David Romer

July 2016

Preliminary – please do not circulate!

Chapter 10

FINANCIAL MARKETS AND FINANCIAL CRISES

The previous two chapters study the behavior of households and firms in partial-equilibrium settings. In Chapter 8, households divide their income between consumption and saving taking the set of available assets and the distribution of their rates of returns as given. In Chapter 9, firms decide how much investment to undertake taking how future profits are valued as given.

Financial markets are where these saving and investment decisions meet. In the absence of asymmetric information, externalities, and other imperfections, they play a central role in getting the economy to its Arrow-Debreu outcome. The signals sent by asset prices and state-contingent returns are what frame the partial-equilibrium problems that households and firms face. They therefore determine how households allocate their resources among consumption and holdings of various risky assets and what investment projects are undertaken. And it is the interaction of the demand and supply of risky assets that determines their prices. General equilibrium occurs when households and firms are optimizing taking prices as given, and where prices cause asset markets to clear. Section 10.1 presents a model of perfectly functioning financial markets in general equilibrium that shows this interplay between saving and investment decisions.

The main reason that economists are so interested in financial markets is that they do not appear to function in this idealized way. There are at least four distinct issues.

The first is whether there are important macroeconomic propagation mechanisms operating through financial markets. With perfect financial markets, asset prices passively summarize all available information. But if there are imperfections in financial markets that cause departures from first-best outcomes, those distortions may change endogenously in response to economic developments. As a result, they can magnify the macroeconomic effects of various types of shocks to the economy.

Sections 10.2 and 10.3 investigate this idea. Section 10.2 presents a microeconomic model of investment in the presence of asymmetric information between outside investors and entrepreneurs and examines the determinants and microeconomic effects of the resulting distortions. It then shows that there are several channels that cause those distortions to be greater when the economy is weaker. Since the distortions reduce investment, this means that their endogenous response to the state of the economy magnifies the macroeconomic effects of shocks—a mechanism that is known as the *financial accelerator*. Then in Section 10.3, we will examine some of the microeconomic evidence about the importance of financial-market imperfections to investment and the macroeconomic evidence about the financial accelerator.

The second issue concerns whether departures of financial markets from the perfect case of Section 10.1 can not only magnify the effects of other disturbances, but be an independent source of shocks to the economy. In particular, Section 10.4 is devoted to the issue of possible *excess volatility* of asset prices. In perfect financial markets, the price of any asset is the rational expectation given the available information of the present value of the asset's future payoffs using the stochastic discount factor that arises from agents' marginal utilities of consumption; and the price of the asset changes only if there is new information about its payoffs or about the stochastic discount factor. Section 10.4 examines the possibility that this assumption might fail. It shows that the forces pushing asset prices toward fundamentals if they depart are not unlimited, and analyzes several factors that limit their strength. It also describes how movements in asset prices not driven by fundamentals can affect macroeconomic outcomes.

A third macroeconomic subject raised by financial markets and the possibility of financial-market imperfections is financial crises. One might expect that a large financial system in an economy with millions of participants would change smoothly in response to economic developments. In fact, however, they are subject to convulsions not only at the level of individual assets and financial institutions, but at the level of broad swaths of the financial system. One notable episode occurred in the Great Depression, when the economic downturn and repeated waves of panics led to the failure of one-third of U.S. banks. For decades, many believed that such worldwide financial crises were a thing of the past. But in the fall of 2008, we learned that this optimistic view was wrong. Lehman Brothers, a major investment bank, declared bankruptcy in September. In the aftermath, equity prices fell by more than 25 percent in just four weeks; spreads between interest rates on conventional but slightly risky loans and those on the safest and most liquid assets skyrocketed; many borrowers were unable to borrow at any interest rate; and the world economy suffered a severe recession.

Financial crises are the subject of Sections 10.6 and 10.7. The first presents the classic Diamond-Dybvig model of the possibility of a self-fulfilling run on a financial institution that would otherwise be solvent. The second addresses the issue of how financial market disruptions and failures can spread among financial institutions.

