

A CRISIS

of

BELIEFS

INVESTOR PSYCHOLOGY

and FINANCIAL

FRAGILITY



NICOLA GENNAIOLI

and ANDREI SHLEIFER

A Crisis of Beliefs

INVESTOR PSYCHOLOGY AND FINANCIAL FRAGILITY

Nicola Gennaioli and Andrei Shleifer

PRINCETON UNIVERSITY PRESS
Princeton and Oxford

Introduction

The collapse of the investment bank Lehman Brothers on Sunday, September 14, 2008, caught almost everyone by surprise. It surprised investors, who dumped stocks and brought the market index down by 500 points on Monday. It surprised policymakers, who rushed to rescue other financial institutions after declaring for months that there would be no government bailouts. It also surprised economic forecasters. Only six weeks before the Lehman bankruptcy, in early August 2008, both the Federal Reserve and professional forecasters predicted continued growth of the U.S. economy. Contrary to that prediction, the U.S. financial system nearly melted down after the Lehman bankruptcy, and the economy slid into a deep recession. This happened despite extraordinary—and ultimately successful—government efforts to save the financial system after Lehman.

Why was the Lehman crisis such a surprise? After all, fragility has been building up in the financial system for quite some time. In the mid-2000s, the U.S. economy went through a massive housing bubble. As home prices rose, households levered up to buy homes with mortgages. Banks and other financial institutions levered up to hold mortgages and mortgage-backed securities. As the bubble deflated after 2006, the financial system

experienced considerable stress, as reflected in runs on financial institutions, followed by bankruptcies, rescues, and mergers. Yet the system and the economy stayed afloat until the fall of 2008, supported by successful interventions by the Federal Reserve aimed to avoid a financial panic. By mid-2008, investors and regulators expected that, despite the deflating housing bubble, the situation was under control. On May 7, 2008, Treasury Secretary Henry Paulson felt that “the worst is likely to be behind us.” On June 9, 2008, Fed Chairman Ben Bernanke stated that “the danger that the economy has fallen into a ‘substantial downturn’ appears to have waned.”

The relative quiet before the storm, expressed in both the official and private-sector forecasts of the economy and the speeches of government officials, gives us important clues as to why Lehman was such a surprise. It surely was not the news of Lehman’s financial weakness per se, since the investment bank was in trouble and expected to be sold for several months prior to its September bankruptcy. U.S. banks more generally were making large losses for several months as the housing and mortgage markets deteriorated, and no major economic news surfaced that weekend. Nor can the surprise be attributed to the government reiteration of its “no bailout” policy. For if that were the reason for the collapse, the markets would have bounced back as soon as it became clear on Monday that bailouts were back in. In fact, markets bounced around a bit but continued their slide as the financial system deteriorated over the next several weeks, despite all the bailouts.

The evidence on the beliefs of investors and policymakers instead tells us that the news in the Lehman demise was the extreme fragility of the financial system compared to what was previously thought. Despite consistently bad news over the course of 2008, investors and policymakers came to believe that they had dodged the bullet of a major crisis. The pressures building up from home price declines and mortgage defaults

were attenuated by the belief that the banks' exposure was limited and alleviated by effective liquidity support from the Fed. The risks of a major crisis were neglected. The Lehman bankruptcy and the fire sales it ignited showed investors and policymakers that the financial system was more vulnerable, fragile, and interconnected than they previously thought. Their lack of appreciation of extreme downside risks was mistaken. The Lehman bankruptcy had such a huge impact because it triggered a major correction of expectations.

Ten years after Lehman, economists agree that the underestimation of risks building up in the financial system was an important cause of the financial crisis. In October 2017, the University of Chicago surveyed a panel of leading economists in the United States and Europe on the importance of various factors contributing to the 2008 Global Financial Crisis. The number-one contributing factor among the panelists was the "flawed financial sector" in terms of regulation and supervision. But the number-two factor among the twelve considered, ranking just below the first in estimated importance, was "underestimation of risks" from financial engineering. The experts seem to agree that the fragility of a highly leveraged financial system exposed to major housing risk was not fully appreciated in the period leading to the crisis.

These judgments are made with the benefit of hindsight. The world, however, has witnessed an extensive history of financial bubbles, expanding credit, and subsequent crises as the bubbles deflated. Errors in beliefs appear in multiple narratives. Classic studies such as Kindleberger (1978), Minsky (1977), and more recently Reinhart and Rogoff (2009) argue that the failure of investors to accurately assess risks is a common thread of many of these episodes. Rajan (2006) and Taleb (2007) stressed the dangers from low probability risks to financial stability. Even before the Lehman bankruptcy, Gerardi et al. (2008) drew attention to expectation errors in the developing subprime crisis. Since the 2008 crisis, a great deal

of new systematic evidence on credit cycles, both for the United States and worldwide, has been developed, starting with the pioneering work of Greenwood and Hanson (2013). Much of this work points to errors in expectations over the course of the cycle. Here we take this point of view further and put inaccurate beliefs at the center of the analysis of financial fragility.

To this end, we seek in this book to accomplish three goals. First, we would like to show that survey expectations data are a valid and extremely useful source of information for economic research. Expectations in financial markets tend to be extrapolative rather than rational, and this basic feature needs to be integrated into economic analysis.

Second, we seek to provide an empirically motivated and psychologically grounded formal model of expectation formation that can be used across a variety of domains, from lab experiments to studies of social beliefs to dynamic analyses of financial and macroeconomic volatility. In economics, nonrational beliefs have been typically formalized using so-called adaptive expectations, which describe mechanical extrapolation of past trends into the future. This approach has been criticized on the grounds that individuals are forward-looking in that they react to information about the future, not only to past trends. We develop a more realistic nonmechanical theory of belief formation, building on evidence from psychology. In this theory, decision makers react to objectively useful information, but in a distorted way.

Third, we use this model of expectation formation to account for the central features—including both market outcomes and beliefs—of the 2008 crisis both before and after Lehman and to explain credit cycles and financial fragility more generally. With the model of expectations we propose, many empirically established features of financial markets emerge in otherwise standard dynamic economic models. Getting the psychology right allows us to shed light on the conditions under which

financial markets are vulnerable to booms and busts. It may also help in thinking about the role of economic policy.

Expectations Data

A natural starting point for assessing the significance of financial “instability from beliefs” is to analyze the beliefs themselves. This entails not only directly measuring expectations of market participants and systematically testing whether these beliefs are rational, but also characterizing the type of mistakes (if any) that investors make.

This enterprise is feasible because a wealth of available survey data reports the beliefs of investors, corporate managers, households, and professional forecasters. These data offer important insights on whether, in 2008 and in other historical episodes, investors appreciated the risks building up before the crisis or alternatively failed to see the trouble coming. More generally, survey data help identify regular patterns in beliefs during economic fluctuations, needed to develop better theories of expectation formation and credit cycles.

Our approach is a natural extension of the long-standing research agenda in behavioral finance. Traditional behavioral finance tests the rationality of beliefs *indirectly*, by looking at the predictability of security returns. Because returns should be mostly unpredictable when markets are efficient, the consistent findings of predictability are taken to be evidence that expectations are not rational. Here we take the next step and argue that actual expectations data should become a direct target of investigation. These data can shed additional light on what investors think and how they trade, but also on market behavior. The focus on beliefs is pivotal in high leverage situations, such as the study of credit cycles, because changes in expectations can trigger massive dislocations in the financial system, as we saw after the Lehman bankruptcy.

Although rather obvious, the use of survey expectations as direct targets of economic analysis has been quite controversial in economics, for an important methodological reason. Over the past forty years, macroeconomics has been dominated by the Rational Expectations Hypothesis (REH), and finance by its close relative, the Efficient Markets Hypothesis. These theories, which represent important intellectual achievements of twentieth-century economics, hold that economic agents are rational and, as such, form their expectations about the future in a statistically optimal way, given the structure of the economy. This view has one profound consequence. It implies that expectations are dictated by the structure of the economy itself, so that survey data on expectations are redundant and noisy information. The weakness of this approach is that the REH, like any other hypothesis, cannot be just assumed to hold. Rather, as forcefully argued by Charles Manski (2004), it should be subject to empirical tests. Assessing the statistical optimality of survey data on beliefs is a natural place to start.

For the period leading to the 2008 crisis, we have a good deal of data on the expectations of homebuyers about future home price growth, on investor beliefs about the risk of home price declines and mortgage defaults, and on forecasts of economic activity made by both private forecasters and the Federal Reserve. We also have a variety of contemporaneous documents and speeches of policymakers, as well as discussions at the Federal Open Market Committee (FOMC) meetings, which shed light on the beliefs of policymakers. We can then ask directly: What were homebuyers, banks, investors, and policymakers thinking as the events leading up to the crisis unfolded?

The answers to this question cast doubt on the “too big to fail” theory of the crisis, which holds that the banks knew the risks but gambled on bailouts. The expectations of bank executives and employees seem to be very similar to those of other investors. Bankers were optimistic about

housing markets and made loans as well as personal home purchases accordingly. There is no evidence that bankers understood the risks better than anybody else.

