

Why Does the Economy Fall to Pieces after a Financial Crisis?

Robert E. Hall

The worst financial crisis in the history of the United States and many other countries started in 1929. The Great Depression followed. The second-worst struck in the fall of 2008 and the Great Recession followed. Commentators have dwelt endlessly on the causes of these and other deep financial collapses. Less conspicuous has been the macroeconomists' concern about why output and employment collapse after a financial crisis and remain at low levels for several or many years after the crisis. This article pursues modern answers to that question. It focuses on events in the United States since 2008.

Existing macroeconomic models account successfully for the immediate effects of a financial crisis on output and employment. I will lay out a simple macro model that captures the most important features of modern models and show that realistic increases in financial frictions that occurred in the crisis of late 2008 will generate declines in real GDP and employment of the magnitude that occurred. But this model cannot explain why GDP and employment failed to recover once the financial crisis subsided—the model implies a recovery as soon as financial frictions return to normal. At the end of the article I will mention the ideas that are in play to explain the persistent adverse effects of temporary crises, but these ideas have not made their way into the mainstream model.

This article cites only a few of the many important contributions to the mainstream model. My paper Hall (2009) contains many citations and the forthcoming new volume of the *Handbook of Monetary Economics* discusses the literature fully.

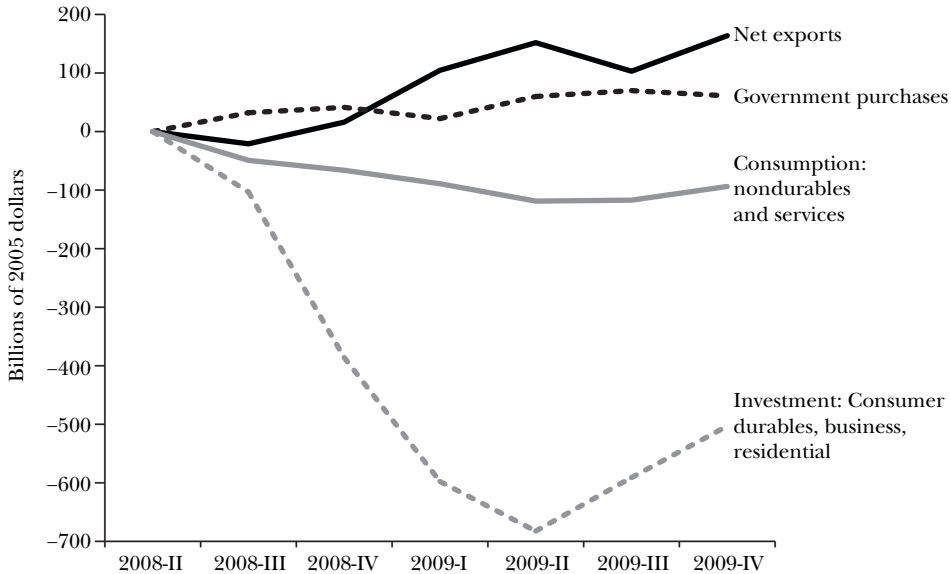
Figure 1 shows what happened to four components of real GDP after the second quarter of 2008. Three components were only slightly affected. Net exports and

■ Robert E. Hall is the Robert and Carole McNeil Joint Hoover Senior Fellow and Professor of Economics, Stanford University, Stanford, California. His e-mail address is rehall@stanford.edu.

doi=10.1257/jep.24.4.3

Figure 1

Changes from the Second Quarter of 2008 in Four Components of Real GDP during the Crisis




Source: U.S. National Income and Product Accounts, Table 1.1.6.

government purchases rose—the latter reflects in part the fiscal stimulus authorized in February 2009. Consumption of nondurables and services sagged a little immediately, before beginning a recovery in the later half of 2009. Essentially the entire large decline in real GDP was in investment, broadly conceived. This measure includes investment-type purchases by consumers (cars, appliances, furniture), business investment in plant, equipment, and inventories, and residential investment. The crisis did not cause a general contraction in spending. Rather, the contraction is essentially entirely in investment.

All components of investment rely on financial markets for funds. Residential investment relies on it the most—homebuilders finance construction with bank loans and homebuyers almost always finance a substantial fraction of the price of a newly built home. The majority of new car buyers take out loans and the car-making industry depends heavily on borrowing in the bond market. Across all industries, borrowing to finance plant, equipment, and inventories is common. It is not surprising that when investment declines after a financial crisis these flows dry up. What is surprising is that almost all of the huge decline in output in the U.S. economy following the crisis was confined to investment.

Standard principles of macroeconomics hold that interest rates are the regulator of investment and saving. When demand is strong, interest rates are high, so investment projects with lower returns fail to make the bar; the claimants on output

Figure 2
Interest Rates



Source: St. Louis Federal Reserve Bank FRED database (Federal Reserve Economic Data).

with the highest value receive goods and services while those with lower values defer their demands until interest rates return to normal. In a slump, interest rates fall, stimulating demand by lowering the bar for investment projects. Not only does this process operate on its own, but the Federal Reserve encourages equilibration of sources and uses of output by raising its short-term rate in booms and lowering it in slumps. Prior to 2008, many economists had begun to think that this process worked pretty well. Though recessions occurred in 1990 and 2001, they were widely separated and saw only mild contractions in output.

