

# Psychology and the Financial Crisis of 2007-2008

Nicholas Barberis<sup>1</sup>

This draft: August 2011

## Abstract

I discuss some ways in which ideas from psychology may be helpful for thinking about the financial crisis of 2007-2008. I focus on three aspects of the crisis: the surge in house prices in the years leading up to 2006; the large positions in subprime-linked securities that many banks had accumulated by 2007; and the dramatic decline in value of many risky asset classes during the crisis period. I review a number of psychology-based mechanisms, but emphasize two, both of which have already been extensively studied in behavioral finance and behavioral economics: over-extrapolation of past price changes; and belief manipulation.

---

<sup>1</sup> Yale School of Management. This essay is in preparation for *Financial Innovation and Crisis* (MIT Press, Michael Haliassos, ed.). It is based on an invited talk at a conference in 2009 on the occasion of the awarding of the Deutsche Bank Prize in Financial Economics to Robert Shiller. I am grateful to Malcolm Baker, James Choi, Kent Daniel, Michael Haliassos, Jonathan Ingersoll, Andrew Metrick, and Andrei Shleifer for very helpful comments.

The field of behavioral finance investigates whether certain financial phenomena are the result of less than fully rational behavior on the part of some agents in the economy. For guidance on *how* people deviate from full rationality, it advocates a close reading of research in psychology. The field has focused, with some success, on three areas of application: the pricing of financial assets; the portfolio choice and trading decisions of investors; and the behavior of firm managers.

Can behavioral finance offer a useful perspective on the financial crisis of 2007-2008? In particular, can ideas from psychology help us to make sense of the crisis? I suspect that they can, but it is still too early to be sure. The process of gathering and analyzing data from the crisis period is far from over. Researchers may eventually conclude that psychological factors *were* important during the crisis; but they may also conclude that they were not of first-order significance. In this short essay, I speculate about some ways in which concepts from psychology *may* be helpful for understanding the crisis. I do not attempt a comprehensive discussion, but simply sketch a few specific ideas.<sup>2</sup>

## **Bubbles**

A central element in many discussions of the financial crisis is the idea that there was a real estate “bubble”: that, by 2006, due to a friction of some kind, or to irrational thinking, real estate prices had been pushed up to unsustainably high levels. According to a common narrative, the bubble burst, triggering widespread defaults on subprime loans, dragging down the value of banks’ subprime-linked holdings, and setting off a run on the banking system.<sup>3</sup>

While many commentators are very confident that we witnessed a bubble, it is hard to know for sure that this was the case. Still, it is a reasonable hypothesis. And the

---

<sup>2</sup> See Akerlof and Shiller (2009) and Shefrin (2009) for more extensive psychology-based accounts of the crisis.

<sup>3</sup> The term “bubble” is widely used but rarely defined, probably because it is hard to define. A working definition for the purposes of this note is: a bubble is an episode in which irrational thinking or a friction causes the price of an asset to rise to a level that is higher than it would be in the absence of the friction or the irrationality; and, moreover, the price level is such that a rational observer, armed with all available information, would forecast a low long-term return on the asset.

fact that a bubble may have played a critical role in recent events has, not surprisingly, led many observers to call for more research on why bubbles form.

I agree with this call for action. At the same time, it is important to note that there is *already* a lot of research, some of it done by behavioral finance scholars, about the formation of bubbles -- in short, about why an asset class might become overvalued. The problem may not be that we lack theories of bubble formation, but rather that we have too many such theories. As a result, rather than rushing to develop entirely new theories of bubbles, we should perhaps first test and refine the theories we already have.

It may be useful to list some of the theories of asset market overvaluation that already exist in the behavioral finance literature. The theories can be categorized based on whether they focus on investor *beliefs* or on investor *preferences*.