Beliefs are more in line with the classical analyses of Kindleberger (1978) and Minsky (1977) that emphasize excessive optimism before crises. Homebuyers were unrealistically optimistic about future home price growth. Investors in mortgages and in securities backed by these mortgages, including financial institutions, considered the possibility that home prices might fall but did not fully appreciate how much and what havoc these declines would wreak. And macroeconomic forecasters from both the private sector and the Federal Reserve did not, in forming their expectations, recognize the risks facing the U.S. financial sector and the economy as late as the summer of 2008. The evidence does not suggest that investors or policymakers were totally naïve or oblivious to the risks in the financial system. Rather, they did not fully appreciate tail risks until the Lehman collapse laid them bare.

The data on beliefs prior to the Lehman crisis point to two key patterns: the extrapolation of past home price growth into the future, and the neglect of unlikely downside risks. Extrapolation of past home price growth sheds light on the housing bubble. Neglected downside risk explains how the financial system became so leveraged. This levering up of both households and financial institutions was most plausibly supported by the widely shared beliefs that the prices of homes were unlikely to collapse and that financial institutions were protected from bad shocks by diversification and hedging.

Neglect of downside risk explains how it took a year between initial bad news and the Lehman bankruptcy to ignite a financial panic. As home prices started falling, beliefs began deflating as well, leading to an unwinding of unwanted risk exposures. Starting in the summer of 2007, this unwinding led to mortgage defaults, foreclosures, fire sales of assets,

liquidations, runs on some financial institutions, and other correlates of distress. But markets did not collapse, despite the deflating housing bubble, and the financial system held together for over a year. In part, this was due to successful liquidity interventions from the Fed. But it was also due to the continued belief that banks were not vulnerable to extreme tail risks, even if home prices fell. The Lehman bankruptcy was a massive surprise precisely because it laid bare these extreme downside risks. Investors learned that they were wrong in thinking that the situation was under control. This was the making of the financial crisis and of the Great Recession that followed, driven by erroneous beliefs.

Beliefs tie together the transmission mechanisms of the crisis, which are well understood by economists (Brunnermeier 2009). Prior to Lehman, the financial system already faced significant instability, such as asset fire sales, bank runs, and rescues of failing institutions, but there was no major disruption because investors did not anticipate a full meltdown. After Lehman, the very same amplification mechanisms could no longer be controlled without capital injections, and the financial system nearly collapsed before the government injected capital to prevent massive insolvencies. Lehman was an eye opener. It proved that financial institutions were much more exposed to risk than previously thought. To understand the pivotal role of the Lehman bankruptcy in the crisis, one needs to understand the evolution of beliefs.

Looking at beliefs data also sheds light on financial fragility more broadly, beyond the 2008 crisis. A great deal of survey data on investor and professional forecaster expectations about not only stock markets, individual stocks, and credit markets, but also the real economy, are available and can be examined. The evidence presented in this book—both new and summarized from earlier studies—suggests that extrapolation of past trends is in fact a common feature of expectations held by investors, corporate managers, and professional forecasters. This is in line with

studies of other bubble episodes (Kindleberger 1978; Glaeser 2013; Greenwood, Shleifer, and You 2018). Neglect of risk, as pointed out by Rajan (2006) and Taleb (2007), is pervasive as well. The neglect of downside risk is present in several documented instances of financial innovation, such as portfolio insurance and index options (Coval, Pan, and Stafford 2014) and other episodes of credit expansion (Baron and Xiong 2017). The kinds of patterns we see in 2008 appear in other financial and economic episodes.

The expectations data actually tell us something deeper. Across many economic domains, forecast errors are predictable, even among professional forecasters. Expectations are too optimistic in good times and too pessimistic in bad times. This stands in contrast to the Rational Expectations Hypothesis, which holds that statistically optimal forecasts should use all available information, thereby avoiding predictable errors. The failure of standard economic models to account for expectations is a major gap in the analysis because it assumes away a potentially critical source of instability. There is enough evidence to take the “instability from beliefs” hypothesis seriously.

The Psychology of Expectations

The empirical challenges that expectations data present to the REH are only the beginning of the story. It takes a theory of expectations to replace the existing theory. Naïve theories of irrational beliefs cannot explain how extrapolation and neglected downside risk are connected and how they come and go, around 2008 or in general. Adaptive expectations, a theory of mechanical extrapolation of past trends, can explain the growth of the housing bubble but not why the system stayed afloat after the bubble started deflating in 2006 or why a single event such as the failure of Lehman induced such a drastic revision of expectations.

This brings us to the second goal of this book: to move from the analysis of expectations data to a new theory of expectation formation that can account for the facts. We present one psychologically founded theory of expectation formation, which we call Diagnostic Expectations. We have developed this theory over the past several years together with Pedro Bordalo and have taken it both theoretically and empirically to a number of different domains with Katherine Coffman, Yueran Ma, and Rafael La Porta. The theory is surely not the last word in modeling expectations, but it suggests that one can make some progress in understanding the reality of financial markets by moving away from the REH in a psychologically realistic direction.

In developing this model of expectations, we are guided by four principal considerations. First, we would like a theory of beliefs to be biologically and psychologically plausible, and in particular based on the evidence on human judgment obtained in experimental data. Psychologists have for decades studied judgment under uncertainty and the biases it entails, and a theory of beliefs might as well start with this evidence.

Second, we would like a theory of beliefs to be portable in the sense proposed by Matthew Rabin (2013). That is, we would like the same theory to explain evidence in psychological experiments, social judgments individuals make, financial markets, and perhaps other domains. There is no compelling reason to think that belief formation in financial markets is different from that anywhere else. One can argue, of course, that in financial markets, unlike in other domains, rational arbitrageurs profitably trade to eliminate the effects of belief distortions of irrational “noise traders” on security prices. Yet this objection has long been rejected in finance: Arbitrage is limited by capital constraints and risk aversion of arbitrageurs, and it typically does not eliminate inefficiencies in market prices (DeLong et al. 1990; Shleifer and Vishny 1997).

Third, we would like a theory in which beliefs are forward-looking. Before the rational expectations revolution, economists relied on models of adaptive expectations, in which decision makers mechanically extrapolate the past rather than react to news. These models were effectively criticized by Robert Lucas (1976), who argued, using both logic and evidence, that economic agents react to news about the future in forming their beliefs. The evidence from psychology also shows that humans do not update mechanically. They revise their beliefs about the probabilities of different events on the basis of information. The question the Lucas critique leaves open is not whether economic agents are forward-looking and react to information but rather whether they do so by the right amount. In our approach they do not.

Finally, we would like a theory of belief formation to be testable using survey evidence on beliefs. The available evidence shows that survey expectations are not noise and that both investors and managers make decisions in line with their stated beliefs. To us, these beliefs are as significant a component of empirical data that economic models need to explain as any other. A successful model of belief formation must as a start account for measured beliefs.

The model of expectations we describe builds on the famous representativeness heuristic of human judgment under uncertainty initially proposed by psychologists Daniel Kahneman and Amos Tversky in 1972. According to Kahneman and Tversky (1983), "an attribute is representative of a class if it is very diagnostic, that is, if the relative frequency of this attribute is much higher in that class than in a relevant reference class." Representativeness entails a judgment error of overestimating the likelihood of representative attributes in a class.

To illustrate, suppose someone is asked to predict the most likely hair color of an Irish person. In several informal surveys we conducted,

many people said red. It is absolutely the case that red hair is objectively more common among the Irish than among other humans: 10 percent of the Irish have red hair, compared to 1 percent elsewhere. But because red hair is a representative attribute of the Irish, people tend to believe that the Irish are even more likely to have red hair than they actually do. Judgments by representativeness contain a kernel of truth in that they respond to information in the objectively correct direction. However, they do so excessively. For this reason, people overestimate the percentage of Florida residents older than age 65 or the share of African Americans who live in poverty, and underestimate the likelihood of unrepresentative types, such as Republicans supporting abortion.

Judgment by representativeness is a universal decision heuristic, which accounts for many striking experimental findings. It is also tied to the biology of memory, which accounts for mechanisms of selective recall. Representativeness is the foundation of our theory of expectations. It creates a belief distortion that social psychologists call “the kernel of truth”: Beliefs exaggerate true patterns in the data, or, in dynamic contexts, they overreact to information. This implies that both beliefs and their errors are predictable from the underlying reality. To see this, go back to the example of the red-haired Irish. Here, people overreact to the information that a person is Irish in estimating the person’s hair color. Without the knowledge of a person’s nationality, they might have estimated that the hair color of a random person is dark, which is the most common hair color in humans. But once they learn a person is Irish, the recall of red hair is immediate because in the data, red-haired types are *relatively* much more prevalent in the Irish population than elsewhere. As a consequence, when thinking about the Irish, the probability mass shifts too far toward red hair.