Figure 2 shows what happened to the interest rates that private decisionmakers faced around the time of the crisis. For comparison, it also shows the rate that the U.S. Treasury paid for its 10-year note. The Treasury rate fell dramatically in late 2008 as the crisis took hold and the bottom fell out of the U.S. economy. Had private borrowers enjoyed the same decline in borrowing cost, some of the adverse effect of the financial pounding would have been offset. Only in mortgages did the interest rate fall along with the Treasury rate. Here the Federal Reserve helped in an unusual way—it bought hundreds of billions of dollars worth of mortgages packaged as securitized mortgage-backed bonds. The decline in observed mortgage interest rate also conceals a large increase in down-payment and credit-rating requirements for borrowers. Would-be homebuyers with little cash and lower credit ratings, who

could have borrowed in 2005, were rationed out of the mortgage market in 2008, especially when the crisis struck in the early fall.

The crisis caused an *elevation* of the Baa corporate bond rate, a borrowing cost at the heart of the business borrowing market. Far from serving as an equilibrating factor, the Baa rate worsened the collapse of investment. Plant and equipment spending and inventory investment were further discouraged even as idle resources grew. That interest rate failed to perform its recessionary duty.

Credit-card borrowing rates—influential for many consumer purchases of appliances and furniture—stayed constant during the height of the crisis, then rose as the crisis itself receded in early 2009 but the recession worsened. Car-loan rates also failed to help stimulate car purchases at a time when the human and capital resources for making cars were mostly idle. In January 2009, General Motors did not build a single passenger car, an absolutely unprecedented collapse.

Financial Friction

I use the term “friction” to mean a cost to one side of a transaction that is not a benefit to the other side. A friction usually arises because some intermediary is taking a cut from the transaction. A tax creates a friction where the intermediary is the government. In debt markets, where the intermediary is a bank or other financial institution, a friction drives up the borrower’s cost of funds without raising the payoff that the supplier of funds receives. The dominant view among macroeconomists today is that a financial crisis causes real economic activity to collapse by raising frictions. The canon of this line of thinking is a *Handbook of Macroeconomics* chapter by Ben Bernanke (*the Ben Bernanke*), Mark Gertler, and Simon Gilchrist (1999). Bernanke had a long history of scholarship on the Great Depression which pursued that exact topic. His 1999 chapter with Gertler and Gilchrist fully refutes the widespread notion that macroeconomics lacked the tools for connecting financial turmoil to economic collapse.

The friction in Bernanke, Gertler, and Gilchrist (1999) arises from an agency cost in the contracting relationship between a lender and a borrower. The lender hopes for a trouble-free debt relationship in which the producer will prosper and repay the face amount of the debt. The huge efficiency of debt is that if the user of capital repays debt in full, the lender has no reason to spend any resources figuring out if the borrower has completed all duties under the contract. To put this point another way, an equity investor has to monitor an investment in all states of the world, while a debt lender only has to monitor if the borrower fails to repay in full. This basic setup, created by Townsend (1979), is called the “costly state verification” model. The pricing of a loan in this setup depends on the likelihood that the borrower will not repay in full and that the subsequent investigation costs, usually involving bankruptcy, will be a further deduction from the lender’s profit. The expected verification cost is the financial friction, because it has to be charged to the borrower but is not a benefit to the lender—it is an actual resource cost.

The likelihood of borrower default depends on the borrower's own wealth and on the distribution of random outcomes that might result in such a low profit that repayment of the loan in full becomes impossible. Events that cause depletion of borrowers' wealth or that raise the likelihood of low profit result in a widening of the agency friction. The application to the 2008 crisis is obvious—all institutions and businesses that held real estate or financial claims on real estate lost wealth when real-estate prices declined.

Agency frictions can occur at every link in the chain from households that save to businesses that produce output. Savers have to worry about default by financial institutions who take their wealth as deposits or investments. Those institutions have to worry about default by the producing businesses who borrow from the institutions. Modern economies have many such links separating saving from producing. Most of the time, the agency frictions are tiny because defaults are truly rare. But when an adverse shock depletes wealth in many of the links, the frictions can suddenly become large.

The Sensitivity of Real Activity to Financial Frictions

How much difference does it make to macroeconomic equilibrium if financial frictions widen when, say, real-estate values decline? The answer requires delving into some of the big issues in modern macroeconomics—it depends on answers to questions about the sources of booms and recessions. With regard to the aggregate influence of other driving forces of fluctuations—monetary shocks, movements of government purchases, and changes in taxes—the research generally shows that in standard neoclassical models, with normal preferences and technology and competitive markets, shifts of realistic magnitude fail to deliver anything like the volatility seen in the U.S. economy.

A recent paper of mine, Hall (2009), explains this point in detail with respect to changes in government purchases. In that case, the multiplier—the increase in GDP caused by each dollar of government purchases—is the natural metric. Neoclassical models have multipliers around 0.4. One of the reasons for the modest effect is that in these models consumption falls by 0.6 when government purchases rise by 1. But U.S. experience with large changes in government purchases, mainly in wartime, contradicts these predictions. GDP rises by a large fraction of the increase in purchases and consumption does not fall.