On the beliefs side, there are perhaps three main theories. The first argues that a bubble forms when investors disagree sharply about an asset's future prospects *and* there are short-sale constraints (Miller, 1977; Harrison and Kreps, 1978; Scheinkman and Xiong, 2003; Hong and Stein, 2007). The logic is straightforward. Suppose that some investors are very bullish about an asset's prospects, while others are very bearish. In the presence of short-sale constraints, the price of the asset will only reflect the views of the bullish: bearish investors will stay out of the market. In other words, the asset will be overvalued.<sup>4</sup>

Another belief-based theory of overvaluation argues that bubbles arise because investors extrapolate past outcomes -- returns, earnings growth, or default rates -- too far into the future (Lakonishok, Shleifer, Vishny, 1994; Barberis, Shleifer, Vishny, 1998; Greenwood and Hanson, 2010). This assumption is usually motivated by Kahneman and Tversky's (1974) representativeness heuristic. According to this heuristic, people expect even small samples of data to reflect the properties of the parent population. As a result, they draw overly strong inferences from these small samples, and this can lead to over-extrapolation. The accompanying article by Barberis and Shleifer (2003) presents a

---

<sup>4</sup> There is a second channel through which disagreement and short-sale constraints can lead to overvaluation. If there is disagreement about an asset's future prospects, the asset's price will be higher than its holders' valuation of its future cash flows because these holders believe that they can sell to other more optimistic investors in the near future. This dynamic channel is the focus of Harrison and Kreps (1978) and Scheinkman and Xiong (2003).

model of bubble formation based on over-extrapolation of past returns, itself motivated by representativeness.

A third belief-based theory of bubble formation is based on overconfidence -- specifically, on the idea that people overestimate the precision of their forecasts (Daniel, Hirshleifer, Subrahmanyam, 1998). According to this theory, when investors, in an effort to estimate an asset's fundamental value, gather and analyze information, they become overconfident about the usefulness of this information. For example, if they uncover favorable information about the asset, their overconfidence about how reliable the information is leads them to push the price of the asset up too high.

While most models of bubble formation are belief-based, there are also some preference-based models. One theory, for example, argues that, after investors experience gains in their holdings of an asset, they become less risk averse because of a "house money" effect: in short, having experienced gains, they are less concerned about future losses because any losses will be cushioned by the prior gains. Their reduced risk aversion leads them to buy the asset even more enthusiastically, thereby pushing its price up even further (Thaler and Johnson, 1990; Barberis, Huang, Santos, 2001).

Another, quite different preference-based model of overvaluation argues that bubbles are particularly likely to occur in stocks related to a new technology (Barberis and Huang, 2008). The reason is that investors view these stocks as lottery-like: should the new technology deliver on its early promise, some of the stocks may experience huge increases in value. Given that many people have a strong preference for lottery-like payoffs -- perhaps because, as Kahneman and Tversky (1979) argue, the brain overweights low probabilities -- they may overvalue these stocks. A theory of this type may be particularly suited to thinking about the high valuations of U.S. technology stocks in the late 1990s.

Of all these models, the one that may be most useful for understanding the recent behavior of the real estate market is the second type of belief-based model: the model that argues that bubbles occur because, perhaps due to the representativeness heuristic, people over-extrapolate the past when making forecasts about the future. On one level, we can apply this idea to home buyers and say that, when forecasting the future growth in house

prices, they over-extrapolated the past growth in these prices. This led them to overpay for their new homes and to take out loans with excessively high loan-to-value ratios.

In order to generate a real estate bubble, however, it is not enough to assume that households were over-extrapolating. Since homes are usually purchased with the help of outside financing, we need to argue that the people involved in the provision of this outside financing were *also* over-extrapolating. In more detail, the story might go as follows. A real estate bubble formed because of an oversupply of credit to home buyers, principally in the form of subprime loans. This, in turn, occurred because, through the process of securitization, the subprime loans could be used to manufacture securities that investors were very enthusiastic about, namely securities with AAA ratings. Crucially, investors were *too* enthusiastic about these securities, because the AAA ratings were often not truly deserved.

It is here, in the rating agencies, that over-extrapolation may have had its greatest impact. The agencies may have given AAA ratings to securities that did not truly deserve them because they extrapolated the past growth in home prices too far into the future, which, in turn, led them to severely underestimate the level of future subprime defaults. The over-extrapolation may have occurred because analysts naively applied the representativeness heuristic; but it may also have occurred because they *wanted* to believe that house prices would keep rising, a belief that the representativeness heuristic made it particularly easy to embrace. I return to this last idea below.<sup>5</sup>

In summary, then, while an account of the recent behavior of the real estate market sounds different, in many of its details, from accounts of other perceived bubbles – the U.S. stock market in the 1920s and the 1990s, the Japanese real estate and stock markets in the late 1980s, not to mention the South Sea bubble of 1720 and the tulip mania of the 1630s – all of these episodes may nonetheless have at least one important driving force in common: a tendency on the part of some market participants to extrapolate past price increases too far into the future.