Applied to expectations in macroeconomics and finance, representativeness has some distinctive implications. The kernel of truth principle

implies that people tend to overweight future outcomes that become more likely in light of incoming data. Just as they overreact to the news that a person is Irish in estimating the color of their hair, they react to macroeconomic news in the correct direction but excessively. Good macroeconomic news makes good future outcomes more representative, and therefore overweighted, in judgments about future states of the world. The converse is true for bad macroeconomic news. The same principles of belief formation that apply to lab experiments and social judgments translate one-for-one into our model of diagnostic expectations.

Under some conditions, diagnostic expectations tie together extrapolation and neglect of tail risk. News pointing to higher likelihood of economic growth causes high-growth scenarios to be representative and recessions to be unrepresentative, leading investors to both neglect downside risk and to display excess optimism about average conditions. News pointing to reduced volatility renders extreme shocks unrepresentative, leading investors to neglect risk. Diagnostic expectations also generate systematic reversals of optimism and pessimism in the absence of news. When trends in news cool off, no particular outcome is representative and expectations revert toward rationality. If the corrective news is bad enough, the left tail becomes representative and investors display excess pessimism. These movements in beliefs are entirely due to investors' overreaction to objectively useful information, not to their mechanical extrapolation of the past.

A formal model of diagnostic expectations satisfies our four criteria for a theory of expectations. It is based on extensive psychological evidence. It is portable in that the same model is applicable to lab experiments, to human social judgments such as stereotypes, and to financial markets. It offers testable predictions about the evolution of expectations in economic and financial contexts. And it is forward-looking in that it is first and foremost a theory of how people react to information. But

unlike in the case of rational expectations, the reaction to news is not statistically optimal. Rather, it is distorted by a basic principle of human judgment.

Diagnostic Expectations and Financial Fragility

Diagnostic expectations provide a useful unifying account of the 2008 crisis. They can serve as a foundation of extrapolative beliefs that characterized the housing bubble, which can be seen as updating and overreacting to repeated good news about home prices and general economic conditions. But they can also account for the neglect of downside risk, due to good news both about economic conditions (which rendered the left tail unrepresentative) and about the safety of financial institutions. More subtly, diagnostic expectations can account for the quiet period between the first tremors in housing and financial markets in the summer of 2007, which the Fed contained so successfully, and the eventual Lehman crisis. Even though the housing bubble was deflating and expectations about economic conditions were revised downward, the perception of tail risks remained dampened due to Fed policies and to the “diversification myth,” an exaggerated faith in the new insurance mechanisms. Diagnostic expectations may thus explain why both Federal Reserve and private-sector forecasts of future economic activity made as late as August 2008 point to a widely shared—and exaggerated—belief that, despite the early tremors, the situation was under control.

The theory also accounts for the extreme reaction to the Lehman bankruptcy, as the tail risks to the financial system came out into the open and market participants reacted. The Lehman bankruptcy revealed that the situation was far from being under control, that financial institutions were highly interconnected, so that systemic risk was much higher than previously expected. As a consequence, the previously neglected left tail

became representative, causing beliefs to overweight the black swan of a financial meltdown. The market panic, asset fire sales, runs on financial institutions, mergers to avoid bankruptcy, and of course government rescues can be viewed as reflecting—at least in part—the massive revision of diagnostic expectations about financial fragility. The Lehman crisis was a crisis of beliefs.

To be sure, this is just a narrative of one important episode. To evaluate the theory more thoroughly, one needs to bring it to more systematic data. Fortunately, since the 2008 crisis, economists have assembled a good number of credit cycle facts that describe the relationship between credit growth, credit valuations, and economic activity. In brief, this research finds that private credit expansion, and especially the expansion of household credit, predicts increased likelihood of future economic crises, as well as low stock returns for banks that extended that credit. This predictable fragility is reflected in debt markets as well. A high share of risky debt issuance in total debt issuance, as well as low yield spreads between risky and safe debt, are indicators of substantial market appetite for taking risk. In the data, these bullish indicators anticipate a reversion in credit conditions, meaning that high risk appetite predicts low returns and subsequent declines in issuance of risky debt. But they also predict slowdown in economic growth. This evidence points with some reliability to predictable credit cycles: Frothy credit market conditions predict both financial and economic trouble ahead. The Lehman crisis and the Great Recession were extremely dramatic, but far from unique.

The predictability of economic outcomes and security returns from credit market conditions suggests that rational expectations models might not be the most natural way to explain the data. Rather, the evidence suggests that instability comes from expectations themselves. Diagnostic expectations formalize this hypothesis, offering a unified account of

credit cycles. Good economic news, such as growth in home prices or improvement in economic conditions more generally, makes right-tail outcomes representative. This leads investors to both overestimate average future conditions and to neglect the unrepresentative downside risk, causing overexpansion of both leverage and real investment. When good news stops coming, investors revise their expectations down, even without adverse shocks. These revisions cause credit spreads to revert, the lenders to perform poorly, and economic and financial conditions to deteriorate, leading to deleveraging and cuts in real investment. A severe crisis occurs if arriving news is sufficiently bad as to render left-tail outcomes representative and hence overstated.

The fragility of the financial system here comes entirely from beliefs. Without diagnostic expectations, there surely will be market volatility in response to shocks but no major market stresses without major shocks. Diagnostic expectations change this calculus both because they point to predictable reversion in economic conditions coming from predictable reversion in expectations and because they lead to extreme financial fragility when left-tail outcomes become representative. In an intuitive way, diagnostic expectations supply a theory of a panic driven by shifting representativeness.

Roadmap

In the rest of this book, we develop these ideas more systematically and formally. Chapter 1 summarizes the basic facts about the financial crisis of 2008 and draws attention to several key facts that need to be explained, such as the year-long delay between the onset of bad news about housing and the Lehman bankruptcy. Chapter 2 then presents a variety of evidence, both from surveys and from policymaker speeches and narrative accounts, which summarizes the beliefs during this period. The main

message of this chapter is that extrapolation and the neglect of left-tail risk were empirically measurable features of beliefs during the housing bubble and that downside risk was neglected in 2007–2008, even as the bubble was deflating. Rational expectations models, such as “too big to fail” or “bank runs,” are not consistent with these facts as long as one takes expectations data seriously.

Chapter 3 introduces a model of a financial system that is standard except that downside risks are neglected. The goal of the model is to assess whether this belief distortion can account for the key features of the crisis. In chapter 3, we simply postulate neglected risk rather than derive it from a micro-founded model of belief formation. We show how the neglect of downside risk can lead not only to excessive expansion of debt that investors perceive to be safe, but also to the levering up of banks that issue such debt to investors and the need to find assets to issue profitable liabilities. When neglected risks resurface, investors realize that credit expansion is excessive and risks are misallocated between them and the banks. The unwinding of original positions causes debt prices to fall severely, triggering major financial instability. The model delivers some stylized facts about a financial crisis without any nonstandard assumptions other than an error in beliefs.

The first three chapters leave us with two questions. First, are extrapolation and neglect of risk general features of beliefs in other episodes of market volatility? The single episode of 2007–2008 allows limited inference, so assessing errors in beliefs more broadly is critical to evaluate their importance for financial fragility. Second, the model of belief distortions in chapter 3 makes several assumptions about the conditions under which beliefs display extrapolation and neglect of risk. Are these assumptions realistic and applicable in other contexts? Here we see the importance of starting with more systematic psychological foundations. These two questions motivate the rest of the book.

In chapter 4, we summarize some of the rapidly growing empirical evidence on expectations. This evidence comes from many sources, including expectations of aggregate stock market returns, expectations of earnings growth of companies, and expectations of credit market conditions. Chapter 4 establishes three major facts that are central to our analysis. First, in cases when multiple data sources about the same expectations are available, survey expectations are extremely consistent across these sources. The traditional dismissal of survey expectations data as noise is rejected by this evidence. Second, investors and managers make decisions that are in line with their stated expectations. When investors expect stock market returns to be high, they put money into equity mutual funds. When corporate managers expect high earnings growth, they plan more investment. On every count we consider, survey expectations appear to be a credible measure of the beliefs of market participants. Third, in several domains, survey expectations are extrapolative. Investors expect high returns on the stock market when past returns have been high. Corporate managers and analysts expect high earnings growth when past earnings growth has been high. Fourth, expectations are not mechanically extrapolative. They depend on the features of the underlying economic process. Beliefs exaggerate the frequency of outcomes that have become objectively more likely. Updating is also more aggressive when the underlying series is more persistent, which is again inconsistent with mechanical extrapolation. Such “sophisticated” extrapolation is a central feature of beliefs demanding explanation.

We also use chapter 4 to summarize the evidence on credit cycles. This evidence most importantly points to predictability of economic conditions and security returns from indicators of credit market frothiness, such as the growth of credit, credit spreads, and risky debt share in total

issuance. Such predictability points directly to the need for a theory of expectations.