In the effort to bring the macro model into alignment with actual experience, two possible departures from the neoclassical benchmark have attracted the attention of researchers: 1) introducing variations in market power in product and labor markets and 2) taking account of unemployment. Hall (2010) contains a formal version of what follows.

Countercyclical Market Power

Market power is another friction. The standard measure of market power in a product market is the ratio of selling price to marginal cost; the ratio is 1 under perfect

competition and exceeds 1 in the presence of market power. When a seller charges more than marginal cost for a product, it implies that the buyer is giving up more than the suppliers of resources receive. The producer or retailer with market power is the agent who is driving a wedge into the transaction, in the same way that the government drives a tax wedge. Frictions in the labor market take the form of a wedge between the marginal product of labor and the marginal value of workers' time.

Another key chapter in that earlier *Handbook of Macroeconomics*, Rotemberg and Woodford (1999), explains the importance of variations in market-power friction and tries to document the friction empirically. A spontaneous reduction in market power is a stimulus to output and employment, just as a reduction in a tax rate encourages more production and work effort. Although exogenous movements in market power could be a powerful driving force on their own, the nearly exclusive focus of thinking in this topic has been how endogenous reductions in market power can amplify other expansionary forces. For example, if one adds to the neoclassical model an effect running from higher government purchases to reductions in market power, the multiplier rises substantially.

Why might market power decline when an exogenous force expands the economy? An early answer was that cartels are more likely to fall apart when demand is temporarily strong than at other times, because the payoff to defection is greatest when the defector can grab most of the larger market by taking rivals by surprise with a price cut. However, that idea has not become a significant part of the rationalization for countercyclical margins.

The dominant explanation in current thinking rests on price and wage stickiness. The idea is that a supplier sticks to a price or wage despite forces that would cause the supplier to seek a higher price or wage. A product supplier sticking to an earlier price in the face of an expansionary impulse experiences an increase in marginal cost by moving up its short-run marginal cost schedule and by facing increases in the prices of inputs, so the ratio of price to marginal cost falls. A worker sticking to an earlier wage who is now working more hours will have a higher marginal value of time, so the labor wedge falls as well. The original expansionary impulse is amplified as a drop in friction leads to further expansion.

Price and wage stickiness is hard to introduce into a model—see the canon of the New Keynesian model, Christiano, Eichenbaum, and Evans (2005), for the gory details. These details are arguably needed for studies of inflation and central bank policy, which are a major application for that class of models. But from the perspective of the analysis of real shocks, such as financial crises or changes in government purchases, a shortcut works fine. This shortcut introduces a constant-elastic relation between output and the price/cost markup ratio, with an elasticity of around -1 . Thus, a force that raises output from its normal level by 1 percent lowers the price/cost ratio by 1 percent.

Unemployment

Over the past 30 years, macroeconomics has gained a coherent, fully articulated formal theory of unemployment that comes to grips with the facts about joblessness and flows in the labor market. The theory retains a key idea from

earlier, less-formal macro thinking in that unemployment rises when product demand falls, but the theory contains a much fuller description of the rationing of scarce jobs during slumps.

Aggregate models are just beginning to incorporate unemployment explicitly. Most models, notably those of Christiano, Eichenbaum, and Evans (2005) and their co-authors and followers, use the sticky-wage specification to make the supply of total hours of work effectively fairly elastic with respect to the wage. Without adopting that specification, one can come reasonably close to mimicking the behavior of a full, modern, labor market sub-model by setting the wage elasticity of labor supply to a value about twice what micro labor-supply studies report. I elaborate this point in Hall (2009) with some citations that include my own work in this area.

Elastic labor supply interacts with countercyclical markups so that they reinforce each other. A model with both features has a government purchases multiplier of 0.98. Without the elastic labor supply, the multiplier drops to 0.40, and without the countercyclical markup, it drops to 0.60. Both features are needed to get the kind of amplification that seems to exist in the U.S. economy.