---

<sup>5</sup> Some countries, of course, experienced dramatic increases in real estate prices even in the absence of much securitized finance. In these cases, the over-extrapolation hypothesis, while essentially the same, may take a simpler form: that banks extrapolated the low past rates of mortgage default too far into the future and, as a result, made many imprudent loans.

## **Cognitive dissonance and risks in the banking system**

The recent plunge in real estate prices was far more devastating to the U.S. economy than the plunge in technology stock prices a few years earlier. One possible reason for this is that, in one case, the banking system was largely unaffected, while in the other, it was severely compromised. In particular, over the course of several years leading up to 2007, banks built up large holdings of subprime loans and of subprime-linked securities. When house prices started falling, the value of these holdings also fell, triggering what Gorton (2010) describes as a crippling run on the banking system. This, in turn, led to a reduction in the supply of credit to the economy. By contrast, when technology stock prices collapsed, banks were barely affected: their exposure to these stocks was relatively small.

How can we understand banks' large holdings of subprime loans and of related securities? On one level, we can try to explain these holdings by saying that they were simply the inventory that inevitably accumulates over the many weeks it can take to complete a securitization deal; that the securities were the "skin in the game" that banks were required to have in order to satisfy investors; or that, since the securities were earning a return that was higher than their funding cost, it was profitable to hold them, at least in the short term.

The problem with these explanations, of course, is that banks' subprime-linked holdings also carried some very significant risks. If house prices were to fall, the value of these holdings would drop precipitously. Given that the banks were highly levered, often with short-term debt, this could have severe consequences. A crucial puzzle therefore remains: Why, in spite of the risk, did banks take on the exposures they did?

There are perhaps three broad answers to this last question. The first, which I label the "bad incentives" view, posits that people on the mortgage desks of banks were *aware* that, through their activities, they were exposing their institutions to significant risk; but that they simply did not care, because their compensation schemes did not force them to face the consequences of the risks they were taking (Acharya et al., 2009). In

many cases, they were compensated largely on the size of the deals they were structuring, and not on the long-term performance of those deals.

The second explanation for banks' large holdings of subprime-linked securities can be labeled the "bad models" view. It says that the people on the mortgage desks of banks were genuinely unaware of the risk embedded in their subprime holdings, and that this was due to faulty reasoning. For example, they, too, may have extrapolated the past growth in real estate prices too far into the future. The models they used to value their positions incorporated this faulty belief, and, as a result, did not reveal any alarming risks.

In appealing to faulty reasoning, the "bad models" view is implicitly assuming that in 2006, say, a *rational* individual with the right incentives would have known that banks' subprime-linked holdings were very risky. A third view, the "bad luck" view, disputes this. According to this view, expressed by Vassalou (2011) in this volume, a rational individual, even one with the right incentives, would *not* have assigned a high probability, *ex-ante*, to the poor subsequent performance of subprime-linked securities. This poor performance was simply bad luck: the realization of a state of the world that a rational observer in 2006 would have deemed very unlikely.

While all three of these views are defensible, I am skeptical of the "bad luck" view. If a rational observer had carefully and exhaustively examined the quality of the subprime loans being extended in the run-up to the crisis, it seems likely that he would have raised at least a few red flags.<sup>6</sup>

The "bad incentives" and "bad models" views, by contrast, seem more plausible. At the same time, even these hypotheses are not quite satisfactory. After all, the mortgage desks of the largest banks were generally staffed by highly intelligent and capable

---

<sup>6</sup> There is some preliminary evidence consistent with the claim that the poor performance of subprime-linked securities was predictable through careful analysis. In an interview reported by McLean and Nocera (2010), a prominent hedge fund manager states that a significant number of fixed-income hedge funds bet against subprime securities. And according to several accounts of the crisis, the two banks with perhaps the most respected risk management organizations – Goldman Sachs and J.P. Morgan – reduced their exposure to subprime loans before the worst of the crisis hit. Much more data on these issues is clearly needed, however.

individuals. How could they allow sloppy reasoning to mislead them about the risks they were taking? In other words, how plausible is the “bad models” view?

