

7529121-behavioral-finance-at-40

SUMMARY: This article discusses the progress and applications of behavioral finance, which has evolved into a vibrant area of financial economics over the past four decades. It focuses on three main areas: beliefs about future earnings and returns, and prospect theory, and their applications to understanding asset prices and investor behavior. The author reviews recent progress, open questions, and new directions in the field, including cognitive foundations, non-traditional decision processes, and artificial intelligence.

METHOD: The article synthesizes recent research on behavioral finance, covering studies on beliefs about future earnings and returns, prospect theory, and their applications to asset prices and investor behavior. It also discusses new directions in the field, such as cognitive foundations, non-traditional decision processes, and artificial intelligence.

KEY FINDINGS: - Behavioral finance has expanded into a significant area of study, with three main applications: asset prices, household finance, and corporate finance.

- Recent research has provided support for behavioral approaches to asset prices, particularly through data on people's expectations about future returns and earnings.
- The field has seen significant progress in household finance, with new datasets and research on structural approaches, household liabilities, and improving financial decision-making.
- Behavioral corporate finance research has focused on two frameworks: one for security issuance and capital structure, and another for acquisition activity, with managers often being irrational in their decision-making.

IMPLICATION FOR TRADING: Behavioral finance has practical implications for traders and investors, as it helps explain why financial markets and investor behavior often deviate from rational expectations. Understanding the psychological biases and decision-making processes can aid in identifying mispriced assets and predicting market movements. Additionally, it can inform investment strategies and risk management practices, leading to more informed and effective trading decisions.

Behavioral Finance at 40

Nicholas Barberis

September 2026*

Abstract

Behavioral finance seeks to improve our understanding of markets, investors, and firms by adopting a psychologically realistic description of how people think. While behavioral finance ideas have been part of human discourse for decades if not centuries, the modern research effort in the field began in the 1980s. Forty years later, behavioral finance is a vibrant area of financial economics. In this article, I explain why by reviewing some of its recent progress. I focus on three topics – beliefs about future earnings, beliefs about future returns, and prospect theory – and on their applications to understanding asset prices and investor behavior. I close by commenting on some new directions in the field – cognitive foundations, non-traditional decision processes, and artificial intelligence.

Keywords: behavioral finance, mispricing, belief biases, prospect theory, cognitive foundations

JEL Codes: G11, G12, G14, G40, G50

*The author's affiliation is Yale School of Management, Yale University, New Haven, Connecticut, USA; e-mail: nick.barberis@yale.edu. This article draws on, but also differs in content from, the author's Annual Lecture at the 2025 meeting of the American Finance Association, available on the Association's YouTube channel and titled "Behavioral Finance at 40: Progress, Open Questions, and New Directions."

1 Introduction

The field of behavioral finance argues that we can improve our understanding of financial markets, investor behavior, and corporate finance by adopting a psychologically realistic description of how people think – one that, in particular, allows for thinking that is not fully rational. Behavioral finance ideas have been part of human discourse for decades if not centuries, but scientific research on these ideas began in the 1980s with studies such as Shiller (1981), De Bondt & Thaler (1985), and De Long et al. (1990).

Today, 40 years later, behavioral finance is more vibrant than at any time in its modern history. In this essay, I explain why by reviewing recent progress, listing some open questions, and commenting on new directions. The article is intended to be helpful both to specialists and to people with little knowledge of the field. Behavioral finance is now a large area of study – three long chapters were needed to survey it for the 2018 *Handbook of Behavioral Economics* – so, in this short review, I will have to focus on a few specific topics. To provide context, I will start with a brief overview.

Most behavioral finance papers can be placed somewhere in the matrix in Figure 1. Each row of the matrix corresponds to an aspect of human psychology and each column to an empirical fact. For example, Benartzi & Thaler (1995), which uses loss aversion to explain the equity premium puzzle, would be in the “loss aversion” row (under “preferences”) and in the “equity premium” column (under “asset prices”). The matrix shows that, when we think about behavioral finance, or when we teach it, we can use either a row-based approach – we can take each psychological concept in turn and consider all of its possible finance applications – or a column-based approach: for each empirical fact in turn, we can review all of its possible psychological explanations.

The matrix also indicates that behavioral finance has three main areas of application – asset prices, household finance, and corporate finance – and involves three aspects of psychology: beliefs, preferences, and the decision-making process. It will be helpful to make brief remarks about each of these.

