran, respectively). Similarly, the 1992 crisis in the European exchange rate mechanism and the recent Mexican crisis did not have clear triggers. In the case of the Asian crisis, foreign exchange traders might argue that there was a clear trigger: a run on the Thai currency. However, a run on the currency is simply a visible expression of a loss of confidence, so the question of what caused the dramatic drop in investor confidence remains.

Why did the crisis happen?

Below, I list some proposed explanations for the Asian crisis in increasing order of plausibility. First, it is tempting to blame "rapacious foreign exchange speculators." (This was the theme of Malaysian Prime Minister Mahathir Mohamad's well-publicized tirade against U.S. financier George Soros.) Certainly, speculative attacks on the currency of a crisis country are often the first visible evidence that a crisis is underway. However, the dramatic falls in foreign exchange values were actually a symptom rather than a cause of the crisis. According to Paul Krugman, "the currency crises were only part of a broader financial crisis, which had very little to do with currencies or even monetary issues per se" (Krugman, 1998). If investors doubt the profitability of investments in a particular country, they will start withdrawing capital. This reversal of capital flows causes foreign exchange rates to collapse. Of course, astute currency speculators may foresee such a reversal and bet against the currency. However, the causality runs *from* the perception of declining profitability *to* the collapse of currency markets, not vice versa.

A second theory, outlined in a recent publication of the Organization for Economic Cooperation and Development,[13] attributes the Asian crisis to "excessive optimism, throughout the region itself and in financial institutions in Europe, North America, and Japan;" "insufficient weighting of downside risks;" a tendency toward "overheating;" and "excessive concentrations of capital in particular areas." These words imply investor irrationality. There is a long history of explaining systemic risk in this way. Kindleberger (1978), Minsky (1982), and Feldstein (1991) attribute systemic risk to an irrational piling-on of debt. According to this explanation, firms become overly optimistic about their debt capacity during expansions. They finance capital investment with ever-increasing leverage ratios. Eventually, this debt level becomes "unsustainable," leading to an inevitable collapse: "Greed overcomes

fear and individual investors take greater risks in pursuit of greater returns. A shock occurs and the market prices of assets begin to collapse. Bankruptcies ... follow. ... The resulting failure of the payments mechanism and the inability to create credit bring on an economic collapse" (Feldstein, 1991).

I regard theories based on investor irrationality to be methodologically unsound. While investors do not *always* behave rationally, it is too easy to ascribe anomalous events to pervasive and systematic mistakes by investors who clearly have every incentive to avoid such mistakes. However, even if one were to accept this idea in principle, the policy implications are not credible: Presumably the government should impose regulations to keep these exuberant firms from over-leveraging. In other words, the firms whose money is at stake cannot determine the proper level of debt, but government bureaucrats can.

A third theory, however, provides an explanation for "excessive debt" that is completely consistent with investor rationality: the problem of *moral hazard* induced by a mispriced government-provided safety net. Moral hazard arises whenever an insured party, by virtue of being insured, fails to take precautions to prevent the event being insured against. In financial markets, moral hazard can arise if investors or bank depositors believe that the government will bail them out if their investments (or the bank) fail. They have less incentive to monitor the firms or banks to ensure that these institutions make prudent decisions. In particular, they lend more than they would in the absence of the government guarantee.

Under this explanation, "systemic risk" is a state of overleveraging due to the existence of the safety net, and "systemic crises" occur when, due to an external adverse shock, there are excessive defaults (relative to the number of defaults one would expect in the absence of the safety net). In the context of the Asian crisis, foreign investors may have believed that they would be protected by Asian governments (and perhaps ultimately by the International Monetary Fund [IMF]) in the event of a crisis. As a result, lenders were less diligent in monitoring their borrowers than they would have been if they did not anticipate a bailout, and the quality of investment projects chosen was poorer with higher levels of risk.

It is easy to demonstrate that an investment project is poor once the project fails. However, for a proper test of the moral hazard hypothesis, one needs to show that, *before* the failure, investors exercised less diligence than they should have in making their investment choices. This is a far more difficult exercise. Corsetti, Pesenti, and Roubini (1998) cite evidence that investment quality was low: A substantial fraction of the new investment was directed toward real estate, as opposed to increased manufacturing capacity. The rate of nonperforming loans before the crisis was above 15 percent in Thailand, Korea, Indonesia, and Malaysia. In 1996, 20 of the 30 largest Korean conglomerates showed a rate of return on invested capital below the cost of capital. While certainly not definitive, these patterns are consistent with a lack of prudence on the part of investors.

The moral hazard hypothesis can explain why rational foreign investors might lend short term to finance long-term risky projects. Furthermore, it is consistent with the view that the Asian crisis was triggered by the Thai government's unwillingness to bail out failed finance companies. However, it still can't account for the rapid reversal in capital flows. Did news casting doubt on implicit guarantees arrive so fast that investors executed a 180-degree reversal in a matter of days? In the standard model, a gradual inflow of adverse news would induce gradual change in the optimal level of investment. This is not what happened in Asia.

This lack of a triggering event commensurate with the scale of the crisis presents a real challenge to theorists. In the standard model, there is a unique "fundamental value" of a firm: the present value of the net cash flow generated by the firm. The market value of the firm should be the best guess, given current information, about this fundamental value. Similarly, the unique fundamental value of a national economy is the present value of the cash flow generated by the total economic activity of that country. A country's currency represents a claim to the current and/or future output of that country's economy. The value of this currency should in some sense reflect the fundamental value of the country's productive capacity. A sudden, precipitous drop in the value of a country's asset markets or currency would then mean that:

- The market's pre-crisis assessment of fundamental value was too high; or
- The market's post-crisis assessment of fundamental value was too low; or
- The crisis was precipitated by the arrival of *new* information that caused rational investors to change their assessment of this fundamental value.