The final issue concerns the social value of financial markets. In the presence of imperfections, private marginal products are potentially either bigger or smaller than social marginal products. For example, someone who channels funds from small savers to a poor entrepreneur with the potential to become the next Thomas Edison or Steve Jobs will likely capture only a tiny part of the social value of the resulting inventions that would not have occurred otherwise. On the other hand, someone who makes an enormous profit by buying an undervalued asset whose price would have risen a few seconds later in any event probably has almost no effect on any actual investment decisions, and so has negligible social product. Likewise, it is far from obvious that it is socially optimal that many of the most talented college graduates in the United States pursue careers in finance rather than, for example, working to create new products, better production processes, or new knowledge.

Although this last issue is fascinating and important, we will not pursue it. McKinnon (1973) and others argue that the financial system has important effects on overall investment and on the quality of the investment projects undertaken, and thus on economies' growth over extended periods. Because the development of the financial system may be a by-product, rather than a cause, of growth, this argument is difficult to test. Nonetheless, there is at least suggestive evidence that financial development is important to growth (for example, Levine and Zervos, 1998, Rajan and Zingales, 1998, and Jayaratne and Strahan, 1996).

10.1 A Model of Perfect Financial Markets

This section presents a model of perfectly functioning financial markets to show how the interplay of consumer preferences and the set of possible investments determine what investment projects are undertaken and how claims on the projects' output are valued.

Assumptions and Equilibrium Conditions

The economy lasts for two periods. A representative household has an endowment E of the economy's sole good in period 1 and no endowment in period 2. It maximizes the expected value of its lifetime utility, which is given by

$$V = U(C_1) + \beta U(C_2), \quad \beta > 0, \quad U'(\cdot) > 0, U''(\cdot) < 0, \tag{10.1}$$

where C_t is the household's consumption in period t.

All period-2 output comes from investments undertaken in period 1. There are N possible investment projects. The output of each project is potentially uncertain. Specifically, there are S possible states of the world in period 2. If quantity K_i of period-1 output is devoted to project i, it produces $R_{is}K_i$ in period 2 in state s. We let π_s denote the probability of state s occurring; the π_s's satisfy $\pi_s \geq 0$ for all s and $\sum_{s=1}^{S} \pi_s = 1$. The K_i's cannot be negative. It is convenient to think of each investment project as being undertaken by a distinct firm.

Finally, the economy is perfectly competitive: households and firms are price-takers.

It is straightforward to write down the conditions that characterize the equilibrium of this economy. Because there are complete markets and no imperfections, we can describe the equilibrium in terms of Arrow-Debreu commodities—that is, claims on period-2 output in the various states of the world. Specifically, let q_s be the price, in units of period-1 output, of a claim on one unit of perod-2 output in state *s*. Then equilibrium is a set of prices, $\{q_s\}$, investment decisions, $\{K_i\}$, and consumption decisions, C_1 and $\{C_2^s\}$, with three properties.

First, households must be maximizing their utility subject to their budget constraint. The budget constraint is

$$C_1 + \sum_{s=1}^{S} q_s C_2^s = E. \tag{10.2}$$

The Euler equation for the experiment of reducing C_1 by a small amount and using the savings to increase C_2^s is

$$U'(C_1) = \frac{1}{q_s}\pi_s \beta U'(C_2^s) \quad \text{for all } s. \tag{10.3}$$

We can rearrange this as

$$q_s = \pi_s \beta \frac{U'(C_2^s)}{U'(C_1)} \quad \text{for all } s. \tag{10.4}$$

That is, in equilibrium the price of a claim on output in state *s* equals the product of the probability of the state occurring and the marginal utility of consumption in that state relative to consumption today.

Second, there must be no unexploited profit opportunities. The cost of investing marginally more in project i, in terms of period-1 consumption, is just 1. The payoff is the revenues from selling the state-contingent output, which is $\sum_{s=1}^{S} q_s R_{is}$. If a strictly positive amount is invested in the project, the payoff to investing marginally more must equal the cost. And if nothing is invested in the project, the payoff from the first unit of investment must be less than or equal to the cost. Thus we have

$$\sum_{s=1}^{S} q_s R_{is} \quad \begin{cases} = 1 & \text{if } K_i > 0 \\ \\ \leq 1 & \text{if } K_i = 0 \quad \text{for all } i. \end{cases} \tag{10.5}$$

Notice that since there is a full set of Arrow-Debreu commodities, there is no risk in undertaking the project: although the amount that the project produces depends on the state, claims on output in all states are sold in period 1.