Application: Asset prices. This is the first topic that was studied in behavioral finance, beginning in the 1980s. It was a natural topic to take on because there were well-established facts about stock prices that were challenging to understand from a rational-investor perspective. However, this focus plunged behavioral finance into controversy from the start: at the time, many economists believed that, due to the force of arbitrage, irrational investors were unlikely to have much impact on asset prices. The tense debates on this issue quieted down only after the work on “limits to arbitrage” in the 1990s, which argued that it is difficult for rational investors to fully correct mispricing (Shleifer & Vishny 1997). Today, four decades later, asset pricing is once again a vibrant area of research in behavioral finance, primarily because data on people’s expectations about future returns and earnings provide support for behavioral approaches to asset prices.

Application: Household finance. This field studies the financial decisions of households. Since these decisions are likely to be suboptimal, it is not surprising that behavioral ideas have thoroughly permeated household finance, and have done so in an uncontroversial way. An early, and very behavioral, line of work in this area was Brad Barber and Terrance Odean’s research on the stock trading of individual investors (Odean 1998, Barber & Odean 2000). In the 2000s and 2010s, with the arrival of new datasets on household behavior, the field developed rapidly. Recent topics of interest include structural approaches (Choukhmane 2025); household liabilities (Lee & Maxted 2026); and improving people’s financial decision-making (Thaler & Sunstein 2009, Choi 2022).

Application: Corporate finance. Research on behavioral corporate finance began in earnest in the 2000s and focused on two frameworks. In the first, which is often used to think about security issuance and capital structure, managers respond rationally to the excessive exuberance or pessimism of outside investors (Baker & Wurgler 2000, 2002). In the second framework, often used to think about acquisition activity, managers themselves are irrational – for example, overconfident (Malmendier & Tate 2005, 2008). There were fewer breakthroughs in behavioral corporate finance in the 2010s, but in the 2020s, the field began exploring new directions, such as the social preferences of managers (Guenzel, Hamilton & Malmendier 2025) and the health effects of CEO stress (Borgschulte et al. 2025).

Psychology: Beliefs. The first wave of behavioral finance research in the 1980s and 1990s focused on investor beliefs. There was less discussion of beliefs in the 2000s, but starting in the 2010s, data on expectations about earnings and returns spurred a new wave of research on the topic, one that continues through to the present day.

Psychology: Preferences. Behavioral finance research on investor preferences began in the 1980s. From the start, it has been heavily influenced by Kahneman & Tversky’s (1979) prospect theory, the most prominent framework in all of behavioral economics. Research on preferences continues in behavioral finance today although, over the past decade, there has been a greater focus on investor beliefs.

Psychology: The decision-making process. Most papers in financial economics assume that people make decisions by first forming beliefs about future outcomes and then choosing among different possible actions based on their preferences. However, decades of research in psychology suggest that decisions may sometimes be made by way of a different process, one that does not involve explicit formation of beliefs. I discuss this in more detail in Section 5.2.

In this essay, I will focus on behavioral finance research on asset prices. This has been a particularly vibrant area of study over the past 10 years. Its findings are striking and challenge models with rational investors – so much so that we may be on the cusp of a significant shift in academic thinking about asset prices, away from the rational-investor models which

have dominated for the past 40 years, and toward models with a more behavioral flavor.¹

The space allotted for this essay does not allow for a full review of behavioral asset pricing. Even within this area, I have to be selective. To understand my choices, it is helpful to start with a list of prominent behavioral frameworks on asset prices:

- limits to arbitrage
- (beliefs) disagreement
- (beliefs) biases in forecasts of future earnings (**)
- (beliefs) biases in forecasts of future returns (**)
- (preferences) gain-loss utility and prospect theory (**)

The work on limits to arbitrage and on how disagreement, when coupled with a short-sale constraint, can lead to overpricing, are now both part of the canon of financial economics; however, there have not been any major breakthroughs on these topics over the past 10 years. In this essay, I will therefore focus on the three other frameworks, those marked with asterisks. In Section 2, I discuss biases in beliefs about firms’ future earnings. In Section 3, I turn to biases in beliefs about assets’ future returns. In Section 4, I review frameworks in which investors derive utility from gains and losses and that incorporate elements of prospect theory. For each of these topics, I list several open questions. In Section 5, I comment on three new directions – cognitive foundations, alternative decision-making processes, and artificial intelligence. Section 6 concludes.