Chapter 5 presents perhaps the central innovation of this book, which is the model of expectation formation that we developed with Pedro Bordalo. This model builds expectations from the psychological foundations up. We begin by summarizing Kahneman and Tversky's research on representativeness and then describe a formal model of judgment by representativeness and the way it transforms true probability distributions of future outcomes into diagnostically expected ones. We show how this model can account for the findings of famous experiments but more generally explains distorted beliefs such as stereotypes in social domains. We then use the model to account for the stages of the 2008 crisis: the exuberance around the growth of housing prices, the quiet period after the first bad news started to come in, and the collapse of markets after Lehman. Unlike other theories, a model with diagnostic expectations accounts not just for market outcomes during this period but also for the beliefs of the participants.

In chapter 6, we put the various pieces together by presenting a full-blown dynamic economic model of credit cycles with diagnostic expectations. In the model, good news leads to excessively optimistic expectations and overexpansion of credit and economic activity. Critically, such overexpansion is self-correcting even without adverse news, as exaggerated beliefs revert to fundamentals. The model accounts for the credit cycle facts described in chapter 4 but also suggests how representativeness can lead to sharp movements in expectations and activity when tail outcomes become representative.

Chapter 7 summarizes the book. We hope to accomplish three goals. The first is to provide a new narrative of the 2008 financial crisis that assigns a central role to beliefs in accounting both for periods of relative

quiet and those of extreme volatility. Second, we hope to demonstrate that survey expectations are a valid and extremely useful source of data for economic analysis. They can be used to test conventional null hypotheses, such as rational expectations. Perhaps more important, expectations data can be used to test and distinguish alternative models of economic fluctuations. Third, we propose to show that beliefs can be modeled in economic analysis starting with first principles of psychology, such as representativeness. Our proposed model of such beliefs—diagnostic expectations—can then be incorporated into standard dynamic economic models, such as that of credit cycles. This model is both rigorous and testable. It disciplines the analysis, as do rational expectations, but fits the evidence on beliefs considerably better.

Perhaps, however, there is a more holistic message to all of this. What we hope to show is that behavioral economics has grown up. It no longer needs to be merely a critique of neoclassical economics. A researcher can start with fundamental psychology, build formal models of beliefs, incorporate those models into standard models of markets, and bring predictions to the data. Tractable and psychologically founded models of beliefs can be integrated into the standard analysis in finance and macroeconomics. This approach may address the concerns about the discipline of the analysis while bringing to it a healthy dose of realism.

The Financial Crisis of 2008

The financial crisis of 2008 is one of the most dramatic economic events in modern U.S. history. It featured an effective meltdown of the U.S. financial system, massive government rescues of financial institutions, and a deep and long-lasting recession. Although a decade later the U.S. economy has recovered almost completely, the crisis has left behind fundamental questions about the fragility of the financial system and its impact on the real economy. In this chapter, we present the basic narrative of the origins of the financial crisis, the crisis itself, and its immediate aftermath.

We begin with the 2008 crisis for two reasons. First, despite being dramatic, this crisis is actually fairly typical, sharing the essential features of many others. It grew out of a bubble in the housing market, which was to a large extent financed with mortgage debt. The growth of mortgages led to a huge expansion of household leverage, but also leverage of banks and other financial institutions that were exposed to housing. The crisis was most directly caused by the deflation of the housing bubble beginning in 2006, leading to mortgage defaults, collapses in values of bank assets, losses of financial institutions, and ultimately fire-sale liquidations and massive distress.

The second reason to look closely is that, because the crisis is so recent, and so heavily investigated and discussed by academics, journalists, and policymakers, we know a great deal about what various participants were thinking in real time as the events unfolded. What were the beliefs of households buying homes and taking out mortgages in markets with rapidly rising prices? What was the reasoning of the commercial banks, investment banks, and insurance companies that took on massive and highly leveraged exposures to mortgage debt as home prices rose? What were the beliefs of investors and policymakers *starting* in the summer of 2007, as it became clear that the housing bubble was deflating? Why was the bankruptcy of Lehman Brothers such a pivotal event that radically changed beliefs about the markets and the economy? We have enough data on the crisis of 2008 to begin asking questions about beliefs and expectations underlying the observed behavior.

In this chapter, we summarize the basic facts about the crisis, while in chapter 2 we focus on the beliefs of the participants in this episode. The rest of the book then presents our conceptual take on these issues. Our goal in this book is to interpret historical events, and to this end we need to start with agreed-on facts.

The Housing Bubble and Housing Finance

Every narrative of the 2008 crisis begins with the housing boom. As argued by Robert Shiller, U.S. housing prices rose only moderately, in inflation-adjusted terms, for a century prior to the mid-1990s. Yet, as figure 1.1 shows, during the decade beginning in 1996, housing prices doubled in real terms. Prices rose even more dramatically in parts of the United States with inelastic housing supply, such as major cities (Saiz 2010). Figure 1.1 shows also how this dramatic decade-long increase was

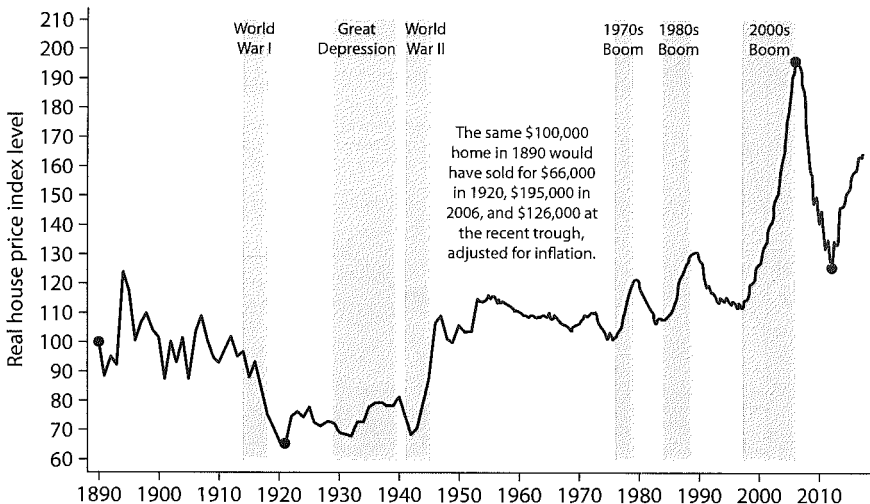


FIGURE 1.1. U.S. Home Values over Time. Robert Shiller's inflation-adjusted index of sale prices of standard existing homes in the United States. The index is set to 100 in 1890. *Source:* Shiller, Robert J. 2016. *Irrational Exuberance: Revised and Expanded Third Edition*. Princeton, NJ: Princeton University Press. Reprinted by permission of Princeton University Press.

followed by price declines that started in 2006 and accelerated in 2007. Within five years, home prices gave up much of their earlier gains.

Owner-occupied housing in the United States is largely financed with mortgages, with 80 to 90 percent of purchase prices typically paid with such debt. The rise of home prices was accompanied by equally dramatic growth in household debt, as figure 1.2 shows. Some of that growth came from home purchases by people who had never owned homes before: The rate of home ownership in the United States rose from 64 percent in 1994 to 69 percent at the end of 2006 (Federal Reserve Bank of St. Louis 2017a). Some of the rise in debt came from existing homeowners moving to more expensive residences. Some of it took the form of second mortgages, as homeowners took on additional debt to finance home improvements. Some of the debt reflected mortgage refinancing,




FIGURE 1.2. U.S. Household Debt. Total debt outstanding for U.S. households in trillions of dollars. *Source:* Board of Governors of the Federal Reserve System. 2017. “Debt Outstanding, Domestic Nonfinancial Sectors, Households—LA154104005.Q: 1945–2017.” December. <https://www.federalreserve.gov/releases/z1/default.htm>.

as households took out larger mortgages against homes whose prices had recently appreciated, so as to “cash out” and perhaps spend some of the housing equity. Last but not least, a key feature of the housing boom was investors buying multiple homes for speculative purposes, hoping to resell as home prices went up (DeFusco, Nathanson, and Zwick 2017). For all these reasons, household debt in the United States rose from \$5.7 trillion in 1998 to \$14.0 trillion at the end of 2008, and the ratio of debt to household income rose from about 0.9 to 1.3 over the same period, a truly enormous increase (Board of Governors of the Federal Reserve System 2017; Federal Reserve Bank of St. Louis 2017b).

Two additional features of U.S. housing finance during this decade deserve special mention. First, many households previously unable to get a mortgage because of poor credit scores were now able to do so. In previous times, the share of such subprime mortgages was under 10 percent.




FIGURE 1.3. U.S. Subprime Mortgage Originations. Total U.S. subprime mortgage originations are broken down into those securitized and those non-securitized. U.S. subprime mortgage originations are also given as a percentage of total U.S. mortgage originations. *Source:* Financial Crisis Inquiry Commission, 2011. *The Financial Crisis Inquiry Report: Final Report of the National Commission on the Causes of the Financial and Economic Crisis in the United States.* Washington, DC: U.S. Government Printing Office.