The Static Model

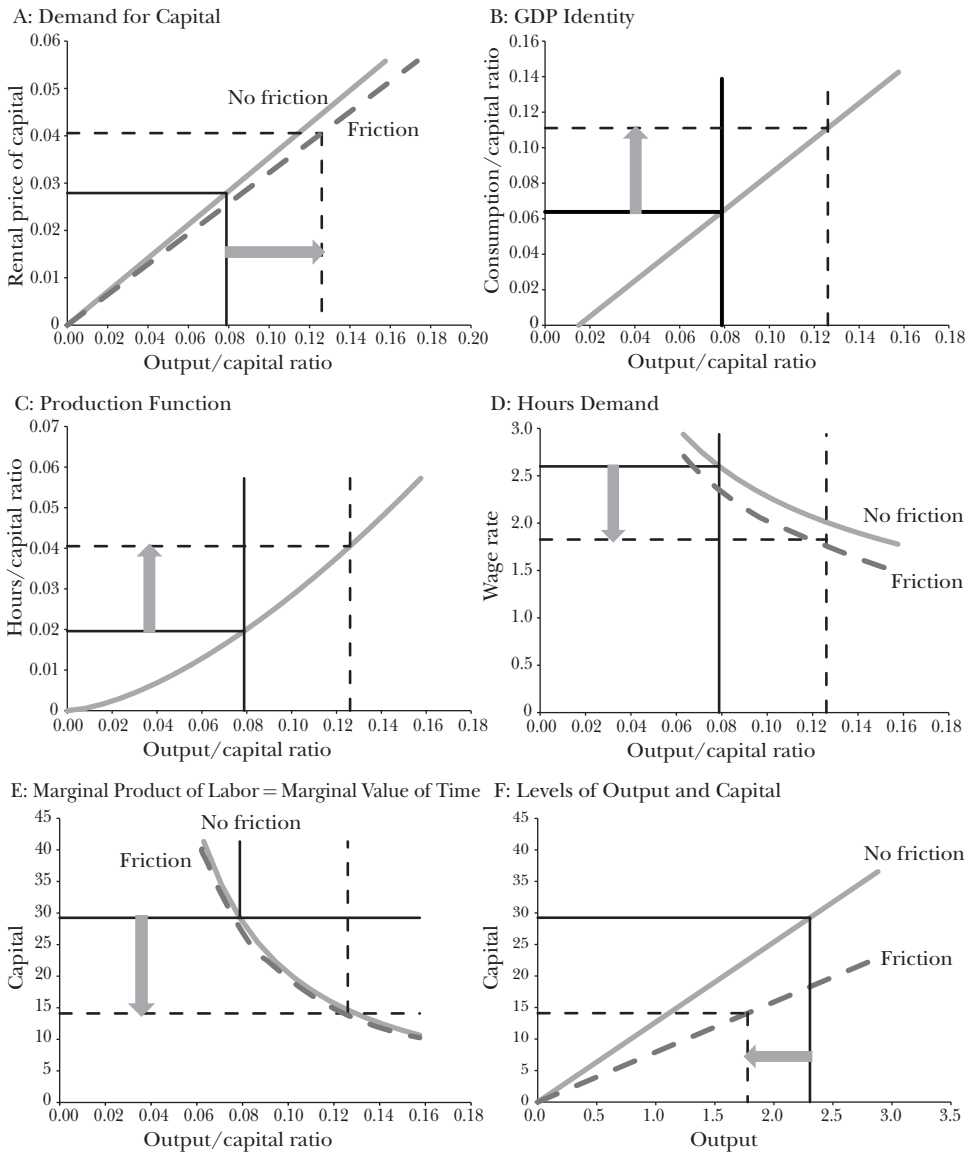
With these pieces in place, it's now possible to build a macroeconomic model that describes how an increase in financial frictions can serve as a shock that depresses the real economy. At present, such models are highly stylized. They are useful for illuminating the connections of economic theory, and by inserting plausible parameters, one can determine whether it is at least possible that the appearance of a financial friction of a plausible size can lead to a substantial drop in output.

There are two versions of this model. Through the panels of Figure 3, this section presents a basic static model; a later section discusses a dynamic model. It turns out that the behavior of the static model is surprisingly informative about the dynamic response and involves all the economic considerations that govern the dynamic response. The quantitative results described here are the result of a simulated version of the model, based on more-or-less realistic parameters.

The basic static model is an economy with a single production sector that has two inputs, labor and capital. It uses a Cobb–Douglas production function, with a labor elasticity of 0.65. There is a single household, which has a utility function that includes consumption and leisure. The utility function is consumption raised to the power -1 (corresponding to an intertemporal elasticity of substitution of 0.5) less hours of work raised to the power 1.5 (corresponding to a Frisch elasticity of labor supply of 1.9). (I boosted the elasticity of labor supply as a rough-and-ready way to take account of unemployment, as explained in Hall (2009).) The rate of depreciation of capital is 6 percent per year. Capital in this model has both a direct cost of production and also a cost of deferring consumption. The rate of time preference is 5 percent per year, which is also the real interest rate absent a financial friction.

Figure 3

Effects of a Financial Friction in the Basic Model



The model has no economic growth, no uncertainty, and it is a closed economy. The friction in this model takes the form of a property tax on capital, which adds 4 percent per year to the rental price of capital. (Proceeds of the tax are returned to the public as a lump sum.) The crisis also raises the price/cost markup ratio from 1.0 to 1.1. The simulation of this model shows that these modest alterations—the added financial friction and the rise in the price/cost markup ratio—have large

effects, cutting the capital stock in half and depressing output by 23 percent and consumption by 16 percent.

In the six panels, the effect of an increase in financial friction appears as a shift of a structural relationship, from a solid to a dashed line. An arrow shows the corresponding movement of one of the variables of the model. The static model is not simultaneous. Rather, one can follow the effects from variable to variable without circling back. The dynamic model, not surprisingly, is deeply simultaneous.

The upper left panel of Figure 3 shows the starting point. The original no-friction line shows a relationship, based upon the demand for capital, in which an increase in the rental price of capital causes a lower quantity of capital to be used in production, and thus leads to a rise in the output/capital ratio. When a financial friction is added, then the effect of a rise in the price of capital leads to a slightly larger rise in the output/capital ratio (slightly less capital is used to produce any given output).