It is also unclear how plausible it is that traders *knowingly* exposed their banks and the broader financial system to risk simply because of bad incentives, in other words, simply because they wanted a larger end-of-year bonus. A fundamental idea in social psychology is that people do not only want to make money -- they also want to feel good about themselves, and it is hard to feel good about oneself if one is knowingly doing something that is potentially ruinous to others. So if a trader was aware that his business model posed serious risks to his firm, he might limit the scale of his activities, even if he could earn more money by expanding it further.

If the “bad incentives” and “bad models” views do not tell the whole story, how *can* we understand the large subprime positions that banks built up? Here is an alternative hypothesis. Under this hypothesis, traders on mortgage desks were vaguely aware that their business model might entail serious risks. However, by *manipulating their beliefs*, they deluded themselves into thinking that their business model was *not* risky, but rather, worth pursuing.

One way to put this idea on firmer psychological footing is through the concept of cognitive dissonance. Cognitive dissonance is the discomfort we feel when we take an action that conflicts with our typically positive self-image. Of particular importance is what people often do to remove the feeling of discomfort: they manipulate their beliefs.

For example, smokers often experience cognitive dissonance. A smoker will say to himself: “I am a sensible person – so why am I doing something that is bad for my health?” To reduce the dissonance that he feels, he can stop smoking – but that is hard to do. Instead, he manipulates his beliefs, and convinces himself that smoking is *not*, after all, as risky as some say. He may, for example, remind himself of the 85-year old man who lives down the street and who, despite smoking for much of his life, seems to be doing just fine.

How can we use cognitive dissonance to formalize the story I told above? If a trader on the mortgage desk of a bank begins to sense that the holdings of subprime securities he is building up may pose serious risks to his institution and to the broader

financial system, this will threaten his positive self-image – specifically, his self-image as an upstanding person whose work is valuable to society – and will therefore create uncomfortable dissonance. After all, he does not want to believe that, while enriching himself, he is putting many others at risk. To remove the dissonance, he could resign his position – but that would be financially costly. Instead, he manipulates his beliefs, telling himself that his business model is *not* that risky. For example, he might stop himself from inspecting the quality of the subprime loans he is working with too closely, lest he stumble on some disturbing information.<sup>7</sup>

A similar mechanism may have been at work in the credit rating agencies. On the one hand, an analyst at a rating agency who was being asked, by an issuing bank, to give a subprime-linked product a AAA rating, had a strong financial incentive to do so, even if the rating seemed undeserved: by rating the product AAA, he would avoid losing the business to another rating agency, thereby allowing both him and his firm to earn more money that quarter. On the other hand, the analyst would also want to be able to maintain a positive self-image: to be able to think of himself as a responsible person providing a useful service to society. Giving a AAA rating to a product that did not deserve one would make it hard to maintain a positive self-image and would immediately induce dissonance.

As with the traders on the mortgage desks of banks, the analyst may have reacted to the uncomfortable feeling of dissonance by manipulating his beliefs: by telling himself that the product he was analyzing was perhaps not that risky after all, and therefore deserving of the AAA rating. For example, he may have told himself that, since house prices had been rising for years, they were likely to keep rising, thereby ensuring that subprime defaults would remain low. The representativeness heuristic would have made this argument seem quite plausible: after all, according to that heuristic, people have a natural tendency to believe that past trends *will* continue into the future.

---

<sup>7</sup> I am by no means the first to propose that belief manipulation played a role in the crisis. Many accounts of recent events, both academic and non-academic, have also suggested it. In particular, see Benabou (2009) for a formalization of ideas related to those I present here. For a discussion of the research in psychology on cognitive dissonance, see Kunda (1999) and Chapter 6 of Aronson, Wilson, and Akert (2005).

It is worth noting that, for at least two reasons, subprime securitization may have lent itself particularly well to belief manipulation. The first reason is that subprime-linked products were often complex. Given their intricacies, it would have taken considerable effort to disprove the claim that they were relatively safe. This may have made it easier for people to delude themselves about their risks.