As figure 1.3 illustrates, at the peak of the housing boom from 2004 to 2006, between a fifth and a quarter of new mortgage originations by dollar value were subprime. In 2006, households took out \$600 billion of subprime mortgages, and the total over the decade added up to over \$2 trillion. Because they ended up defaulting at high rates, subprime mortgages play a crucial role in the narrative of the crisis.

Figures 1.1–1.3 already describe the basic elements of extreme financial fragility. When home prices fall 30 percent or more, and homeowners have borrowed 80 or 90 percent of the price, many households end up with negative equity on their homes and default on their mortgages, especially if they are experiencing financial hardship such as unemployment. Even if they do not lose their homes, their spending capacity falls sharply. Investors who hold these mortgages also lose money when homeowners

delay payments or default. The question is what happens to the financial system and to the economy when so much wealth is destroyed and this wealth is held by heavily leveraged entities.

To address this question for the 2008 crisis, we need to describe securitization, a crucial feature of mortgage finance during the housing boom. With securitization, mortgages were packaged and financially reengineered into tradable securities, which could be sold off rather than kept on the banks' balance sheets. Through impressive financial innovation, a large share of these securities were believed to be extremely safe. For this reason, securitization enabled massive increases in leverage and the fragility of the financial system.

Securitization

In the United States during this period, a bank or another financial intermediary that issued a mortgage could do one of three things with it. First, it could keep the mortgage on its own books as an asset. Second, the bank could sell mortgages to the government-sponsored housing finance enterprises (GSEs), Fannie Mae and Freddie Mac, whose liabilities were implicitly guaranteed by the U.S. government. Mortgage-originating banks did this to a substantial extent, but Fannie and Freddie typically only bought conforming mortgages (meaning not too large, not too risky, and sufficiently well documented). Fannie and Freddie sometimes retained these mortgages on their own balance sheets as assets, but more often securitized them, as we describe below. Third, the bank could sell mortgages to private investment banks, which would then securitize and sell them to investors, who thereby took on the risk of mortgage defaults.

Although all kinds of loans can be securitized, including auto loans and student loans, the most rapid growth of securitization occurred in the mortgage market. Such securitization is a financial operation that creates

mortgage-backed securities (MBS) in two steps: pooling and tranching. Pooling refers to building a portfolio or pool of mortgages out of individual ones. It creates safety by diversifying away the uncorrelated risks of individual mortgages. Tranching refers to the creation of claims of different seniority out of the pool of mortgages. This means that, to the extent that mortgages in the portfolio default or delay payments, the losses are allocated to the junior tranches first and to the senior ones only after the junior tranches are worthless. Tranching builds in additional safety for the most senior securities. The combination of pooling and tranching is so powerfully risk-reducing that senior tranches were seen as nearly completely safe. Although securitization was not invented during the housing boom, it achieved unprecedented magnitude in the creation of new MBS products, including both agency MBS issued by GSEs and private-label MBS issued by private originators.

The exact reasons for the enormous demand for securitization during this period continue to be debated, but it certainly originated in a huge demand for assets that investors perceived to be safe. Some of that demand came from foreign investors, especially from high-saving Asian countries, who poured trillions of dollars into U.S. markets over this period (Bernanke et al. 2011). Some of the demand came from the increasingly wealthy U.S. households and corporations wishing to keep a share of their wealth in safe assets. By manufacturing AAA-rated senior tranches of MBS, securitization met this demand. Depending on the estimates, 60 percent or more of MBS by dollar value created through the process ended up being rated AAA (Fitch Ratings 2007). By comparison, only about 1 percent of corporate bonds by dollar value receive such ratings (Fitch Ratings 2007).

One might ask what happened to the junior tranches of MBS created in securitizations—the tranches designed to absorb the initial losses when some mortgages in the pool default. Some of these tranches were retained

by the financial intermediaries, exposing them to concentrated mortgage risks. But even more were themselves securitized, creating so-called collateralized debt obligations (CDOs). For this, an intermediary would assemble the junior tranches of MBS into another pool and then tranche that pool to create a variety of claims. The senior tranches of CDOs, impressively, were also rated AAA. In the mid-2000s, nearly 70 percent of CDOs by dollar value—not just MBS engineered from original mortgages but CDOs engineered from risky tranches of MBS—were rated AAA. CDOs ended up being held disproportionately by financial institutions and became a major source of their losses in the crisis.

The demand for securitization products was met with supply from both quasi-public and private sources. GSEs securitized conforming mortgages into agency MBS, and then sold them off into the market, keeping some on their own books. Because of implicit government guarantees, agency MBS were perceived as a close substitute for U.S. Treasuries. More consequentially, securitization of mortgages was pursued privately by investment banks using the riskier subprime mortgages as collateral, converting them to a large extent into AAA-rated private-label MBS. These AAA ratings were based not on government guarantees, but on models of default used by the rating agencies.

Securitizations in the mid-2000s reached staggering volumes. By 2007, over \$4 trillion of agency MBS was outstanding, in addition to over \$2.5 trillion of private-label residential MBS. At its peak in 2006, private-label residential MBS issuance exceeded \$1.2 trillion, substantially higher than the issuance of corporate debt. Perhaps even more remarkably, a substantial fraction of these securitizations were backed by subprime mortgages, those too risky to sell to the GSEs. As we saw in figure 1.3, about 75 percent of subprime mortgage originations in the United States were securitized in 2006. In that year, close to half a trillion dollars of subprime mortgages were pooled and tranced into AAA-rated MBS.

Aside from its enormous magnitude, securitization plays an outsized role in the narrative of the financial crisis for three distinct reasons. First, there is the chicken-and-egg problem of whether the growth of credit to subprime borrowers drove the housing bubble or whether alternatively home price growth drove the credit expansion by increasing the value of available collateral. There is growing evidence that loose credit had an influence on home prices—see Mian and Sufi (2017) for a summary of extensive research supporting this view. On the other hand, home prices started growing before the huge expansion of subprime lending in the early 2000s period of very low interest rates, and this growth likely contributed to the expansion of collateral and credit, so there are causal factors running from home prices to credit as well. Our best assessment is that home price growth and credit growth were mutually reinforcing.

Second, the demand for securitized products encouraged fraud and other questionable practices in the subprime mortgage market, including the misrepresentation of borrower income, borrower credit rating, borrower downpayment, and perhaps other information (Keys et al. 2010; Bubb and Kaufman 2014; Piskorski, Seru, and Witkin 2015; Mian and Sufi 2017). There is an ongoing debate whether demand for MBS drove these misrepresentations. How much of the financial crisis can be attributed to such fraud, and how much to home price declines, remains an open question.

Third, and critical from our viewpoint, securitization created a massive misallocation of risk. Investors including financial institutions, hedge funds, sovereign funds, and money market funds that bought AAA-rated MBS or CDOs did so because they demanded safe securities, for reasons of preferences or regulations. When home prices began falling and AAA-rated MBS and especially CDOs lost their value and ratings, these investors not only lost money but ended up holding securities they did not want or even were not allowed to hold for regulatory reasons.

This caused them to sell, leading to major price declines. Risk misallocation is a central feature of our model in chapter 3.

This discussion sums up what securities were issued as a consequence of the housing boom and how they were valued. But who bought this stuff and how did they pay for it?

Buyers of MBS

Figures on the holders of private-label MBS are difficult to come by, but it appears that banks, hedge funds, investment banks, and other leveraged financial intermediaries were significant buyers. They held some of the MBS on their own books. In addition, banks and other intermediaries set up separate legal entities, so-called structured investment vehicles (SIVs), to buy private-label MBS. An SIV would buy a portfolio of predominantly but not completely AAA-rated MBS or CDOs, which would become its assets. The SIV would then finance its assets by issuing short-term liabilities, typically in the form of asset-backed commercial paper (ABCP), which needed to be rolled over every few weeks. As long as money market funds or other buyers of ABCP were willing to roll it over, the bank that owned the SIV would collect the spread between the yield on its assets and the cost of short-term finance. In addition, the banks typically provided SIVs with guarantees of funding in case ABCP funding evaporated. However, banks were required to hold far less regulatory capital against these guarantees and off-balance sheet exposures than would be required if they directly held these assets on their balance sheets. During the crisis, SIVs turned into a major problem for banks, which ended up absorbing SIV losses as MBS and CDO prices plummeted (Acharya, Schnabl, and Suarez 2013).

The International Monetary Fund (IMF) estimates that, prior to the crisis, U.S. banks bought \$1.5 trillion of residential MBS, and another

\$196 billion of commercial MBS, roughly 20 percent of what was outstanding at the time. An even higher fraction, perhaps 25 percent, of CDOs was bought by banks, particularly in less senior tranches. Erel, Nadauld, and Stulz (2014) estimate that at the peak of the crisis, up to \$250 billion of private-label MBS was in SIVs, while Krishnamurthy, Nagel, and Orlov (2014) suggest that 40 percent of AAA-rated private-label MBS went into SIVs. This came on top of other real estate exposures in the form of non-securitized household mortgages and commercial real estate.