The higher output/capital ratio leads to a higher consumption/capital ratio shown in the upper right panel; in fact, they rise by the same amount. This model has no growth, and so stationary investment is always the depreciation rate times the capital stock. Remember, GDP in this economy is just consumption plus investment. Because the investment/capital ratio is constant, a rise in the output/capital ratio necessarily implies a rise in the consumption/capital ratio through the GDP identity. In the simulation, the consumption/capital ratio rises from 0.06 to 0.11.

The higher output/capital ratio also leads to a higher hours-worked/capital ratio. As the economy moves along its constant-returns production function, the greater output needs to be produced with hours worked, as shown in the middle left panel.

The middle right panel shows a relationship based on the labor demand function. A rise in the output/capital ratio, as shown on the horizontal axis, means less capital is being used in the economy relative to output. The economy moves down and to the right along the labor demand curve and the wage falls. The rise in the markup ratio also enters here. It causes a shift inward in the labor demand curve: that is, with a higher price/cost markup, at a fixed wage, the output/capital ratio will be lower. As the panel illustrates, the wage rate falls for two reasons: both because of the increase in the output/capital ratio and also because of the increase in the markup.

The bottom left panel shows the central role of the household in the response. The earlier panels all dealt with intensities—ratios to the capital stock. The household determines the size of the economy because it controls the only primary factor in this economy: hours of work. The two curves—almost on top of one another—show the capital stock on the vertical axis such that the household maximizes its utility by setting equal the marginal product of labor and the marginal value of time (for the fastidious, by the marginal value of time, I mean the marginal rate of substitution between goods consumption and hours of work). The household can be in equilibrium at a variety of places: for example, it can have a low output/capital ratio and a high level of capital, or it can have a high output/capital ratio and a low level of capital. The upper left panel showed that the financial friction raises the output/capital ratio. Hence, the friction lowers the capital stock, which contracts the economy.

The bottom right panel translates the decline in the capital stock into a decline in the level of output, by multiplying the capital decline by the output/capital ratio. The proportional decline in output is less than in capital, because hours of work decline less than does the capital stock.

To summarize, the overall effect of the increased financial friction is to shift the economy into a lower-capital mode. The friction acts as a tax on the use of capital and the economy responds by using less capital and producing less output. In the dynamic model, the transition to a lower capital intensity results in a period of low investment.

Evidence about the Magnitude of the Increase in Frictions from Credit Spreads

Immediately after the critical events of mid-September 2008, the difference between private borrowing rates and the rates the federal government paid rose dramatically. These differences are called credit spreads. The spread has three components: (1) the expected default rate, (2) the difference in financial risk (that is, covariance with the common element of asset returns), and (3) the financial friction. Presumably all three rose during the crisis of late 2008, so the increase in the spread can be viewed as an upper bound on the increase in the financial friction.

Rationing of lending for new borrowing appears to have been important in financial markets in late 2008 and the following years, so the observed spread for new borrowing understates the effective spread facing a given borrower. Thus, it would seem most appropriate to measure spreads from the valuation of existing debt. One of the thickest markets for traded debt of private entities is the corporate bond market, so I use data on the prices of corporate bonds to measure the change in spreads in late 2008. Figure 4 shows the spreads for Baa corporate bonds over 20-year Treasury bonds. The spread began to rise in the summer of 2007 when the first hints of the crisis appeared. It leapt in the fateful month of September 2008, and then gradually fell back toward normal by the end of 2009. For further discussion of spreads in this journal, see Krishnamurthy (2010).


The relation between the spread and the financial wedge is controversial. The spread certainly includes the expected loss from default, which is the amount that bondholders were owed less the amount they recovered after a default. Although expected defaults presumably rose during the crisis, it appears that the rise was a tiny fraction of the total increase in the spread. Financial economists have found risk factors that rose in the crisis and help explain the general pattern of changes in asset prices that occurred in 2008 and other times. The challenge from the perspective of the issues in this paper is to relate those risk factors to the behavior of savers.

I believe that the most successful explanation of the jump in credit spreads in 2008 is segmentation of financial markets resulting from declining real-estate values. In normal times, each class of securities attracts specialists who have deep knowledge of the securities themselves and the markets in which they trade. The

Figure 4

The Difference between the Baa Corporate Bond Rate and the 20-Year Treasury Bond Rate

(in percentage points)



Source: Federal Reserve Bank of St. Louis FRED database (Federal Reserve Economic Data).

specialists deploy wealth put in their hands by financial institutions. If the wealth suddenly disappears, as it did when real-estate-based assets collapsed, the specialists are unable to function. In particular, they lack the means to buy underpriced securities. In the fullness of time, the specialists hook up with other sources of wealth and those affiliated with wealth gain the expertise—at which point the asset-pricing anomalies disappear.

Two episodes in the crisis provide good examples. In August 2007, when hedge funds and other wealth-holders felt the first tremors of real-estate declines, a number found it necessary to wind down similar positions in the stock market to meet cash requirements related to holdings of declining real-estate-based assets. The simultaneous sales of volumes of shares well above normal transaction volume depressed the prices of a vector of stocks without any noticeable effect on the stock market as a whole. If this pattern of sales had occurred regularly in the past, specialists ready to trade against it would have stabilized prices, but there were no such specialists for this isolated event. Similarly, in November 2008 at the height of the crisis, some hedge funds specializing in inflation-protected Treasury bonds had to sell out their positions, causing a temporary depression in the values of these bonds. Not enough

specialists remained with adequate wealth to take advantage of the pricing anomaly. It disappeared in a month or so as the wealth hired the available specialists or the specialists lined up wealth to bid the prices back to normal. Thus, while Wall Street offers no chronic opportunities for excess returns after adjustment for risk, it does offer transitory periods of such returns in the immediate wake of shocks that deprive specialists of the wealth needed to profit from the excess returns.