If, as seems to be the case in many financial crises, the third alternative is difficult to sustain, the standard model forces one to choose between the first and second alternatives, both of which imply serious valuation errors by investors. Otherwise, one must abandon the standard model in favor of one in which there is no unique "fundamental" value.

In the next section, I describe a simple model that, unlike the standard model, admits multiple fully rational equilibria, with a *different* fundamental value associated with each equilibrium. I interpret one of these equilibria as a systemic crisis. In principle, a very small impulse can trigger a shift from the *good* equilibrium to the crisis equilibrium, with large consequences for economic performance.

FIGURE 1

Example of decreasing returns to scale

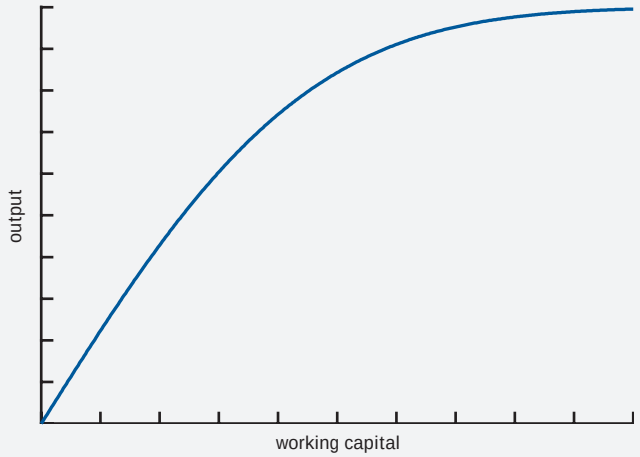

Coordination failure, self-fulfilling prophecies, and systemic risk

A simple model of coordination failure

Here, I propose an approach to systemic risk that builds on Diamond and Dybvig's (1983) idea of coordination failure. In my approach, the fundamental value of a firm or a national economy depends on the state of investor confidence. A formal statement of the model is in technical appendix 1. Suppose a country's economy consists of a large number of identical firms. In the standard model, the output of a firm (or industry, or country) displays *decreasing returns to scale*. That is, as the scale of production grows, the *additional* output resulting from one additional unit of productive inputs (labor, land, capital) decreases. Figure 1 displays these decreasing returns to scale. (For simplicity, the figure has only a single input, which I refer to as *working capital*, denoted L, for "liquidity.") To see why economists make this decreasing returns to scale assumption, consider what would happen if a competitive firm operated where increasing returns prevailed: The firm could always make higher profits by increasing its scale of production, since its revenues would increase faster than its costs. Thus, a profit-maximizing competitive firm would never operate where increasing returns to scale prevail. Furthermore, as long as there are productive inputs such as land or physical capital that are fixed, at least in the short run, then a high enough output level will be in the decreasing returns range.

While this story is plausible, nothing suggests that a firm's production function must have decreasing returns to scale for *all* levels of inputs. In figure 2, there are increasing returns to the scale of production as long as working capital is less than the level designated L^*. Below this level, efficiencies can be captured by increasing the scale of production. For working capital greater than L^*, returns to scale decrease, as in figure 1.

For the simple model, I use a rather extreme version of increasing returns followed by decreasing returns, as displayed in figure 3. L^* is a fixed set-up cost for the firm. If working capital is less than L^*, the firm cannot be completely set up, so production equals zero. If working capital exceeds L^*, the firm produces according to a production function similar to that in figure 1. (All the intuitive results go through with the more general function in figure 2. I use the production function in figure 3 for convenience.)