The growth of assets on the balance sheets of major commercial and investment banks was reflected in the growth of their leverage. Greenlaw et al. (2008) report that by 2007 the three major investment banks (Goldman Sachs, Morgan Stanley, and Lehman Brothers) saw their equity drop to 3 to 4 percent of their assets. For commercial and investment banks, growth in leverage parallels the growth in assets more generally and not just in this period (Adrian and Shin 2010).

In sum, securitization of potentially risky mortgages, and the large exposure of highly leveraged financial institutions to these securities, had the effect of concentrating the risk of the bubble deflating on the financial intermediaries.

The System under Pressure

Home prices peaked in mid-2006 and started declining afterward, as we saw in figure 1.1. This decline infected the financial system. Not surprisingly, changes in home prices showed up first in the rapid increase in default rates on subprime mortgages, illustrated in figure 1.4. In 2006, annual default rates on these mortgages were just 2 percent, justifying the perceptions of safety implicit in the AAA ratings and the pricing of subprime MBS and CDOs. In 2007, these rates doubled to about

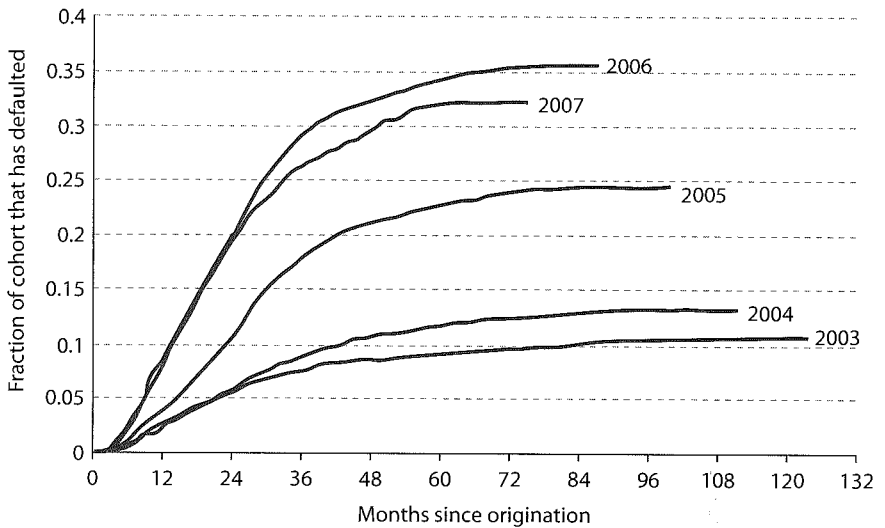


FIGURE 1.4. Cumulative Default Probability by Origination Cohort. The figure tracks the default behavior of mortgages originated in five different years (cohorts). A loan is considered to default when it is indicated to be in foreclosure of real estate owned and as long as it is not later repaid in full. *Source:* Palmer, Christopher J. 2015. “Why Did So Many Subprime Borrowers Default during the Crisis: Loose Credit or Plummeting Prices?” Working paper, University of California–Berkeley, Berkeley, CA, September. <http://faculty.haas.berkeley.edu/palmer/papers/cpalmer-subprime.pdf>. Reprinted by permission of the author.

4 percent, and eventually rose to 16 percent in 2008 and 2009. These defaults were the highest for the most recent vintages of mortgages, those issued in 2006 and 2007. Within five years, close to a third of subprime mortgages issued in these years had defaulted.

Initially, mortgage defaults led to sharp reductions in the prices of junior tranches of MBS. As the critical figure 1.5 shows, these prices had already fallen sharply by mid-2007. But by the summer of 2007, and increasingly toward the end of 2007, the high rates of default brought about declines in the prices of AAA-rated senior tranches as well. Eventually, junior tranches of subprime residential MBS became essentially




FIGURE 1.5. Prices of Residential Mortgage-Backed Securities. The figure shows the Markit ABX Home Equity Index (ABX.HE), which tracks the performance of residential mortgage-backed securities (RMBS) of different issuer ratings. *Source:* Markit. 2017. “ABX Home Equity Index (ABX.HE): 2006–2017.” Accessed August 15, 2017. <https://ihsmarkit.com/products/markit-abx.html>.

worthless, while the senior ones fell to as low as a third of their initial value before recovering starting in 2009.

The financial system came under stress almost immediately once the prices of AAA-rated MBS began to decline, and strains began to be felt by late spring and especially by August 2007. Subprime mortgage lending, and the creation of subprime MBS from these mortgages, came to a halt. In April 2007, a leading subprime mortgage lender, New Century, filed for bankruptcy protection. In June 2007, the investment bank Bear Stearns suspended redemptions from its funds holding MBS. In July, Countrywide Financial, a major underwriter of subprime mortgages, warned of “difficult conditions”; it was eventually acquired in January 2008. Interbank lending rates rose sharply. More worrisome, the declines in MBS values put huge stress on the short-term financing of MBS holdings, especially in SIVs. As figure 1.6 shows, in the summer




FIGURE 1.6. Collapse of the Asset-Backed Commercial Paper Market. The figure shows outstanding asset-backed commercial paper. *Source:* Federal Reserve Bank of St. Louis. 2017c. "Asset-backed Commercial Paper Outstanding—DTBSPCKAM: 2001–2017." Board of Governors of the Federal Reserve System. <https://fred.stlouisfed.org/series/DTBSPCKAM>.

of 2007, the asset-backed commercial paper market froze, so SIVs could no longer maintain their MBS holdings without injections of cash from their sponsors. They had to get this liquidity or, alternatively, liquidate their holdings.

The Federal Reserve responded to the problems in financial markets. Starting in late 2007, it intervened energetically by cutting interest rates, supplying liquidity, and arranging acquisitions of troubled institutions. Critically for what was to happen later, the Fed described the problem as one of liquidity, meaning that banks and SIVs had trouble rolling over their short-term liabilities that financed their MBS holdings. These short-term liabilities needed to be guaranteed or replaced. The cocktail of lower interest rates, collateralized loans to institutions seeking liquidity, and rescues of troubled financial institutions through acquisitions by healthier ones appeared to stabilize the financial system by early 2008.




FIGURE 1.7. VIX. The Chicago Board Options Exchange Volatility Index (VIX) captures the stock market's expectation of volatility. It is computed using S&P 500 index options. *Source:* Federal Reserve Bank of St. Louis. 2017d. "CBOE Volatility Index: VIX—VIXCLS: 1990–2017." Chicago Board Options Exchange. <https://fred.stlouisfed.org/series/VIXCLS>.

Figure 1.7 presents the VIX index describing expectations of stock market volatility implied by S&P 500 options. It shows the index rising substantially in the summer of 2007, but then stabilizing for a year and actually declining in the spring and summer of 2008, prior to the Lehman crisis. The behavior of VIX demonstrates two points. First, the market expected higher volatility after the housing bubble started deflating and the turbulence in the summer of 2007, but it did not expect anything like the extraordinary volatility after the Lehman bankruptcy. Second, while investors did not feel that the financial system was all the way back to normal, the Fed interventions effectively reduced risk perceptions until Lehman.

A prescient paper by Greenlaw et al. (2008) was presented at a University of Chicago–sponsored conference well attended by Fed officials

in late February 2008. The paper pointed to severe stress in markets for short-term financing, but also centrally focused on mortgage losses facing banks. The paper's authors noted that some estimates indicated mortgage losses of \$250 billion by the end of 2007, and some doomsday scenarios forecast total mortgage losses of up to \$500 billion. The authors also quite accurately predicted that such severe losses would lead to the withdrawal of short-term finance, necessitating the liquidation of MBS portfolios in fire sales, leading to further price declines and deterioration of balance sheets of intermediaries. They warned that the U.S. financial system might be facing insolvency and not just illiquidity if capital was not rebuilt. In their view, this posed a serious threat to bank lending and the economy.

During the next six months, home prices fell further, defaults on subprime MBS continued to mount, and MBS prices continued to fall, including now those of senior tranches (figure 1.5). Short-term financing and interbank markets continued to experience severe stress. The Fed massively expanded its liquidity operations. One important event during this period was the rescue of one of the largest investment banks, Bear Stearns, through acquisition by JPMorgan Chase, which required substantial assistance from the Federal Reserve. Remarkably, despite all the stress, the financial system held together through the summer of 2008.

Lehman

The calm lasted until early September. On September 7, the GSEs Fannie and Freddie were placed into government conservatorship, essentially nationalized. This was perhaps the single most significant step in fixing the financial system, but it came too late. Over the weekend of September 13–14, the Fed and the Treasury failed to successfully organize a res-

cue of one of the largest investment banks, Lehman Brothers, which declared bankruptcy on September 14. Rumors that Lehman, with its huge exposure to MBS and real estate, was illiquid and perhaps insolvent circulated for several months, arguably since the rescue of Bear Stearns (Sorkin 2009). But it was widely assumed that Lehman was too big to fail, and as such would be rescued by the government through a sale to a larger and healthier bank, dismemberment, and at least liquidity guarantees. After all, it was a bigger and more financially interconnected bank than Bear Stearns.