The transitory undervaluation of Baa bonds is an important illustration of a financial friction or wedge of the type I consider in this article. The important wedge is between the return that savers earn and the cost of funds to the users of savings. For Baa bonds, the users of savings are the companies issuing the bonds to finance plant and equipment or using the current market value of bonds to make decisions about such investments. Although nothing in principle should have stopped any saver from loading up on underpriced Baa bonds, that did not happen.

Another story that some economists tell about rising credit spreads in financial crises involves a “flight to quality.” Investors suddenly prefer high-quality debt, notably the debt of the U.S. government, to lower-quality claims such as Baa corporate bonds. Quality is not a feature recognized in normal asset-pricing models—these models look to perceptions of payoffs in different future states of the world and to the market prices of those payoffs. If a flight to quality is the result of a misunderstanding by savers of the implications of the crisis, then the rise in the credit spread functions as a friction or wedge. If the flight is based on correct perceptions that mid-rated corporations are in serious trouble but the U.S. government is not, then the rise in credit spreads cannot be considered a friction.

Some theorists have brought the concepts of quality and liquidity into asset pricing through adverse selection or lemons models. The spreads that these models predict generally do have the character of frictions or wedges. DeMarzo and Duffie (1999) define and discuss liquidity.

Dynamic Effects of the Financial Friction

The dynamic version of the model presented here continues to have a single representative household that maximizes utility in choosing between consumption and leisure. However, the dynamic model extends the basic static model in five ways:

First, it endogenizes the increase in the price/cost markup ratio by making the markup a constant-elastic function of output with an elasticity just under -1 . In this simulation, I picked the elasticity to generate a response to government purchases that I believed to be consistent with the empirical findings, with a purely judgmental boost to account for the likely downward bias in those findings.

Second, the dynamic model incorporates a feature of preferences that a number of authors have recognized as important during the past few years: the complementarity of goods consumption and hours of work. The rationalization is straightforward—when people are not working in the market, they have more time to work at home, and home production is a substitute for market purchases. Most

research on this topic has tackled the reduction in consumption that occurs discontinuously when people retire, but the principle applies equally to unemployment. In the U.S. data, the response of consumption to a change in government purchases is close to zero. Incorporating the degree of complementarity between goods and hours of work implied by retirement studies explains this finding nicely. In this way, modern macro is poised—in a way fully consistent with the evidence—between the old Keynesian view based on a simple consumption function that higher government purchases raise consumption because they raise income, and the neoclassical prediction that consumption falls when the government takes more output away from private uses.

Third, the dynamic model has capital adjustment costs, which are in line with empirical evidence.

Fourth, the dynamic model treats homebuilding and consumer purchases of durable goods as investment subject to the same friction as the basic model attributed only to capital. Because the services of houses and consumer durables have a consumption character while those of plant and equipment are an input to production, the two components of investment behave somewhat differently.

Fifth, the dynamic model characterizes the timing of the financial shock and resulting widening of frictions as a large initial widening, followed by a fairly rapid return to normal. The model has two interest rates. The higher one is the rate that borrowers pay to finance plant, equipment, consumer durables, and homes. The lower one is the rate that savers receive. It is also the rate that the government pays. The difference between the two rates is the amount of the wedge from the financial friction. Each quarter, the financial wedge from the crisis falls to 60 percent of its level in the prior quarter.

Like the earlier static model, the dynamic model describes a closed economy, so there is no response of net exports. As noted at the beginning of the article, the actual response of net exports in the recent recession was mildly positive, but it was small in relation to the huge decline in all types of investment. One could also think of the kind of model presented here as describing the world economy, which is an appropriate goal given the rather similar responses of all of the large economies of the world to the global crisis.

Detailed findings for this dynamic model are presented in Hall (2010). The model experiences a transitory rise in the financial friction separating savers from investors in the two categories of capital. Specifically, spread of the short-term borrowing rate for investors over the government's rate jumps 6 percentage points at the outset, then declines with a persistence of 0.6—so in the quarter following the shock, the spread is down to 3.6 percentage points. These figures are derived from an analysis of the short-term spread implicit in the longer-term spread shown in Figure 4. The initial jump of 6 percentage points in the spread is smaller than the one inferred from Baa bonds and larger than the one from AAA bonds.