Finally, markets must clear. The market-clearing condition in period 1 is

$$C_1 + \sum_{i=1}^{N} K_i = E. \tag{10.6}$$

And the market-clearing condition for claims on period-2 output in state s is

$$\sum_{i=1}^{N} K_i R_{is} = C_2^s \quad \text{for all } s. \tag{10.7}$$

The number of equilibrium conditions is 1 (from [10.2]), plus S (from [10.3] or [10.4]), plus N (from [10.5]), plus 1 (from [10.6]), plus S (from [10.7]). The unknowns are the S q_s's, the S c_2's, the N K_i's, and C_2. The number of equations exceeds the number of unknowns by 1 because of Walras's law.

Discussion

From firms' perspective, this model is little different from the partial-equilibrium model of investment we studied in Chapter 9, particularly the model of investment under uncertainty in Section 9.7. And from households' perspective, the model is similar to the model of consumption in the presence of risky assets in Section 8.5. But because both investment and consumption are now endogenous, the marginal utility of consumption in different states, and hence the payoff to investment projects in different states, is now endogenous.

The only assets with net supplies that are strictly positive are claims on the output of the investment projects that are undertaken. Equilibrium requires that the price of a claim on the marginal output of a project equals the marginal cost of undertaking the project. In our simple set-up, where period-1 consumption can be converted one-for-one into any of the investment goods, this means that the price of a claim on the output of 1 unit of any investment projects that are undertaken is 1 (see [10.5]).

Although the net supplies of any other potential financial assets are zero, one can still think of markets where some agents can sell them and others buy them. The price of any asset (including ones without positive net supplies) is determined by the pricing kernel of this economy, which is determined by the marginal utility of consumption in different states. That is, the price of an asset with payoff in state s of x_s is $\sum_{s=1}^{S} q_s x_s$. Two potentially interesting assets are debt and insurance. The price of riskless debt—that is, an asset that pays 1 unit regardless of the state—is $\sum_{s=1}^{S} q_s$. The price of insurance against state s occurring—that is, an asset that pays

1 unit in state *s* and 0 otherwise—is q_s. Of course, since all households are same, we will not observe some agents selling these assets and others buying them. But with heterogeneity in preferences or income, we would.

In this model, financial markets pay a role similar to that of the Walrasian auctioneer in standard microeconomic models: they determine the prices of the economy's Arrow-Debreu commodities that lead to a balancing of supply and demand. All markets are perfectly competitive, information is symmetric, and there are no externalities. Thus the equilibrium is Pareto efficient. If we enriched the model to allow for the arrival of new information (for example, about the probabilities of the different states or the returns to investment projects in different states), there would be changes in asset prices, but they would be efficient. And the only reason for there to be sudden large changes in price would be the sudden arrival of major news.

There are numerous possible extensions that would not affect these central features of the model. Examples include additional periods, heterogeneity among households, adjustment costs in investment, and a role for other inputs into production (most obviously, labor supplied by households). None of these modifications would alter the messages of the simple model that financial markets are where households and firms meet to efficiently share risk and determine the level and composition of investment. Rather than pursuing those extensions, the rest of the chapter turns to settings where financial markets have more significant consequences.

10.2 Agency Costs and the Financial Accelerator

Introduction

In the models of Chapter 9 and Sections 10.1, all parties are equally well informed, and so financial markets function efficiently. Potential investments are valued according to their state-contingent payoffs using the prevailing economy-wide social discount factor. As a result, they are undertaken if their value exceeds the cost of acquiring and installing the necessary capital.

In fact, however, firms are much better informed than potential outside investors about their investment projects. Outside financing must ultimately come from individuals. These individuals usually have little contact with the firm and little expertise concerning the firm's activities. In addition, their stakes in the firm are usually low enough that their incentive to acquire relevant information is small.

Because of these problems, institutions such as banks, mutual funds, and bond-rating agencies that specialize in acquiring and transmitting information play central roles in financial markets. But even they are much less informed than the firms or individuals in whom they are investing their funds. The issuer of a credit card, for example, is usually much less informed than the holder of the card about the holder's financial circumstances and spending habits. In addition, the presence of intermediaries between the ultimate investors and firms means that there is a two-level problem of asymmetric information: there is asymmetric information not just between the intermediaries and the firms, but also between the individuals and the intermediaries (Diamond, 1984).