The second reason why it may have been easy for people to hold distorted beliefs about the risks of subprime securities is because there was a plausible-sounding argument that appeared to justify these beliefs. The argument was simply that, since house prices had been rising for many years, they were likely to keep rising -- and if they did keep rising, then subprime defaults would be low, as would the risks of subprime-linked securities. As I noted above, the representativeness heuristic made this a particularly seductive argument.

The belief manipulation hypothesis can be thought of as an alternative to the “bad incentives”, “bad models”, and “bad luck” views. But it can also be thought of as a foundation for the bad models view. In the belief manipulation view, as in the bad models view, mortgage traders are unaware of the risks they are taking. The belief manipulation view tries to explain *why* they are unaware. In short, they are unaware because they choose to be.<sup>8</sup>

### **Psychological amplification mechanisms**

A striking feature of the crisis was that, during the crisis period, many kinds of risky assets experienced dramatic price declines – price declines that were surprisingly large, given the relatively small delinquencies among subprime loans.

To explain these large price drops, researchers have focused on *institutional* amplification mechanisms. For example, if a bank’s holdings of subprime-related securities decline in value, then, in order to deleverage or to meet more stringent margin requirements, the bank will have to sell some of its risky asset holdings. This will push

---

<sup>8</sup> See Gennaioli, Shleifer, and Vishny (2011) for another possible foundation for the bad models view, namely that market participants neglect unlikely bad scenarios.

down the value of *other* banks' holdings of risky assets, forcing them into sales of their own, thereby pushing the prices of risky assets down even further, and so on.

These loss spirals and margin spirals, described in detail by Shleifer and Vishny (1997), Gromb and Vayanos (2002), and Brunnermeier (2009), among others, were probably important in transforming relatively small subprime losses into much larger price declines on many kinds of risky assets. However, *psychological* amplification mechanisms – specifically, mechanisms related to loss aversion and ambiguity aversion, two concepts that have been extensively studied in behavioral finance – may also have played a role. In short, the idea is that, after suffering losses in their risky asset holdings, both institutional and individual investors experienced increases in loss aversion and ambiguity aversion. This led them to reduce their holdings of risky assets, thereby pushing the prices of these assets down even further.

The idea that an increase in ambiguity aversion was central to the crisis has already been put forward (see, for example, Caballero and Krishnamurthy, 2008, Easley and O'Hara, 2010, and Krishnamurthy, 2010). My goal here is to emphasize that this idea has strong psychological foundations. Many economists are familiar with ambiguity aversion – the notion that people are averse to situations where they do not feel able to assign probabilities to future outcomes – and are aware of the basic evidence for it, the Ellsberg paradox. However, economists are typically much less aware of a literature in psychology – one that is potentially very relevant to finance -- on how ambiguity aversion can *change* over time. Two particularly insightful papers in this literature are those of Heath and Tversky (1991) and Fox and Tversky (1995).

Heath and Tversky (1991) present a theory of ambiguity aversion which they label the “competence hypothesis.” The idea is that an individual can be either ambiguity averse *or* ambiguity seeking, depending on how *competent* he feels at analyzing the situation at hand. Here, “competence” refers to how much the person feels he knows about a situation relative to what could be known. According to the competence hypothesis, if the individual does not feel competent at analyzing some situation, he will be ambiguity averse. Conversely, if he does feel competent at analyzing the situation, he will be ambiguity seeking.

Through a series of ingenious studies, Heath and Tversky (1991) and Fox and Tversky (1995) provide evidence for the competence hypothesis. One of their most striking findings is that they can alter subjects' degree of ambiguity aversion by manipulating their feelings of competence. Specifically, they are able to increase subjects' aversion to an ambiguous situation by reminding them of *another* situation that is easier to analyze; or by telling subjects that another, seemingly more able group of people, is also analyzing the same situation. For example, in one experiment, they tell their undergraduate student subjects at San Jose State University that the situation they are analyzing is also being studied by a group of graduate students at Stanford University. This news significantly increases the San Jose State students' ambiguity aversion.

