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.