Except that in the Lehman case, the government did not save it. The reasons for the decision not to rescue Lehman remain controversial. U.S. Treasury Secretary Henry “Hank” Paulson blamed British officials for a failure to approve the takeover of Lehman by Barclays, a British bank, over the critical September 13–14 weekend. Many believed at the time that the senior management of Lehman was so arrogant and divorced from reality in its failure to sell the bank earlier that a furious Paulson wanted to teach them a lesson (Sorkin 2009). Paulson was adamant, both before and during that weekend, that moral hazard was rampant and that, unlike in the Bear Stearns case, the U.S. government would not be putting up any money to facilitate a takeover of Lehman, possibly discouraging some suitors. He and other government officials faced intense political pressure from Congress not to rescue Lehman.

It is an open question which of these arguments prevailed over that fateful weekend, but a few weeks later, as it became clear that the failure of Lehman precipitated a financial meltdown, the principal policymakers offered an explanation. According to this account, the Fed had no legal authority to rescue Lehman either by buying some of its assets in a takeover by a third party or by guaranteeing its liquidity because it could only do so if confident that Lehman was solvent. Because the officials did

not have such confidence, they were in no position to save Lehman. As Ben Bernanke, then the Federal Reserve chairman, later said with respect to Lehman, “We did everything we could” (Bernanke 2015).

Laurence Ball (2018) questions this account. He argues that the law is ambiguous, and that the Fed and Treasury officials did not try very hard to find out whether Lehman could be legally saved. In fact, the government was adamant for months that it would put no money into a Lehman rescue. Ball suggests that the story that the government had no legal authority to rescue Lehman was made up after the fact, once the catastrophe became apparent. Ball’s story is bolstered in part by the fact that the largest insurance company, AIG, was rescued right after the Lehman bankruptcy, and the fears about its solvency were perhaps as real. Given the incompleteness of the record, it is hard to learn exactly what is true, and we might never know. But the effects of Lehman’s bankruptcy were extreme.

Lehman’s bankruptcy shocked financial markets, and its effects obviously surprised both investors and regulators. Lehman’s failure caused the stock market to drop 500 points on Monday, September 15, and triggered a series of major rescues of financial institutions. As figure 1.7 illustrates, the perception of volatility as measured by VIX exploded. At the same time as it chose not to save Lehman, the government arranged the rescues of an even larger investment bank, Merrill Lynch, through its acquisition by Bank of America, as well as of the large insurance company, AIG, which was selling insurance against the default of mortgage-backed securities. In both instances, the immediate problem was the demand by creditors of the institutions for more collateral, prompted in part by concerns about their solvency.

The failure of Lehman precipitated the second largest financial crisis in U.S. history after the Great Depression. Financial markets were chaotic and disrupted, and short-term financing came to a halt except for that using government securities as collateral. Figure 1.8 shows the mas-

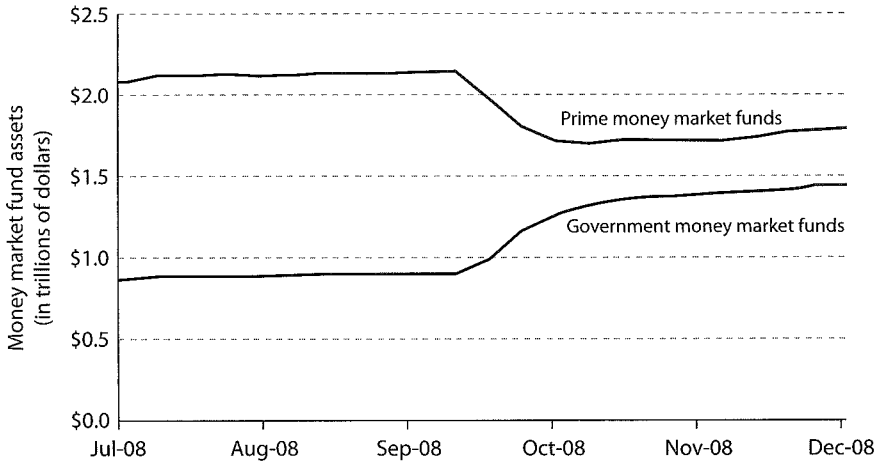


FIGURE 1.8. Money Market Fund Assets. The figure shows how investors ran from prime money market funds that lent to holders of MBS toward government money market funds. Source: Investment Company Institute. 2008–2012. “Weekly Money Market Fund Assets.” www.ici.org/research/stats.

sive reallocation of investment in money market funds from so-called prime funds, which provided short-term financing to investment banks like Lehman for their MBS holdings, toward funds that invested in U.S. Treasuries. In the course of a few weeks, about half a trillion dollars of liquid savings left these types of funds as investors flew to the safety of government debt and away from securities they previously but no longer thought to be safe.

Figure 1.9 presents the evidence on haircuts for repo, the overnight loans collateralized by securities. A haircut of near zero means that a lender is willing to lend almost a dollar overnight backed by a dollar of securities as collateral. A haircut of 40 percent means that the lender is only willing to extend an overnight loan of 60 cents against the same collateral. Clearly the markets saw a severe deterioration of the terms of short-term financing of MBS holdings, which meant that, without large capital commitments, these holdings could not be maintained.

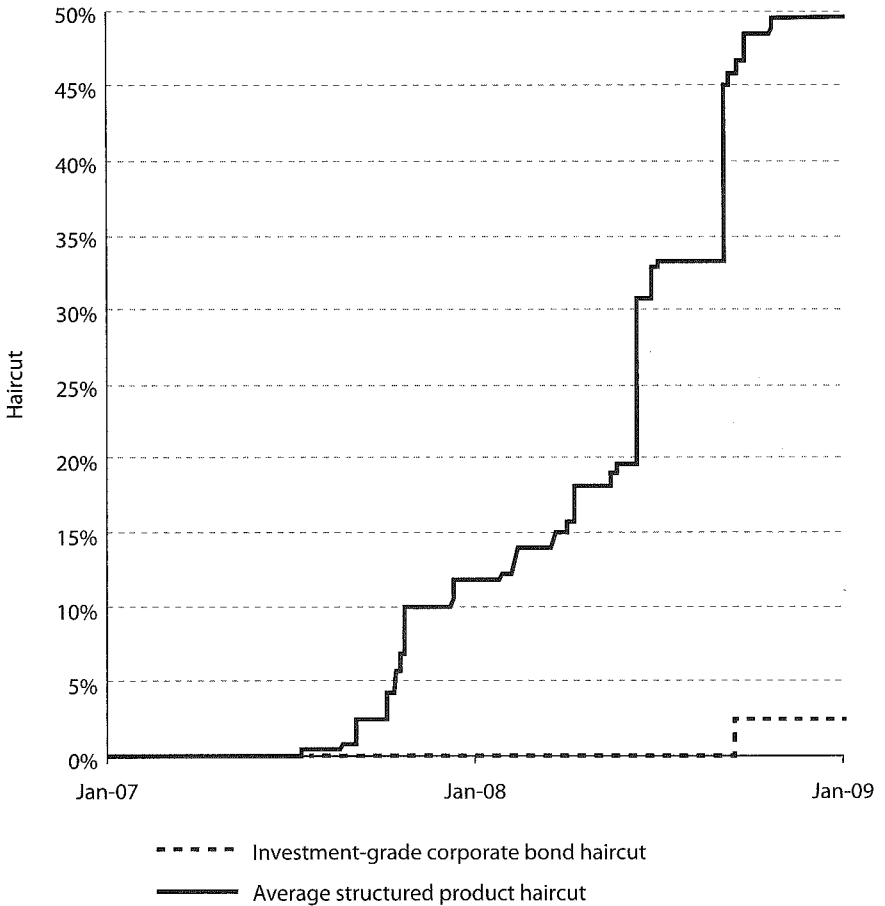


FIGURE 1.9. Repo Haircuts. The figure shows bilateral repo haircuts on structured products and investment-grade corporate bonds. *Source:* Gorton, Gary, and Andrew Metrick. 2010a. “Haircuts.” *Federal Reserve Bank of St. Louis Review* 92 (6): 507–19. Reprinted by permission of the authors.

Why was Lehman so pivotal? Surely there is nothing special about Lehman per se. We would have seen the same market distress if it were AIG or Merrill Lynch. So, what was the real news in the Lehman bankruptcy that caused such a sharp market reaction that surprised both investors and policymakers? What did the markets and policymakers learn that they did not know before? Certainly the news was *not* that

Lehman was in deep trouble, since this was known for months. The policymakers in fact were publicly trying to convince Lehman to merge with another company for close to half a year. Nor was it news that the housing bubble was deflating, since by September this was very old news as well. Nor was it news that financial institutions were losing hundreds of billions of dollars, since that too was widely understood at the time.

Markets appear to have learned two things from the Lehman filing. The first is that the government was willing to let a systemically relevant financial institution go bankrupt. Although Secretary Paulson repeatedly and extremely publicly insisted that there would be no more bailouts, markets did not quite believe him. Most investors thought that the government would blink. Ironically, the government made a U-turn on this policy immediately after the Lehman bankruptcy.