These results, while fascinating in their own right, may also be useful to financial economists because they suggest a way of understanding the large declines in risky asset prices during the crisis. In the language of Heath and Tversky (1991) and Fox and Tversky (1995), once investors suffered some initial losses in their holdings of risky assets – losses that coincided with surprising and confusing developments in the market for subprime-linked securities – they felt less *competent* at analyzing these assets. This made them more ambiguity averse, leading them to reduce their holdings of risky assets, thereby further lowering the prices of these assets.<sup>9</sup>

Loss aversion -- Kahneman and Tversky's (1979) observation that people are much more sensitive to losses than to gains of the same magnitude -- is perhaps even more familiar to economists than ambiguity aversion. Of particular relevance here, however, is some evidence that economists are less aware of, namely evidence that the degree of loss aversion can change over time depending on experienced gains and losses. Specifically, in a series of experiments, Thaler and Johnson (1990) show that subjects who experience a loss subsequently become *more* loss averse, refusing to take gambles that, in the absence of the prior loss, they would take. It is still not fully understood what is driving this effect, but a natural interpretation is that, after suffering through one

---

<sup>9</sup> More generally, the findings of Heath and Tversky (1991) and Fox and Tversky (1995) may be useful for understanding the high empirical return volatility of many risky asset classes. If an asset class performs poorly, investors may feel less competent at analyzing it, increasing their ambiguity aversion and triggering selling and further price declines. Conversely, if an asset class performs well, investors may feel *more* competent at analyzing it, decreasing their ambiguity aversion and leading to purchases and further price increases.

painful loss, people cannot face the idea of going through another loss. In short, their loss aversion increases.

It is clear how changes in loss aversion could have aggravated the collapse in risky asset prices during the crisis. The initial price declines forced many investors to endure painful losses. As suggested by Thaler and Johnson (1990), these losses may have made investors more loss averse, leading them to reduce their risky asset holdings and thereby causing further price declines.<sup>10</sup>

## **Conclusion**

This book is, in part, about financial innovation. In general, the field of behavioral finance takes a favorable view of financial innovation. After all, a major theme of behavioral finance research is that people often make suboptimal financial decisions. If this is the case, then financial innovations can play a useful role in helping people to make better decisions. Indeed, over the past few years, a new branch of behavioral finance has emerged – a branch sometimes known as “prescriptive behavioral finance” – whose goal is precisely to design innovations that can help people achieve better financial outcomes (Thaler and Sunstein, 2008).

While financial innovations can be useful in preventing psychological factors from leading people astray, the discussion above suggests that the same psychological factors can make certain innovations dangerous. This may be particularly true for innovations that are *complex*. If an innovation is complex, it is easier for people supplying the innovation to convince themselves that it is *not* flawed, even if in fact, it is; this may then lead them to market the innovation too aggressively. Moreover, the failure of a complex financial innovation may have large amplifying effects because it may cause investors to feel less competent at analyzing risky assets in general, and hence to drive the prices of these assets down.

---

<sup>10</sup> The argument in this section is related to that of Cochrane (2009). He argues that the initial declines in risky asset prices brought investors’ consumption closer to their habit level of consumption. This increased their risk aversion, leading to further declines in risky asset prices.

Most of the ideas for financial reform that have been proposed over the past few years are aimed at the *institutional* failures that contributed to the crisis. While it is too early to be sure, it is very possible that psychological factors were also central to the crisis. As such, it may be important to think about reforms that can address both the institutional *and* the psychological failures. In short, the financial crisis presents finance researchers, and perhaps behavioral finance researchers in particular, with a challenge: to design a financial system that can mute the impact of irrational thinking, and prevent it from adversely affecting the real economy in the way that it may recently have done. This is a difficult challenge – but it may be one of the most important facing us today.