FIGURE 2

Example of increasing, then decreasing returns to scale

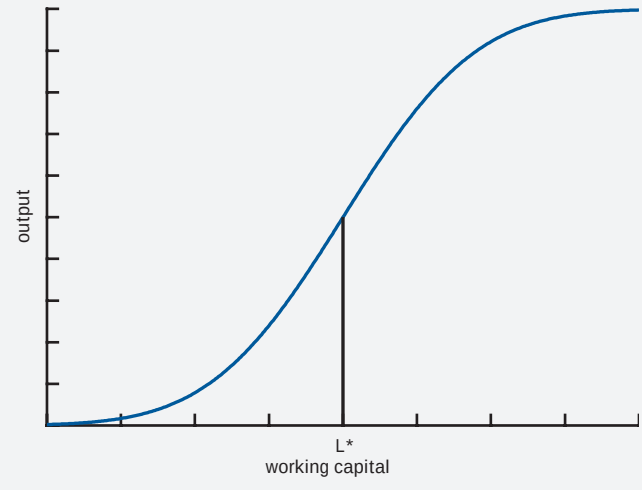

FIGURE 3

Example of threshold working capital requirement

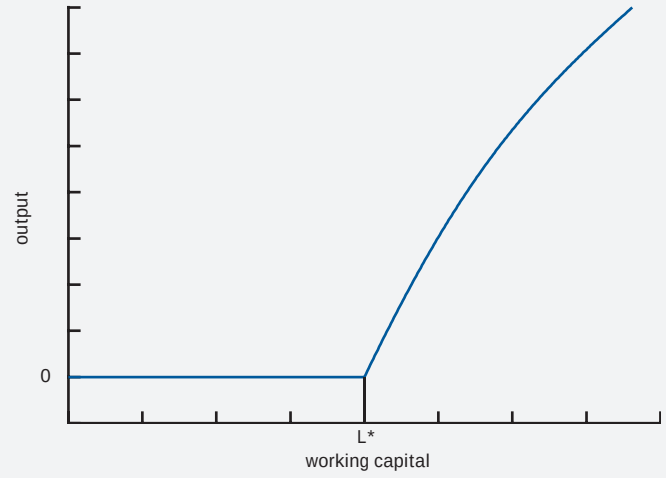

Now, I make two critical assumptions. The first is that inputs must be financed *externally.*[14] That is, the firm cannot provide its own working capital but must borrow from investors. As is typical in models of this type, the return to lenders (denoted r) equals the marginal product of working capital (the additional output produced by one additional unit of working capital). The marginal product corresponding to the production function in figure 3 is given in figure 4. Let us assume that investors have a choice of assets. They can lend working capital to the firm described by figures 3 and 4 or they can invest in a safe low-return investment, paying a gross return of 1.0. This return is the black line in figure 4. I also assume that, for any feasible level of working capital greater than L^*, the return r exceeds 1.0. (Note that the marginal product schedule in figure 4 lies above the black line whenever working capital exceeds L^*.) If total working capital per firm exceeds L^*, then every investor would prefer to invest in working capital loans rather than investing in the low-return asset. Of course, if total working capital per firm is less than L^*, the return to working capital loans is zero, so the low-return asset is preferable.

FIGURE 4

Marginal product associated with figure 3

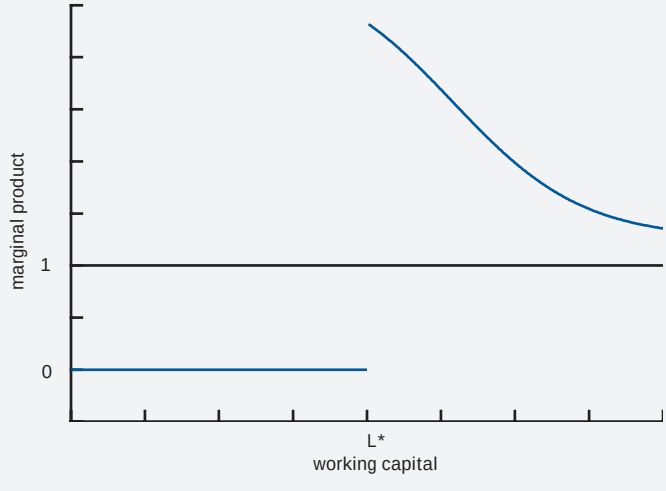

The second critical assumption gives rise to the possibility of coordination failure. I assume that the minimum financing requirement of any given firm is too big for any individual investor to provide. It requires a large number of investors. What does that imply for an individual investor? If an investor is confident that a sufficient number of other individual investors will lend sufficient working capital to the firm to meet its minimum requirement, it would be optimal for the investor in question to lend as well. If all investors think this way, the aggregate working capital will exceed the minimum requirement, the return on the investment will be greater than 1.0, and the investors' investment decisions are indeed justified. This is the high output equilibrium. In this equilibrium, the optimal strategy for all investors is to put all their investable wealth in working capital for the firm, receiving a return r that exceeds 1.0. However, there is another equilibrium with an output of zero. If the individual investor believes that the number of other investors is insufficient to provide the threshold level of working capital, it is optimal for the investor to *refrain entirely* from investing in working capital, and to use the alternative, low-return investment. If all investors think this way, liquidity equals zero, so the return to the lender also equals zero, and, again, the investors' investment decisions are justified *ex post*. Note that both equilibria are fully rational: in neither case do investors have any incentive to deviate from their chosen investment strategy.

To establish the first equilibrium, (the high-output equilibrium) all that is needed is to have enough investors

agree to invest all their wealth in the firm. If this could be done, then the market return on working capital loans would exceed unity, and the remaining investors would also invest in the firm. This second equilibrium corresponds to a systemic crisis. It represents a failure of coordination, since all investors would benefit if they could coordinate their actions and invest in the firm. The only explanation for the low-output result is that each investor believes that *other* investors do not want to invest. In other words, there is a pervasive lack of confidence in the ability of the financial markets to provide sufficient capital to the firm. This state of affairs recalls Keynes's characterization of the stock market as "betting on a beauty contest." What is important is not the state of the firm's technology. (In this example, the firm's technology is the same in both equilibria and is known to all investors.) Rather, what is important is the investor's beliefs about what the other investors will do. Note that, in this model, there are no "bubbles," whereby market values deviate from fundamental value. Rather, the true fundamental value of the firm *depends* on the state of investor confidence.

To apply this model to the Asian crisis, I interpret the "firm" in the model as a developing country in need of dollar working capital loans, the "investors" as large western commercial banks, and the alternative low-return investment as western treasury securities. In the high-output equilibrium, sufficient working capital is provided to the developing country for it to produce at a high level. In the crisis equilibrium, western investors refuse to provide working capital, and the output of the developing country falls precipitously.[15] Under this interpretation, the model can capture the sudden exchange rate disruption that typically accompanies international systemic crises like the recent Asian crisis. Assuming that purchasing power parity holds and both the money supply and velocity are constant in both countries, it is easy to show that the exchange rate is proportional to the ratio of the real output of the developing country to the real output of the western country. The "crisis" equilibrium (where the output of the developing country falls to zero) is then associated with a precipitous fall in the exchange rate.

This model can be interpreted as representing other types of systemic risk. If the firms are banks, the crisis equilibrium is one where households withdraw liquidity from the banking system. This does not necessarily correspond to a bank run in Diamond and Dybvig's (1983) sense. If the depositors withdraw liquidity from the banking system (as in the Great Depression, when the money multiplier fell to a historically low value), banks need not fail as a result. They may be able to accommodate the withdrawals by calling existing loans and refusing to make additional loans. Nonetheless, this economy-wide flight to currency will starve the productive sector of needed liquidity, and the resulting business failures eventually will take their toll on the banking sector.