The model replicates the small negative response of nondurables consumption that actually occurred. Consumers know that the adverse effect of the financial friction is transitory, so they ride through the temporary cutback in income during

Table 1
Comparison of Initial Quarter Response in Model to Actual U.S. Changes from 2008-Q2 to 2009-Q1

	<i>Percent decline</i>	
	<i>US, 2008-Q2 to 2009-Q1</i>	<i>Model, initial quarter</i>
Output	3.9	4.2
Consumption of nondurables and services	1.1	0.8
Investment in plant and equipment	17.7	12.7
Investment in housing and durables	12.5	14.8

Source: National Income and Product Accounts and author's calculations.

the recession by cutting back on cash outlays on durables and new houses, which is enough to finance continuing consumption of nondurables. Table 1 compares the model's immediate response to the shock to what actually happened to the U.S. economy over the period from the second quarter of 2008 to the first quarter of 2009. I chose a longer period for the actual economy than the period in the model because financial stress did not erupt at a discrete moment, as in the model, but built over the period from Bear Stearns in April 2008 to the end of the year. In the model, output falls by a little over 4 percent because of the shock, about the same as the actual U.S. decline. The model roughly tracks the large declines in the two investment categories. The unsurprising principle that a financial friction has its biggest effects on spending categories normally financed by borrowing holds both in reality and in the model.

All the effects decline as time passes after the shock at about the same rate. Macro models of this type lack internal dynamics. As the financial friction disappears, the model economy returns to normal. In this respect, the dynamic model (along with most other models) misses an important feature of the actual U.S. economy—the way in which a slump can persist for years after the visible negative force has disappeared. By mid-2010 in the U.S. economy, credit spreads are back to normal, but unemployment remains high and output is still far below its potential level.

Another Financial Impediment to Full Employment: The Zero Bound

This article has discussed how a financial crisis affects output and employment by acting like a tax on the use of capital. The resulting decline in investment, by businesses in the form of plant and equipment and by households in the form of home-building and acquiring durables, results in recession. Severe financial crises, such as in 1929 and 2008, depress the economy sufficiently to hit another financial snag—the inability of interest rates to drop below zero. The Federal Reserve responded promptly to the crisis by dropping its policy interest rate to close to zero

in late 2008, but it could not continue that particular program of stimulus further, because it lacks the ability to push rates below zero.

The reason for this zero lower bound on the interest rate is simple. If rates offered to investors are negative, they will hold their wealth in currency, which becomes attractive because its nominal interest rate is zero. The Fed stands ready to issue currency to anybody who wants to hold it. Currency holdings would take over the financial system if the Fed succeeded in generating negative rates.

Absent the zero lower bound, the Fed could depress short-term rates on Treasury bills and other safe securities into negative territory, thus lowering the borrowing rates facing private decisionmakers despite large spreads caused by financial frictions. The Fed could engineer a stimulus that would bring the economy back to full employment despite widened spreads from a financial crisis. The zero lower bound worsens the adverse effects of a financial crisis by preventing monetary policy from delivering more than a rather feeble offsetting stimulus.

Observations on the Science of Macroeconomics

Research led by Ben Bernanke about the long period of stagnation following the financial crises that began in 1929 resulted in a reasonably complete understanding of the after-effects of the crisis of 2008. The Bernanke, Gertler, and Gilchrist (1999) study describes the mechanism of the episode in complete, though broad-brush and stylized, terms. A large decrease in the values of the asset holdings of financial institutions resulted in a dramatic intensification of agency problems in those institutions and to a lesser extent in nonfinancial companies. Credit spreads widened and credit rationing became widespread. The diminished ability to finance the acquisition of capital goods resulted in huge cutbacks of all types of investment—plant, equipment, inventories, residential construction, and consumer durables purchases—while other categories of GDP remained roughly constant. All of these events fit existing macro models quite well. In fact, the surprise, if there is one, is just how well certain counterintuitive properties of existing models held up in the data. In particular, stability of consumption of services and nondurables in the face of transitory but extreme adverse shocks flies in the face of what seems like common sense, but it happened and it is a prediction of many models. The macroeconomists who moved away from the life-cycle view of consumption in favor of a belief that a large fraction of families just spend their current incomes have had to ditch that change and go back to the life-cycle, consumption-smoothing model to track the small response of nondurables and services consumption to a large decline in GDP.

The failing of macroeconomics in the last few years was not a lack of understanding of what happens in the overall economy when a financial crisis strikes. Conditional on a large decline in asset values among financial institutions, we got things right. Instead, the failings lay elsewhere.

First, only a few economists, certainly not including this writer, understood that anything that went as high as real-estate valuations in 2005 would have to come down

at some point (as opposed to house values just stagnating for a time, for instance). We did not consider that, if such a decline occurred, financial chaos would ensue.

Second, many macroeconomists failed to foresee a financial crisis as the situation evolved after 2005 because financial markets had handled the huge decline in asset values in 2000–2001 so smoothly. We overlooked the key point that the assets in decline in that episode were business assets, mainly in a sector, high-tech, that uses little debt finance and thus has little leverage. Not only are the companies holding the declining assets not at risk for growing agency frictions, but the entities holding claims on those companies hold equity, not debt, and the entities rarely lever their positions. Business equity resides in large portfolios of rich families, in mutual funds, and in endowments, and these entities rarely borrow against their holdings. Consequently, the large part of the economy built on business assets and equity markets is essentially bullet-proof, as the 2000–2001 crash demonstrated. There wasn't the slightest hint of a financial crisis in the 2000–2001 episode.