The failure of Lehman was probably news in another way. What became clear over that fateful week is that the financial system was in a meltdown that was very hard to control. Investors, and possibly regulators as well, learned that financial institutions facing hundreds of billions of dollars in losses were highly interdependent, both because of their complex derivatives and other contracts with each other and because they were holding similar assets. The bankruptcy of Lehman accelerated the unwinding of derivatives contracts and sales of assets by investors who sought to rebuild their own capital and liquidity. This process rekindled fire sales and made the meltdown even harder to stop. There is little evidence in the pre-Lehman record that policymakers fully appreciated the system-wide liquidation forces that would be unleashed.

In response to rapidly falling asset prices and the evaporation of short-term financing, banks and investment banks had to sell assets to repay their creditors. Such fire sales led to further declines in MBS prices, the further need to liquidate, and precisely the vicious circle accurately described by Greenlaw et al. (2008) in February 2008. Figure 1.10, which




FIGURE 1.10. Residential Mortgage-Backed Securities (RMBS) Fire Sale. AAA-rated RMBS prices (as captured by the ABX.HE RMBS index) exhibit a classic fire sale. *Source:* Markit. 2017. “ABX Home Equity Index (ABX.HE): 2006–2017.” Accessed August 15, 2017. <https://ihsmarkit.com/products/markit-abx.html>.

will be critical to our analysis in chapter 3, shows the price patterns of AAA-rated MBS between January 2008 and January 2013, when prices recovered. From the start of 2008 to the spring of 2009, prices of the safe tranches of MBS collapsed by nearly 70 percent, only to recover by 120 percent in the following few years. This picture looks like a classic fire sale, with the typical V-shaped price pattern.¹

The Fed and the Treasury responded aggressively to the Lehman fallout, recapitalizing financial institutions, making equity injections, arranging mergers, providing loans against risky collateral, and eventually buying up hundreds of billions of dollars of agency MBS and putting

1. Merrill et al. (2014) analyze a data set of MBS transactions by insurance companies, whose risk exposures are heavily regulated, and find direct evidence of fire sales of stressed securities.

them on the Fed's balance sheet. This ended up being sufficient to stabilize the financial system by the spring of 2009, a truly extraordinary accomplishment for which both the Fed and the Treasury officials justly received considerable praise. These accomplishments should not be underestimated; it is easy to imagine the U.S. financial system and economy going into a decades-long decline if the problems were not fixed.

But massive damage was done. The U.S. economy slid into a major recession, which came to be known as the Great Recession, and did not recover for several years. Over a decade after the crisis, the loss of output relative to trend was estimated at \$2 trillion. Housing prices continued to decline; construction virtually stopped and then remained moribund even longer than the rest of the economy. Unemployment peaked at 10 percent in October 2009 (U.S. Bureau of Labor Statistics 2012).

Several studies have explored the causal links from the financial crisis to the Great Recession. Two mechanisms have received substantial support. The first is the standard bank lending channel: As banks' balance sheets deteriorated due to their exposure to housing and MBS, they curtailed lending to firms and households, which in turn cut investment in physical capital and consumption of durables. Studies such as Chodorow-Reich (2014) and Benmelech, Meisenzahl, and Ramcharan (2017) provide supportive evidence for this channel. The second is the household debt channel: As families lost their equity in housing, and even defaulted, they cut their spending, thus deepening the recession. This channel was investigated in a number of compelling studies by Atif Mian and Amir Sufi (2009, 2011, 2014a), culminating in a 2014 book. They suggest that the United States would have had a major recession even without the financial crisis, since the households lost so much wealth after the housing price collapse. There is some evidence that corporate losses also contributed to the recession (Giroud and Mueller 2017), but the evidence here is less extreme than that for financial institutions and households.

Financial System Losses

It is useful to estimate the magnitude and timing of the financial system losses. The answer informs the crucial question of whether the market collapse could have been anticipated before the failure of Lehman. The answer also helps quantify the magnitude of the amplification mechanisms, such as runs and fire sales, which brought the financial institutions close to insolvency.

We rely on the estimates of financial system losses from the International Monetary Fund, which also enables us to see how the magnitudes evolved over time. Already in October 2007, the IMF projected mark-to-market losses for all holders of asset-backed securities (including mostly MBS) and CDOs of about \$200 billion. In March 2008, the IMF provided more detailed estimates. It estimated about \$720 billion of global losses on all securities (including corporate debt), of which about half, or \$340 billion to \$380 billion, would accrue to banks. Most of that, according to the IMF estimates, would come from subprime MBS and CDOs. It also appears that the IMF expected half of those total bank losses to accrue to U.S. banks—roughly \$200 billion. These estimates are similar to the baseline scenario of Greenlaw et al. (2008). Losses on this scale arguably did not constitute an existential threat to banks, but would have created a major hole in bank capital and a major obstacle to future bank lending.

By October 2008, these loss estimates had risen sharply—and these are the numbers before the asset liquidations went into full swing. The total expected loss estimate rose to \$980 billion; the bank share stayed at half of that or roughly \$500 billion. Interestingly, the IMF forecast substantial losses on AAA-rated MBS as well. At that point, the IMF also estimated bank loan write-downs (separately from securities) of up

to \$300 billion. If one attributes half of the losses to the U.S. banks, projected U.S. bank losses would add up to \$400 billion. This gets closer to an existential threat—even before the full effects of Lehman worked through the system.

In April 2009, the IMF estimated total write-downs of \$1.07 trillion on loans, and \$1.64 trillion on securities. The IMF also revised sharply upward the share of these losses accruing to banks, in part because of the massive declines in the prices of CDOs that the banks held disproportionately. Specifically, the IMF projected loan losses of \$600 billion and security losses of \$1 trillion for banks. If we take half of this share to be U.S. banks, the projected losses rise from \$400 billion to \$800 billion. In fact, in October 2009, the IMF estimated the losses of U.S. banks as \$654 billion for loans and \$371 billion for securities, roughly \$1 trillion total.

Several points about these numbers should be mentioned. First, one can say that even losses of \$1 trillion are not huge in the context of the entire U.S. financial system, where the stock market value can move by that amount in the course of a few weeks or even days without major disruptions. But this calculation misses the essential and obvious point that banks are heavily leveraged, and losses of even a few percent of their capital can threaten both their solvency and their ability to do business. Even absent the existential threat, banks heavily exposed to subprime severely curtailed their lending, as Chodorow-Reich (2014) and others show persuasively.

Second, we focused on the risks to U.S. banks rather than the losses of other financial institutions and foreign banks. Since security markets are integrated, to the extent that various leveraged investors sell together, price declines and losses affect them all. While we focused on banks, fire sales bankrupted several hedge funds, mutual funds, and even insurance companies.

Third, while Lehman precipitated a crisis that would have probably destroyed the U.S. financial system without major government interventions, the projected losses were both huge and hugely threatening to both the financial institutions and the economy even before Lehman. These losses represented risks to solvency and survival that were rising over several months, a point which will become central to our interpretation of the evidence.

Fourth, both before but especially after the Lehman bankruptcy, the losses became much more extreme because of fire sales and liquidations, as short-term financing of security holdings evaporated. These fire sales, together with deteriorating fundamentals, explain how losses mounted from \$400 billion to as much as \$1 trillion for U.S. banks alone. In retrospect, the Greenlaw et al. (2008) worst-case scenario in February 2008 turned out to be rosier than reality.

To summarize, what the U.S. economy experienced during this period is an almost classic credit crisis. The housing bubble, and its subsequent deflation, are the obvious cause. This bubble was financed with mortgages, substantially converted into MBS, which led to the tremendous growth in leverage of both households and financial institutions. As home prices started falling, the leverage cycle unwound, leading to massive losses in the financial system, liquidations, and declines in asset prices. As the financial system collapsed, only massive government interventions succeeded in reviving it.

In the period leading to the collapse of the bubble, the economy (including both households and financial institutions) was bearing a great deal of home price risk. Yet market participants appear not to have fully appreciated the extent of this risk, judging by the vast demand for and manufacture of AAA-rated MBS and especially CDOs that assigned a very low price for that risk, as well as the willingness of the financial sector to retain massive exposure to home price risk. Nor does the evi-

dence suggest that either market participants or policymakers appreciated the magnitude of the exposure of financial institutions to housing risk, and the interdependencies between them that posed major systemic risk once home prices started falling. But these are just impressions based on behavior rather than direct evidence on beliefs.

In chapter 2, we take a closer look at the evidence of what various parties said, modeled, and predicted. This evidence on beliefs clarifies the choices of investors and policymakers, as well as the critical market outcomes. What were the expectations sustaining the housing bubble? What was the thinking behind the massive production of AAA-rated assets from risky mortgages? What were the beliefs in 2007–2008 when it became clear that the housing bubble was deflating, and why did the markets remain calm? Why was Lehman so pivotal in triggering the crisis? Our goal is to understand the financial crisis based on much more direct evidence of what the participants were thinking.