One can also interpret the model as a payment system. In this interpretation, the threshold liquidity level L^* for an institution is the amount of liquidity it needs to avoid default and possible bankruptcy. Suppose an institution counts on receiving cash flow during the day in excess of L^* and, due to a disruption in the payments system, it does not receive this cash flow. Other institutions act as the "investors" in the model. They will only provide the needed loans to the stricken institution if they believe the total loan provision will exceed L^*. Many specialist firms on the New York Stock Exchange were in this position the day after the October 1987 crash. Specialist firms had purchased stock during the crash (performing their designated role as market maker). The amount of credit they needed to settle these trades at the end of the day corresponds to L^* in the model. Banks only provided this credit if they thought that enough other banks were participating to insure that the specialist firm would not end up in bankruptcy court, with the attendant costs and delays. Coordination failure in this case was avoided by the active intervention of the government and the Federal Reserve.

Modeling the stock market: A dynamic model

The simple model in the previous section did not account for the declines in asset markets that were a prominent part of the Asian crisis. To incorporate a stock market into the simple model, I incorporate the simple model into a dynamic economic equilibrium. A formal statement of this model is in technical appendix 2. In this model, stock shares are thought of as claims to the firm's physical capital. The price of a claim (that is, the price per share) must represent the present value of the future dividends accruing to the owner of the claim. However, the dividend in any given period depends on the amount of working capital provided by external investors. That is, it depends on *which* equilibrium prevails that period. As a result, the *current* stock price is a function of the probability that a systemic crisis (the *bad* equilibrium) will occur in the *future*. If this probability increases, the stock price falls.[16] Furthermore, if the future systemic crisis in question is expected in the near term, the stock price falls by more than if this

crisis is expected in the more remote future. In other words, current stock values forecast the probability of future systemic crises. Note that the declines in Asian asset markets started before the actual onset of the crisis. (See Krugman, 1998.)

Contagion

The model developed thus far captures a number of the key features of systemic risk. Systemic crises are associated with a sudden withdrawal of liquidity (working capital) from productive firms due to a loss of confidence that seems unrelated to the state of technology. After the fact, this withdrawal of working capital loans seems to indicate an over-reliance on debt, in the sense that a firm would avoid this problem if it could avoid dependence on outside lenders for working capital. There are real costs in economic efficiency to the crisis of confidence associated with the low-output state, suggesting a possible role for government action to bolster investor confidence. Finally, agents in the model are completely rational. Every agent in the economy acts optimally, given the actions of the other agents and the inability to coordinate.

However, the one element that has not been incorporated into this model is contagion: the tendency for the failure of one institution to precipitate a crisis of confidence that harms unrelated firms and leads to reduced economic activity. Most discussions of contagion involve "cascading defaults." As well as having direct exposure to the credit-worthiness of its direct counterparties, an institution has indirect exposure to the credit-worthiness of its counterparties' counterparties and so on. As a result, it is vulnerable to disruptions caused by the failure of institutions to which it has no direct exposure. However, properly understood, this notion of contagion is ordinary credit risk. "Know thy counterparty" requires an assessment of all factors that could impair the counterparty's credit-worthiness, including its exposure to other firms' distress. The fact that a bank may have incomplete information about its counterparties' exposure to other institutions does not change this conclusion. Incomplete information is always associated with risk. Institutions should manage the risk associated with these information gaps in the same way they manage other types of risk.

Here, I consider an alternative notion of contagion: The failure of a firm provides information about the state of investor confidence. As defaults are observed, the investor community cannot tell whether these are purely due to company-specific weaknesses or to individual investors pulling the plug. The default thus may give information that the economy is ready to shift to the bad equilibrium.

To formalize this idea of contagion, I put the simple model developed earlier into a richer context. Recall that the critical determinant of whether a systemic crisis occurs is L, the aggregate level of working capital. In the simple model, L is observed by all investors. Suppose instead that L is not observable. Investors then must decide on how much working capital to provide based on their best estimate of the state of economy-wide confidence, as summarized by L. Any information that causes investors to lower this estimate will increase the estimated probability of systemic crisis and decrease the amount of working capital provided by investors.

One important piece of information with an obvious relevance to estimating L is the failure of an individual firm. (Recall that the Asian crisis was preceded by failures of Korean chaebols and Thai finance companies.) In the context of the simple model, a firm "fails" if its working capital falls short of L^*, the threshold level. Suppose each firm has a different threshold level of working capital. I refer to a firm's threshold working capital requirement as its "technology." Firm technology is not observable to investors. All that investors know is the *distribution* of threshold levels across firms. Since all firms look alike to investors, each firm simply receives the economy-wide average working capital level. A firm can fail for one of two reasons: either it has an unusually high threshold working capital requirement or the economy-wide level of working capital is unusually low. If the former explanation is true, the bankruptcy has no implications for future systemic crises. The problem is idiosyncratic to the firm and does not reflect a weakness in investor confidence. If, however, the latter explanation is true, the bankruptcy is evidence that investor confidence is low. The probability of systemic crisis is now higher. Unfortunately, our investor cannot tell which explanation holds. The investor must solve a *signal extraction problem*, attempting to glean information about the state of aggregate investor confidence from the ambiguous news that a particular firm has failed.

In technical appendix 3, I formalize this intuition. I show how the observation of a firm bankruptcy causes investors to lower their estimate of aggregate investor confidence (as measured by L). If, as is plausible, investors' willingness to roll over short-term loans depends on this estimate, the provision of short-term capital will indeed fall following a bankruptcy. However, the amount of this fall depends critically on the investor's beliefs *before* the bankruptcy is observed. If investors believe that aggregate confidence is strong, observing a single bankruptcy does little to change this opinion. If, however, investors

are less certain that *other* investors will provide sufficient working capital to firms, a single bankruptcy could cause investors to substantially revise their beliefs about the magnitude of aggregate investor confidence. In principle, this could induce investors to reduce the amount of working capital provided to firms and induce nonproductivity ("bankruptcy") in firms *that would have been solvent* had no bankruptcy been observed. That is, a single bankruptcy could tip the economy into systemic crisis.