Asymmetric information creates *agency problems* between investors and firms. Some of the risk in the payoff to investment is usually borne by the investors rather than by the firm; this occurs, for example, in any situation where there is a possibility that the firm may go bankrupt. When this is the case, the firm can change its behavior to take advantage of its superior information. It can only borrow if it knows that its project is particularly risky, for example, or it can choose a high-risk strategy over a low-risk one even if this reduces expected returns. Thus asymmetric information can distort investment choices away from the most efficient projects. In addition, asymmetric information can lead the investors to expend resources monitoring the firm's activities, and the managers or entrepreneurs running the firm to devote less than the socially optimal amount of effort to the firm. Thus again, asymmetric information imposes costs.

This section presents a simple model of asymmetric information and the resulting agency problems, and discusses some of their effects. We will find that when there is asymmetric information, investment depends on more than just interest rates and profitability; such factors as investors' ability to monitor firms and firms' ability to finance their investment using internal funds also matter. We will also see that asymmetric information changes how interest rates and profitability affect investment and magnifies the effects of shocks to the economy.

Assumptions

An entrepreneur has the opportunity to undertake a project that requires 1 unit of resources. The entrepreneur has wealth of W, which is less than 1. Thus he or she must obtain $1 - W$ units of outside financing to undertake the project. If the project is undertaken, it has an expected output of γ, which is positive. γ is heterogeneous across entrepreneurs and is publicly observable. Actual output can differ from expected output, however; specifically, the actual output of a project with an expected output of γ is distributed uniformly on $[0,2\gamma]$. Since the entrepreneur's wealth is all invested in the project, his or her payment to the outside investors cannot exceed the project's output. This limit on the amount that the entrepreneur can pay to outside investors means that the investors must bear some of the project's risk.

To keep things simple, we assume that the entrepreneur and outside investors are risk-neutral, and that there is a technology with no risk or asymmetric information that yields a rate of return of r for sure. We also assume that the outside investors are competitive. These assumptions have several implications. First, the project is socially desirable if and only if the expected rate of return is greater than r; that is, the requirement for a social planner to want the project to be undertaken is $\gamma > 1 + r$. Second, because the entrepreneur can invest at the risk-free rate, he or she undertakes the project if the difference between γ and the expected payments to the outside investors is greater than $(1 + r)W$. And third, competition implies that in equilibrium, outside investors' expected rate of return on any financing they provide to the entrepreneur is γ.

The key assumption of the model is that entrepreneurs are better informed than outside investors about their projects' actual output. Specifically, an entrepreneur observes his or her output costlessly; an outside investor, however, must pay a cost c to observe output. c is assumed to be positive; for convenience, it is also assumed to be less than expected output, γ.

This type of asymmetric information is known as *costly state verification* (Townsend, 1979). It is surely not the most important type of asymmetric information in practice. There are two reasons that it is nonetheless useful to consider. First, it is clearly and strongly grounded in assumptions about the microeconomic environment. Second, as we will see, it shows not only how asymmetric information affects investment outcomes, but also how it shapes financial contracts. However, other types of information asymmetries, such as asymmetric information about the riskiness of projects or about entrepreneurs' actions, have broadly similar implications concerning distortions and the amplification of shocks.

[NOTE TO COPYEDITOR/COMPOSITOR: THE MATERIAL ON PP. 438–443 OF THE 4TH EDITION (SPECIFICALLY, FROM "The Equilibrium under Symmetric Information" ON P. 438 THROUGH WHAT IS CURRENTLY EQ. [9.42] ON P. 443) GOES HERE. THE ONLY CHANGES ARE RENUMBERINGS OF FIGURES (AND FIGURE CALLOUTS) AND OF EQUATIONS (AND REFERENCES TO EQUATIONS). FIGURE 9.15 BECOMES FIGURE 10.1; FIGURE 9.16 BECOMES FIGURE 10.2; FIGURE 9.17 BECOMES FIGURE 10.3. EQUATION (9.38) BECOMES EQUATION (10.8); …; UP THROUGH EQUATION (9.42), WHICH BECOMES EQUATION (10.12).]

In general, either constraint—investors' willingness to lend to the entrepreneur at some interest rate, or the entrepreneur's willingness to undertake the project if a loan is available—can be the relevant one. Suppose, for example, that W is very small. Then undertaking the project is attractive to the entrepreneur even when investors obtain the maximum possible revenues. Thus what determines whether the investment is undertaken is whether investors are willing to finance it. On the other hand, suppose W is substantial, but γ is only slightly above $1 + r$. In this case, there is a level of D that gives investors their needed net revenues, $(1 + r)(1 - W)$, but the agency costs involved may exceed $\gamma - (1 + r)$. In this case, whether the investment is undertaken is determined by whether the entrepreneur is willing to undertake it.