## References

- Acharya, V., Cooley, T., Richardson, M., and I. Walter (2009), "Manufacturing Tail Risk: A Perspective on the Financial Crisis of 2007-2009," *Foundations and Trends in Finance* 4, 247-325.
- Akerlof, G. and R. Shiller (2009), *Animal Spirits*, Princeton University Press.
- Aronson, E., Wilson, T., and R. Akert (2005), *Social Psychology*, Prentice Hall, Fifth Edition.
- Barberis, N. and M. Huang (2008), "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices," *American Economic Review* 98, 2066-2100.
- Barberis, N., Huang, M., and T. Santos (2001), "Prospect Theory and Asset Prices," *Quarterly Journal of Economics* 116, 1-53.
- Barberis, N. and A. Shleifer (2003), "Style Investing," *Journal of Financial Economics* 68, 161-199.
- Barberis, N., Shleifer, A., and R. Vishny (1998), "A Model of Investor Sentiment," *Journal of Financial Economics* 49, 307-343.
- Benabou, R. (2009), "Groupthink: Collective Delusions in Organizations and Markets," Working paper, Princeton University.
- Brunnermeier, M. (2009), "Deciphering the Liquidity and Credit Crunch 2007-2008," *Journal of Economic Perspectives* 23, 77-100.
- Caballero, R. and A. Krishnamurthy (2008), "Musical Chairs: A Comment on the Credit Crisis," *Banque de France Financial Stability Review* 11, 1-3.
- Cochrane, J. (2009), "Asset Pricing After the Crash," Manuscript, University of Chicago.
- Daniel, K., Hirshleifer, D., and A. Subrahmanyam (1998), "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance* 53, 1839-1885.
- Easley, D. and M. O'Hara (2010), "Liquidity and Valuation in an Uncertain World," *Journal of Financial Economics* 97, 1-11.
- Fox, C. and A. Tversky (1995), "Ambiguity Aversion and Comparative Ignorance," *Quarterly Journal of Economics* 110, 585-603.
- Gennaioli, N., Shleifer, A., and R. Vishny (2011), "Neglected Risks, Financial Innovation, and Financial Fragility," *Journal of Financial Economics*, forthcoming.
- Gorton, G. (2010), *Slapped by the Invisible Hand: The Panic of 2007*, Oxford University Press.
- Greenwood, R. and S. Hanson (2010), "Issuer Quality and Corporate Bond Returns," Working paper, Harvard University.
- Gromb, D. and D. Vayanos (2002), "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs," *Journal of Financial Economics* 66, 361-407.
- Harrison, M. and D. Kreps (1978), "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics* 92, 323-336.

- Heath, C. and A. Tversky (1991), "Preference and belief: Ambiguity and Competence in Choice under Uncertainty," *Journal of Risk and Uncertainty* 4, 5-28.
- Hong, H. and J. Stein (2007), "Disagreement and the Stock Market," *Journal of Economic Perspectives* 21, 109-128.
- Kahneman, D. and A. Tversky (1974), "Judgment under Uncertainty: Heuristics and Biases," *Science* 185, 1124-1131.
- Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decision under Risk," *Econometrica* 47, 263-291.
- Krishnamurthy, A. (2010), "Amplification Mechanisms in Liquidity Crises," *American Economic Journal: Macroeconomics* 2, 1-30.
- Kunda, Z. (1999), *Social Cognition: Making Sense of People*, MIT Press.
- Lakonishok, J., Shleifer, A., and R. Vishny (1994), "Contrarian Investment, Extrapolation, and Risk," *Journal of Finance* 49, 1541-1578.
- McLean, B. and J. Nocera (2010), *All the Devils are Here: The Hidden History of the Financial Crisis*, Penguin Group.
- Miller, E. (1977), "Risk, Uncertainty, and Divergence of Opinion," *Journal of Finance* 32, 1151-1168.
- Scheinkman, J. and W. Xiong (2003), "Overconfidence and Speculative Bubbles," *Journal of Political Economy* 111, 1183-1219.
- Shefrin, H. (2009), "How Psychological Pitfalls Generated the Global Financial Crisis," Working paper, Santa Clara University.
- Shleifer, A. and R. Vishny (1997), "The Limits of Arbitrage," *Journal of Finance* 52, 35-55.
- Thaler, R. and E. Johnson (1990), "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice," *Management Science* 36, 643-660.
- Thaler, R. and C. Sunstein (2008), *Nudge: Improving Decisions about Health, Wealth, and Happiness*, Yale University Press.
- Vassalou, M. (2011), "Market Efficiency, Rational Expectations, and Financial Innovation," forthcoming in *Financial Innovation and Crisis*, M. Haliassos, ed., MIT Press.