Note that the source of contagion is *not* that one firm's bankruptcy directly puts pressure on other firms (as in the cascading defaults story). Neither is the source of contagion that a bankruptcy conveys information about the state of other firms' technology. Rather, the bankruptcy gives evidence about the beliefs of the other investors. A bankruptcy tells investors that *other* investors may be less confident than previously thought. Applying this logic to the Asian crisis, the failure of Korean chaebols and Thai finance companies may have induced declines in Indonesian equity and currency markets not because it gave information about the state of the Indonesian economy per se, but because it gave information about the willingness of western investors to lend to Asian emerging economies. In equilibrium, of course, the willingness to lend has profound implications for the performance of these economies.

Conclusion

The framework I have sketched for understanding systemic risk presents a difficult task for policymakers. The framework suggests that systemic crises of the Asian type must be seen as a problem of coordination failure among investors. That is, it is optimal for individual investors to withhold liquidity from the stricken countries, yet all investors would be better off if *everyone* provided this liquidity. One way around this coordination failure is to induce a single large investor to provide the needed liquidity. This was the solution to the Mexican crisis of 1995. The single large investor, of course, was the U.S. government. The Mexican government owed $28 billion in dollar-denominated notes (*tesobonos*), but had only $10 billion in dollar reserves. Thus, the Mexican government was illiquid. Did this indicate insolvency? Clearly not, since the $28 billion represented only 10 percent of the Mexican GDP. There was enough potential economic output to service this debt. Yet, no private creditors would roll over the debt. The logic is precisely that of coordination failure. The solution was for the U.S. government to take over the debt, thereby eliminating the problem of coordination. The U.S. government was repaid on schedule.

A somewhat similar solution is to have a single large private investor provide the needed liquidity to individual firms in the crisis countries. If a viable firm is starved for liquidity and, as a result, has its equity underpriced, this should represent a buying opportunity for some large investor. We see some evidence that this is happening in Asia. GE Capital, an American non-bank financial company, has been taking advantage of low Asian equity values to acquire Asian companies at fire-sale prices. As of early June 1998, its acquisitions include Japanese and Philippine life insurance firms and several Thai finance companies. One analyst forecasts that GE Capital will spend at least $20 billion on Asian acquisitions over the next three years.[17] Presumably, GE Capital has access to the needed working capital to make these acquisitions productive. Policies that help facilitate acquisition of troubled firms by large, cash-rich institutions could, in principle, offset the problem of coordination failure.

However, these mechanisms cannot reliably insure against the possibility of systemic crisis. Perhaps the international community should consider something more formal. Within a domestic economy, the way out of the coordination dilemma is to designate a government institution (usually the central bank) as backup provider of liquidity in the event of a crisis. To control the problem of moral hazard, this institution is given broad regulatory powers over the domestic financial services industry. (In the U.S. banking industry, the problem of moral hazard is controlled by imposing capital adequacy requirements and requiring prompt corrective action should a bank become undercapitalized.)

However, no such institution exists for international transactions. Perhaps the legacy of the Asian crisis will be to encourage the international community to consider establishing a backup provider of liquidity for international liabilities. Such an institution would have to be given regulatory oversight, with the capacity to set and monitor capital levels of borrowing institutions and the ability to take prompt corrective action.[18] This new institution would set up a source of liquidity ex ante (well before any crisis), available to countries that agree to suitable financial reforms including greater transparency. The institution would need to be able to commit credibly *not* to intervene in a crisis *unless* it had established a prior arrangement with the affected country granting it the needed regulatory oversight. Note that the IMF is not structured in this way. It was created as a source of liquidity to help countries adhere to the fixed exchange rates mandated by the Bretton Woods agreements of 1944. Since the demise of Bretton Woods, the IMF has functioned as

a source of foreign exchange to be loaned to troubled countries on an ad hoc basis. It has no *ex ante* regulatory authority, and it does not act according to preset predictable rules.

Why would any sovereign country voluntarily cede regulatory authority to an international body? Quite simply, to reassure potential lenders that a credible lender of last resort in foreign reserves stands ready to step in should a systemic crisis erupt.[19] A credible structure of this kind would rule out the crisis equilibrium. It would also make it difficult for an emerging economy to secure international loans without committing to the institution's regulatory oversight.

TECHNICAL APPENDIX 1

A simple, static model

There are two periods. Agents have perfect foresight. There are N identical firms, which behave competitively, and there are N identical investors (where N is a large number). Each investor is endowed in period 1 with y units of wealth. In this model, it is important that a firm cannot be productive unless many different investors invest in that firm. For simplicity, I assume that investors cannot invest in individual firms. Rather, all invested funds from the N investors are pooled and divided among the N firms equally. In the following, all quantities are expressed as per-household or per-firm magnitudes.

Production

A firm has a fixed quantity of physical capital K. It needs working capital in order to produce. It is assumed that the firm has no internal sources of working capital, so it must borrow working capital from investors.

Output = $F(K,L)$, where

K = physical capital per firm

L = working capital ("liquidity") per firm.