Discussion

Although we have derived these results from a particular model of asymmetric information, the basic ideas are general. Suppose, for example, that there is asymmetric information about how much risk the entrepreneur is taking. In such a situation, if the investor bears some of the cost of poor outcomes, the entrepreneur has an incentive to increase the riskiness of his or her activities beyond the point that maximizes the expected return to the project. Thus there is *moral hazard.*

As a result, asymmetric information again reduces the total expected returns to the entrepreneur and the investor, just as it does in our model of costly state verification. Under plausible assumptions, these agency costs are decreasing in the amount of financing that the entrepreneur can provide (W), increasing in the amount that the investor must be paid for a given amount of financing (r), decreasing in the expected payoff to the project (γ), and increasing in the magnitude of the asymmetric information (c when there is costly state verification, and the entrepreneur's ability to take high-risk actions when there is moral hazard).

Similarly, suppose that entrepreneurs are heterogeneous in terms of how risky their projects are, and that risk is not publicly observable—that is, suppose there is *adverse selection*. Then again there are agency costs of outside finance, and again those costs are determined by the same types of considerations as in our model. Thus the qualitative results of this model apply to many other models of asymmetric information in financial markets.

Likewise, although we have discussed the effects of asymmetric information in the context of an entrepreneur with a potential investment project, the ideas are more general. For example, they carry over to a household with uncertain future income that is considering buying a house or an automobile, and to an established firm with uncertain future profits that is considering increasing its capital stock.

The Financial Accelerator

One critical implication of this analysis is that financial-market imperfections tend to magnify the effects of shocks to the economy. Our model shows that one determinant of agency costs is the resources that agents have to finance investment themselves. When some force reduces output, those resources—the excess of current revenues over current expenses in the case of firms, current incomes in the case of households—fall. Thus the agency costs associated with a given level of investment rise. As a result, investment falls, magnifying the initial fall in output. Bernanke and Gertler (1989) show this formally in the context of technology shocks in a simple real-business-cycle model, but the logic is general. For example, it applies equally to monetary shocks in sticky-price models.

Financial-market imperfections can magnify the effects of shocks not only through the shocks' impact on current resources, but also through their impact on the value of collateral. If we extended the model of this section to include collateral that the outside investor could claim in the event of default, agency costs would be greater when the value of the entrepreneur's collateral was lower. And there are two reasons that asset values tend to fall when economy-wide output falls. First, since declines in output are generally at least somewhat persistent, a fall in output today is typically associated with decreases in expectations of future output. As a result, the expected future marginal products or service flows from assets fall, which translates into lower values of those assets today. Second, the financial-market imperfections themselves act to amplify the falls in collateral values from declines in output. For example, Kiyotaki and Moore (1997) develop a model where a productive asset that can be used as collateral is held by two types of firms: credit-constrained and unconstrained firms. When a fall in economy-wide output raises agency costs, and so reduces the constrained firms' ability to borrow and purchase the asset, more of the asset must be held by the unconstrained firms. With diminishing marginal product, the marginal product of the asset at these firms falls. In the baseline case in which there is no change in the interest rate at which these firms (which can borrow and lend at the prevailing safe interest rate) discount future profits, the value of the asset falls. The result is further reductions in the constrained firms' purchases of the asset, further falls in its value, and so on.

The idea that financial-market imperfections magnify the effects of shocks to economy-side output is known as the *financial accelerator*.[1]

Other Implications

The model of this section has many other implications. As with the financial accelerator, the most important implications arise from financial-market imperfections in general rather than from the specific model. Here we discuss five.

First, financial-market imperfections have a particularly large magnification effect for output movements stemming from changes in interest rates. Consider an increase in the safe interest rate, for example as a result of a tightening of monetary policy in an economy with sticky prices. Agency costs rise not just because of the fall in output, but also because interest rates affect agency costs directly. An increase in r raises the total amount the entrepreneur must pay the investor. This means that the probability that the investor is unable to make the required payment is higher, and thus that agency costs are higher. This acts to reduce investment, and so further magnifies the output effects of the increase in interest rates.