By contrast, the other important asset class in the U.S. economy, real estate, has high leverage everywhere. Most homeowners borrow as much as they can when buying a house; they become unlevered only if they remain in the house and pay down the mortgage. In the 2000s, borrowing as much as you could meant borrowing close to the entire price of the house. The story of the multiple added levels of leverage among financial institutions holding real-estate-related assets has now been told many times. Thus, the economy is severely at risk from even a small decline in real-estate values—following literally the Bernanke, Gertler, and Gilchrist (1999) script—while it is little affected by even large declines in the values of non-real-estate business assets.

Big leaps forward in research are inherently unforecastable, so I won't try. Clearly, our existing models of financial friction are highly stylized. I look forward to improvements in our understanding of many issues, both on the sources and nature of financial frictions and in the big issues of macroeconomic amplification.

On the sources of financial frictions, we understand that principal–agent problems arise in many places in financial systems, so that financial managers are tempted to misbehave in unspecified ways, but there is a huge gap between the models and the experience. We know that adding three or four percentage points to frictions knocks the economy flat, but we are not sure that existing models of these agency problems can deliver frictions of that magnitude.

The two major elements of current thinking about macro amplification will surely see important progress in coming years. With respect to the counter-cyclical price/cost markup, the sticky price model continues to swirl in controversy. Merchants change their prices all the time, but most price changes are temporary cuts and are followed by restoration of an earlier price. Investigation of actual pricing practices has turned up all kinds of fascinating results sufficient to support the conclusion that actual pricing is way more interesting and complicated than it is in textbooks, but without resolving the issue about how much markups decline as output expands. Even straightforward empirical research to measure markups and

relate them to movements in output has been inconclusive, mainly because it is so hard to measure labor cost.

In the labor market, the big issue is the determination of unemployment. We have a really good understanding of the flows in and out of employment and the process in which jobs are rationed by quantity, but we are still struggling with the question of why higher rates of unemployment don't result in much bargaining down of wages. Recent experience with 10 percent unemployment has made this clear. We can force the modern model of unemployment to fit the data only by adopting a bargaining model in which employers are strangely unable to take any real advantage of the increased availability of workers during slumps.

With the financial crisis fading into history while unemployment remains stuck above 9 percent, macroeconomics is starting to ask what aspect of the legacy of the crisis remains so influential. Here the emphasis is on consumers' continuing reluctance to resume buying durables and houses. Households are building their financial assets rather than using their discretionary income to buy cars, houses, and furniture. Existing models of durables purchasing behavior are receiving a new look, with particular attention paid to income uncertainty and to the capacity of households to manage that uncertainty in the post-crisis world of lost liquid assets and drastically reduced ability to borrow to bridge temporary income declines.

Perhaps the most exciting opportunities for the future of macroeconomists' contributions lie in policy design. Beyond making the now-obvious point that equity is safe and debt is dangerous in a financial system, macroeconomists are joining finance economists in thinking up clever ways to allow debt to play an important role in normal times but to lessen its danger in crises. In principle, bankruptcy handles this problem, by converting debt-holders to equity-holders in bad times, but, as we learned in September 2008, traditional bankruptcy is close to impractical for a modern financial institution. One possible answer is debt that converts to equity in times of economy-wide danger.

In the category of blue-sky thinking, a few macroeconomists, including this writer when he has nothing better to do, think about how to work around the zero lower bound on interest rates. The key policy move to eliminate the bound is for the Fed to drop its unlimited willingness to issue currency, given that currency is, in effect, a way that the federal government borrows from the public at above-market interest rates. If the Fed stopped accommodating the swelling demand for currency, the existing stock of currency would appreciate—a \$20 bill would buy more than \$20 worth of merchandise, just as a British pound buys more than a dollar today. We are still pondering how the public would react to this departure from a century and a half of government currency issuance.

References

- Bernanke, Ben S., Mark Gertler, and Simon Gilchrist.** 1999. "The Financial Accelerator in a Quantitative Business Cycle Framework." Chap. 21 in *Handbook of Macroeconomics*, Vol. 1C, ed. John Taylor and Michael Woodford. Amsterdam: Elsevier.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans.** 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113(1): 1–45.
- DeMarzo, Peter, and Darrell Duffie.** 1999. "A Liquidity-Based Model of Security Design." *Econometrica*, 67(1): 65–99.
- Hall, Robert E.** 2009. "By How Much Does GDP Rise if the Government Buys More Output?" *Brookings Papers on Economic Activity*, Fall, pp. 83–231.
- Hall, Robert E.** 2010. "The High Sensitivity of Economic Activity to Financial Frictions." The Hahn Lecture, Royal Economic Society. (Forthcoming, *Economic Journal*)
- Krishnamurthy, Arvind.** 2010. "How Debt Markets Have Malfunctioned in the Crisis." *Journal of Economic Perspectives*, 24(1): 3–28.
- Rotemberg, Julio J., and Michael Woodford.** 1999. "The Cyclical Behavior of Prices and Costs." Chap. 16 in *Handbook of Macroeconomics*, Vol. 1B, ed. John B. Taylor and Michael Woodford, pp. 1051–1135. Amsterdam: Elsevier.
- Townsend, Robert M.** 1979. "Optimal Contracts and Competitive Markets with Costly State Verification." *Journal of Economic Theory*, 21(2): 265–93.