It is assumed that there is a threshold amount of working capital needed to make the firm productive. This amount, denoted L^*, must be paid up front. If the threshold is not met, the firm produces zero output. That is, F takes the form:

$$1) \quad F(K,L) = \begin{cases} 0 \text{ if } L < L^* \\ f(K, L - L^*) \text{ otherwise,} \end{cases}$$

where f is a non-increasing-returns-to-scale production function satisfying

$$2) \quad f(K, y - L^*) - f_L(K, y - L^*)L > 0.$$

Let r denote the rental rate on working capital. The first-order condition of the firm's profit maximization problem implies:

$$3) \quad r = F_L(K,L) = \begin{cases} 0 \text{ if } L < L^* \\ f_L(K, L - L^*) \text{ otherwise.} \end{cases}$$

To ensure that r always exceeds unity (the return to "storage," described below) whenever $L \geq L^*$, I assume that

$$4) \quad \min_{x \in [0, y - L^*]} f_L(K,x) > 1.$$

Investors

Each investor has two investment options: 1) lend wealth to the firm to provide working capital, receiving a gross rate of return r; or 2) invest in a storage technology that pays a gross return of unity.

The investor only consumes in the second period. The consumption of the ith investor (denoted c_i) satisfies

$$c_i = rL_i + y - L_i,$$

where L_i denotes the amount of wealth invested by investor i in working capital loans. Assume that $L^* < y$, so there is sufficient wealth in the economy to pay the firm's threshold capital requirement. The ith investor solves the utility maximization problem

$$\max_{L_i \in [0, y]} U(rL_i + y - L_i).$$

The first-order conditions are:

$$U'(rL_i + y - L_i)\cdot[r-1]\begin{cases} = 0 \text{ if } 0 < L_i < y \\ \le 0 \text{ if } L_i = 0 \\ \ge 0 \text{ if } L_i = y \end{cases}.$$

Equilibrium

If N is sufficiently large, there are two equilibria:

1. $L_i = y, \forall i, r = f_L(K, y - L^*) > 1;$
2. $L_i = 0, \forall i, r = 0.$

The first is the high-output equilibrium. All investors, seeing a market return to working capital loans that exceeds the alternative return of unity, invest all their wealth in loans to the firm. The output per investor is $f(K, y - L^*) > y$. The second equilibrium is the low-output equilibrium. All investors, seeing a market return to working capital loans equal to zero, invest all their wealth in the storage technology that pays a gross return of unity. The output per investor is simply y (the output of the storage technology).

TECHNICAL APPENDIX 2

A dynamic model

The basic model in technical appendix 1 is incorporated into an overlapping generations structure. Each investor lives two periods and consumes in the second period only. In the first period of life, the investor receives an endowment of y units of wealth and can invest this wealth in three types of assets: 1) lend wealth to the firm to provide working capital in the current period, receiving a gross rate of return r (paid in the *following* period); 2) invest in shares of stock, which give the investor a claim to the profits of the firm in the following period; or 3) invest in a storage technology that pays a gross return of unity in the following period. The first period budget constraint is:

5) $L_t + p_t e_t + s_t = y,$

where e_t = number of shares of stock purchased at date t by the typical young investor, p_t = stock price, s_t = wealth invested in the storage technology at date t by the typical young investor, and L_t = working capital provided at date t by the typical young investor.

In the second period of life, the investor receives: 1) the earnings from working capital investments the previous period; plus 2) the dividends from her stockholding; plus 3) the proceeds from selling her stock to the current young investors at the current market price; plus 4) whatever was invested in the storage technology. The second period budget constraint is:

6) $c_{t+1} \le e_t d_{t+1} + L_t r_t + p_{t+1} e_t + s_t,$

where r_t is the return to working capital loans at date t, and d_t is the dividend per share at date t, which is given by:

7) $d_t = F(K, L_t) - r_t L_t.$

Substituting equation 5 into equation 6, evaluating both at equality, one obtains

8) $c_{t+1} = e_t(d_{t+1} + p_{t+1} - p_t) + L_t(r_t - 1) + y$

For simplicity, assume that the investor is risk neutral. The investor's problem is then:

9) $\max_{L_t, e_t} E_t[c_{t+1}]$

subject to equation 8 and the following short-sale restrictions:

10) $L_t \ge 0;\ e_t \ge 0;\ s_t \ge 0.$

The total number of shares of stock per investor is normalized to unity, so

11) $e_t = 1.$

The Kuhn–Tucker conditions for the maximization problem (equations 8–10), along with the market clearing condition (equation 11) imply the following characterization of the stock price p_t:

1. If $L_t > 0$, the economy is in the high-output state, and

12) $p_t = \dfrac{E_t(d_{t+1} + p_{t+1})}{r_t}.$

2. If $L_t = 0$, the economy is in a state of systemic crisis, and

$$13)\quad p_t = \begin{cases} E_t(d_{t+1} + p_{t+1}) & \text{if } s_t > 0 \\ y & \text{if } s_t = 0 \end{cases}.$$

Equation 12 and the top part of equation 13 simply say that the current stock price equals the discounted future stock payoff per share $d_{t+1} + p_{t+1}$ (although the discount rate depends on whether L_t equals or exceeds zero). The case represented by the bottom part of equation 13 must be considered because the short-sale restrictions on s_t and L_t preclude a levered stock portfolio.

Equations 12 and 13 can be solved forward recursively to imply that the current stock price p_t is the present discounted value of the future dividend stream:

$$14)\quad p_t = E_t\left\{\sum_{\tau=t}^{\tilde{T}-1} \pi_t(\tau+1)\frac{E_t\left[d_{\tau+1} \mid L_{\tau+1} > 0\right]}{\prod_{j=0}^{\tau-t}\tilde{r}_{t+j}} + \frac{y}{\prod_{j=0}^{\tilde{T}-t}\tilde{r}_{t+j}}\right\},$$

where I use the following notation:

$$\pi_t(\tau+1) = \text{prob}(L_{\tau+1} > 0 \mid \text{information at date } t)$$

$$\tilde{r}_{\tau+1} = \begin{cases} r_{\tau+1} \text{ if } L_{\tau+1} > 0 \\ 1 \text{ otherwise} \end{cases}$$

$\tilde{T}$ denotes the first date after t when both $L_t = 0$ and $s_t = 0$.