Second, in the case of monetary and other aggregate demand disturbances, there is yet another channel through which financial-market imperfections make the real effects of shocks larger: as described in Section 6.7, they increase the degree of real rigidity. Recall that firms' incentives to change prices are an important determinant of the real effects of nominal shocks. Greater real rigidity—that is, a smaller responsiveness of profit-maximizing real prices to

[1] Since the financial accelerator has nothing to do with an increasing rate of change, a more logical name would be the financial amplifier.

changes in aggregate output—reduces incentives for price adjustment in response to nominal shocks, and so increases the shocks' real effects. Costs of obtaining financing are one component of firms' costs. Thus, the fact that agency costs are higher when aggregate output is lower acts to mute falls in firms' overall costs when output is lower, and so reduces the fall in their profit-maximizing prices. If firms face costs of adjusting their nominal prices, the result is that they adjust their prices by less in response to a monetary shock, and so the effects of the shock on aggregate output are greater.

Third, and more prosaically, because the agency costs arising from asymmetric information raise the cost of external finance, they reduce the level of investment. Under symmetric information, investment occurs in our model if $\gamma > 1 + r$. But equation (10.12) shows that when there is asymmetric information, investment can occur only if $\gamma > 1 + r + A(c,r,W,\gamma)$. Thus the agency costs reduce investment at a given safe interest rate.

Fourth, many variables that do not affect investment when capital markets are perfect matter when they are imperfect. Entrepreneurs' wealth provides a simple example. Suppose that γ and W are heterogeneous across entrepreneurs. With perfect financial markets, whether a project is funded depends only on γ. Thus the projects that are undertaken are the most productive ones. This is shown in Panel (a) of Figure 10.4.[2] With asymmetric information, in contrast, since W affects the agency costs, whether a project is funded depends on both γ and W. Thus a project with a lower expected payoff than another can be funded if the entrepreneur with the less productive project is wealthier. This is shown in Panel (b) of the figure.

Two other examples of variables that affect investment only when capital markets are imperfect are average tax rates and idiosyncratic risk. If taxes are added to the model, the

[2] NOTE TO COPYEDITOR/COMPOSITOR: THIS IS THE CURRENT FIGURE 9.18, ON P. 446 OF THE FOURTH EDITION.

average rate (rather than just the marginal rate) affects investment through its impact on firms' ability to use internal finance. And risk, even if it is uncorrelated with consumption, affects investment through its impact on agency costs. Outside finance of a project whose payoff is certain, for example, involves no agency costs, since there is no possibility that the entrepreneur will be unable to repay the investor. But, as our model shows, outside finance of a risky project involves agency costs.

Finally, and critically, our analysis implies that the financial system itself can be important to investment. The model implies that increases in c, the cost of verification, reduce investment. More generally, the existence of agency costs suggests that the efficiency of the financial system in processing information and monitoring borrowers is a potentially important determinant of investment.

This observation has implications for both long-run growth and short-run fluctuations. With regard to long-run growth, McKinnon (1973) and others argue that the financial system has important effects on overall investment and on the quality of the investment projects undertaken, and thus on economies' growth over extended periods. Because the development of the financial system may be a by-product, rather than a cause, of growth, this argument is difficult to test. Nonetheless, there is at least suggestive evidence that financial development is important to growth (for example, Levine and Zervos, 1998, Rajan and Zingales, 1998, and Jayaratne and Strahan, 1996).

With regard to short-run fluctuations, our analysis implies that disruptions to the financial system can affect investment, and thus aggregate output. Recall that the transformation of saving into investment is often done via financial intermediaries, creating a two-level asymmetric information problem. This creates a potentially large propagation mechanism for shocks.

Suppose some development—for example, the crash of the stock market in 1929 and the contraction of the economy in 1930, or the fall in house prices in 2007 and 2008—lowers borrowers' wealth. This not only reduces their ability to borrow and invest; it also weakens the position of financial intermediaries, and so reduces their ability to obtain funds from ultimate wealthholders. This reduces their lending, further depressing investment and output. This amplification can be compounded by links among intermediaries. In the extreme, some intermediaries fail. If they have specialized knowledge about particular borrowers, those borrowers' investment collapses. The end result can be catastrophic. Precisely these type of financial amplification mechanisms were at work in the Great Depression (Bernanke, 1983b), and they were central to the crisis that began in 2007.