Note the dependence in equation 14 of p_t on $\pi_t(\tau)$, for $\tau > t$. If investors' assessment of the probability of a future crisis increases (that is, $\pi_t(\tau)$ goes down for some future dates τ), then the stock price p_t declines. This effect is more pronounced if the crisis is expected in the near term, since $\pi_t(\tau)$ is discounted by the prevailing short-term investment rate compounded from t to τ. The stock market thus becomes an early warning signal of the market consensus probability of future crises of confidence.

TECHNICAL APPENDIX 3

An example of contagion

The example below illustrates how the failure of an individual firm can change investors' estimates of the aggregate confidence level of the economy. In this example, there is a continuum of firms, indexed by j. The firms differ in their threshold working capital requirement, denoted L_j^*. A firm's value of L_j^* is unknown to the investors. As a result, investors treat all firms alike, and the same amount of working capital (denoted L) is provided to each firm. Variable L summarizes the state of aggregate investor confidence. Finally, it is common knowledge that the cross-sectional distribution of L_j^* is uniform over $[0, \overline{\theta}]$, where $\overline{\theta}$ is a strictly positive parameter.

At the beginning of the first period, the ith investor has a prior belief that L has a uniform distribution on $[0,\lambda_i]$, where λ_i is a positive parameter. In principle, λ_i can differ across investors. After L is determined, but before production takes place, a single firm is selected at random to be audited. If $L < L_j^*$ for this firm, the auditor makes a public announcement that a bankruptcy has occurred. Let B denote an indicator variable for the event of bankruptcy, so this announcement corresponds to $B = 1$. If $L \geq \theta_j$, the auditor announces that no bankruptcy has occurred, in which case $B = 0$. The investors then update their prior beliefs about L using the realization of B. In a fully specified version of this model, the investor would be permitted to withdraw working capital from the firm after observing B, but before production takes place. In this way, the realization of B could affect the actual amount of working capital available to firms. However, here I only discuss how investors update their beliefs about L after observing the realization of B. The ith investor's prior distribution on L is:

$$\text{prob}_i(L < \tilde{L}) = \frac{\tilde{L}}{\lambda_i},$$

so this investor's expected value of L is $\frac{1}{2}\lambda_i$. Using Bayes's Rule, one can derive the corresponding posterior distribution conditional on observing B:

$$\text{prob}_i(L < \tilde{L} \mid B = 1) = \begin{cases} \dfrac{\tilde{L}}{\lambda_i}\left(\dfrac{2\bar{\theta} - \tilde{L}}{2\bar{\theta} - \lambda_i}\right), \text{ if } \lambda \le \bar{\theta}, \\ \dfrac{\tilde{L}(2\bar{\theta} - \tilde{L})}{\bar{\theta}^2}, \text{ if } \lambda > \bar{\theta} \text{ and } \tilde{L} < \bar{\theta}, \\ 1, \text{ otherwise.} \end{cases}$$

$$\text{prob}_i(L < \tilde{L} \mid B = 0) = \begin{cases} \left(\dfrac{\tilde{L}}{\lambda_i}\right)^2, \text{ if } \lambda \le \bar{\theta}, \\ \dfrac{\tilde{L}^2}{\bar{\theta}(2\lambda_i - \bar{\theta})}, \text{ if } \lambda > \bar{\theta} \text{ and } \tilde{L} < \bar{\theta}, \\ \dfrac{2\tilde{L} - \bar{\theta}}{2\lambda_i - \bar{\theta}}, \text{ otherwise.} \end{cases}$$

Note the chain of first-order stochastic dominance: For any L,

$$\text{prob}(L < \tilde{L} \mid B = 1) > \text{prob}(L < \tilde{L}) > \text{prob}(L < \tilde{L} \mid B = 0).$$

That is, the event of a bankruptcy shifts probability mass to the lower tail of the distribution of L, while the event of no bankruptcy shifts probability mass to the upper tail. Furthermore, the amount of this shift in probability mass is non-decreasing in λ_i: a higher value of λ_i implies a bigger revision of the prior distribution when B is revealed.

The ith investor's expected value of L, conditional on B, is:

$$E_i(L \mid B = 1) = \begin{cases} \dfrac{\lambda_i(3\bar{\theta} - 2\lambda_i)}{3(2\bar{\theta} - \lambda_i)}, \text{ if } \lambda_i \le \bar{\theta}, \\ \dfrac{\bar{\theta}}{3}, \text{ if } \lambda_i > \bar{\theta}. \end{cases}$$

$$E_i(L \mid B = 0) = \begin{cases} \dfrac{2}{3}\lambda_i, \text{ if } \lambda_i \le \bar{\theta}, \\ \dfrac{\lambda_i^2 - (1/3)\bar{\theta}^2}{2\lambda_i - \bar{\theta}}, \text{ if } \lambda_i > \bar{\theta}. \end{cases}$$

As expected, the mean value of L conditional on $B = 0$ exceeds the unconditional mean, which in turn exceeds the mean conditional on $B = 1$. That is, if no bankruptcy occurs in period 1, investors increase their estimate of the level of investor confidence. If a bankruptcy is reported in period 1, they decrease their estimate of the confidence level.

NOTES

[1]Recently, several efforts have been made to catalogue the various proposed characterizations of systemic risk, along with the related (and equally ambiguous) notions of "financial fragility" and "financial bubbles." Among these are Davis (1992), a collection of essays edited by Kaufman (1995), and a special issue of the *Journal of Financial Services Research* (Benston et al. 1995).

[2]See Minsky (1982), Kindleberger (1978), and Feldstein (1991).

[3]See Benston and Kaufman (1995), Krugman (1998), and Corsetti, Pesenti, and Roubini (1998).

[4]See Bank for International Settlements (1990) and Eisenbeis (1995).

[5]See U.S. General Accounting Office (1994).

[6]Schwartz (1986)

[7]See Benston and Kaufman (1995), Meltzer (1982), and Crockett (1995).

[8]Diamond and Dybvig (1983).

[9]See Radelet and Sachs (1998a,b) and Corsetti, Pesenti, and Roubini (1998). A detailed chronology of the events in Asia can be found on the Web site maintained by Nouriel Roubini, Internet at http://www.stern.nyu.edu/~nroubini/asia.

[10]See Cline and Barnes (1997).

[11]See Radelet and Sachs (1998a).

[12]See Institute of International Finance (1998).

[13]Organization for Economic Cooperation and Development (1998), p. 9.

[14]One could weaken this assumption to "a substantial fraction of inputs requires external financing."

[15]This working capital shortage in the Asian crisis was profound. For example, the *Financial Times* of March 11, 1998, reported that the Indonesian poultry industry was in dire straits due to an insufficiency of foreign exchange needed to purchase imported chicken feed.

[16]A formal statement of this relationship is in equation 14 of technical appendix 2.

[17]See *The Economist* (1998), pp. 72–77.

[18]A proposal along these lines was recently made by Mishkin (1998). Similar proposals were made by Marcus Noland in his statement before House International Relations Committee Subcommittees on Asian and Pacific Affairs and International Economic Policy and Trade on February 3, 1998, and by Soros (1997).

[19]There is an analogy with countries such as Argentina that establish currency boards, thereby relinquishing the right to conduct monetary policy. They are willing to do so in order to reassure investors that the government won't expropriate earnings via high inflation.

REFERENCES

Bank for International Settlements, 1990, "Report of the Committee on Interbank Netting Schemes of the Central Banks of the Group of Ten Countries," Basle, Switzerland.

Benston, George J., Franklin R. Edwards, Robert A. Eisenbeis, Paul M. Horvitz, and George G. Kaufman (eds.), 1995, *Journal of Financial Services Research*, Vol. 9, No. 3/4.

Benston, George J., and George G. Kaufman, 1995, "Is the banking and payments system fragile?," *Journal of Financial Services Research*, Vol. 9, December, pp. 209–240.

Cline, William, and Kevin Barnes, 1997, "Spreads and risks in emerging markets lending," Institute of International Finance, research paper, No. 97-1, November.

Corsetti, Giancarlo, Paolo Pesenti, and Nouriel Roubini, 1998, "What caused the Asian currency and financial crisis?," New York University, Stern School of Business, working paper.

Crockett, Andrew, 1995, "Financial fragility: Sources, prevention, and treatment," in *Coping with Financial Fragility and Systemic Risk*, Harald A. Benink (ed.), Boston: Kluwer Academic Publishers.

Davis, E.P., 1992, *Debt, Financial Fragility and Systemic Risk*, Oxford: Oxford University Press.

Diamond, Douglas, and Philip Dybvig, 1983, "Bank runs, liquidity, and deposit insurance," *Journal of Political Economy*, Vol. 91, June, pp. 401–419.

Economist Newspaper Limited, The, 1998, "Vulture: GE Capital in Asia," *The Economist*, June 6, pp. 72–74.

Eisenbeis, Robert A., 1995, "Private sector solutions to payments system fragility," in *Coping with Financial Fragility and Systemic Risk*, Harald A. Benink (ed.), Boston: Kluwer Academic Publishers.

Feldstein, Martin, 1991, *The Risk of Economic Crisis*, National Bureau of Economic Research Conference Report, Chicago: University of Chicago Press.

Institute of International Finance, Inc., 1998, "Capital flows to emerging market economies," report, January 29.

Kaufman, George G., (ed.), 1995, *Research in Financial Services*, Vol. 7, Greenwich: JAI Press, pp. 47–52.

Kindleberger, Charles P., 1978, *Manias, Panics, and Crashes*, MacMillan: London.

Krugman, Paul, 1998, "What happened to Asia?," Massachusetts Institute of Technology, working paper.

Meltzer, Allan H., 1982, "Rational expectations, risk, uncertainty, and market responses," in *Crisis in the Economic and Financial Structure*, P. Watchel (ed.), Salomon Bros. Series on Financial Institutions and Markets, Lexington, MA: Lexington Books.

Minsky, Hyman P., 1982, "The financial instability hypothesis: Capitalist processes and the behavior of the economy," in *Financial Crises: Theory, History, and Policy*, Charles Kindleberger and Jean-Pierre Laffargue (eds.), Cambridge: Cambridge University Press, pp. 13–38.

Mishkin, Frederic S., 1998, "Systemic risk, moral hazard, and the international lender of last resort," Columbia University, Graduate School of Business, working paper.

Organization for Economic Cooperation and Development, 1998, *OECD Economic Outlook*, Vol. 63, June, p. 9.

Radelet, Steven, and Jeffrey Sachs, 1998a, "The onset of the East Asian financial crisis," Harvard University, Harvard Institute for International Development, working paper.

_____, 1998b, "The East Asian financial crisis: Diagnosis, remedies, prospects," Harvard University, Harvard Institute for International Development, working paper.

Schwartz, A., 1986, "Real and pseudo-financial crises," in *Financial Crises and the World Banking System*, Forrest Capie and Geoffrey E. Wood (eds.), London: Macmillan, pp. 11–31.

Soros, George, 1997, "Avoiding a breakdown," *Financial Times*, December 31.

U.S. General Accounting Office, 1994, "Financial derivatives: Actions needed to protect the financial system," Report to congressional requesters, No. GAO/GGD-94-133, May.