2 Biased beliefs about earnings

Over the past decade, we have learned a lot about market participants’ beliefs about firms’ future earnings. One way to summarize these findings is through the lens of “Coibion-Gorodnichenko” (CG) regressions. Suppose that the realized value of an economic variable x at time t is denoted x_t , and let $F_s x_t$ be an individual’s forecast at time $s < t$ of the value of x at time t . In a CG regression, we regress forecast errors on forecast revisions,

$$FE_{t+1} = \alpha + \beta FR_{t-1 \rightarrow t} + \varepsilon_{t+1},$$

¹Readers who would like to learn more about behavioral household finance can consult Beshears et al. (2018); for behavioral corporate finance, see Baker, Ruback & Wurgler (2007) and Malmendier (2018). Barberis (2018) is an earlier review of behavioral perspectives on asset prices, while Barberis & Thaler (2003) is an even earlier survey that covers all three areas of application.

where $FE_{t+1} = x_{t+1} - F_t x_{t+1}$ and $FR_{t-1 \rightarrow t} = F_t x_{t+1} - F_{t-1} x_{t+1}$. For a rational individual, the forecast error FE_{t+1} should not be predictable by any variable in his information set at time t ; in particular, it should not be predictable by the forecast revision, which means that β should equal zero. If we instead obtain $\beta > 0$, this is interpreted as underreaction, and $\beta < 0$ as overreaction. To see the logic, take the case where $\alpha = 0$. In this case, if $\beta > 0$, then, even after the individual revises up his forecast, the realized value of x systematically exceeds the revised forecast. The upward revision is therefore insufficient, on average.

Two important findings from the past decade are that, when we run the above regression for Wall Street analysts' forecasts of *near-term* earnings, we obtain $\beta > 0$ (Bouchaud et al. 2019); but when we run it for forecasts of *long-term* earnings growth, we obtain $\beta < 0$ (Bordalo et al. 2019). Forecasts of near-term earnings therefore appear to underreact to information, while forecasts of long-term earnings growth appear to overreact. These results hold whether we run the regression using consensus or individual analyst forecasts. Moreover, if we run the regression using analysts' forecasts of *economy-wide* long-term growth, constructed by aggregating firm-level forecasts, we again obtain $\beta < 0$ (Bordalo et al. 2024).

These findings have attracted attention because they are helpful for understanding important facts about asset prices. The underreaction in forecasts of near-term earnings can explain why, in the data, a firm's recent earnings momentum, recent return, and profitability predict the firm's subsequent stock return in the cross-section with a positive sign. A positive change in earnings, a good return, and high profitability are all positive signals about a firm's near-term future earnings; but if investors underreact to them, the firm's stock will initially be undervalued and will earn high returns later as the undervaluation is corrected. Bouchaud et al. (2019) provide empirical support for this mechanism: they show that the above three patterns of return predictability are particularly pronounced for stocks covered by analysts who underreact more when making forecasts of near-term earnings.

The overreaction in long-term growth forecasts can also account for important facts about asset prices, such as the excess price volatility of the aggregate stock market, the negative relationship between the market's price-to-earnings (P/E) ratio and subsequent market returns, and the value premium in the cross-section of stock returns – the fact that low P/E “value” stocks have higher average returns than high P/E “growth” stocks. If, following good news about a firm, investors become too optimistic about its long-term earnings growth, they will push its P/E ratio up too high; subsequent earnings growth will disappoint on average, bringing low returns. This can explain the value premium, and a similar mechanism at the level of the aggregate market can explain the excess volatility and time-series predictability of market returns.

Bordalo et al. (2024) provide evidence for this mechanism in the case of the aggregate stock market. A key finding is that analysts' consensus forecast of long-term economy-wide earnings growth (LTG) predicts the market's subsequent return with a negative sign. This is consistent with the LTG forecast overreacting to news and hence being too extreme – and

with investors using this overly extreme belief to price the stock market. Delao & Myers (2021) also link aggregate stock prices to beliefs, but argue for a larger role for forecasts of near-term earnings. Bordalo et al. (2025) present results for the cross-section of stock returns. They show that a large part of the high return to value stocks can be traced to realized earnings systematically exceeding analysts’ overly pessimistic long-term growth forecasts, and to analysts’ gradually revising up these pessimistic forecasts.

For the past 30 years, academic thinking about asset prices has been dominated by models of rational investors that focus on discount rates. To explain aggregate market facts, these models assume that discount rates vary over time due to changing risk or risk aversion. To explain facts about the cross-section of returns, they posit that certain types of stocks have higher average returns because they are riskier and therefore require a higher discount rate. The research discussed in this section offers a different perspective: that asset prices are driven not by discount rates – the denominator in the present-value formula – but by the numerator, biased expectations of cash flows. This perspective goes back decades – see Barberis, Shleifer & Vishny (1998) – but the expectations data provide new support for it.

2.1 Sources of under- and over-reaction

Since the biases in analyst beliefs – the underreaction in near-term earnings forecasts and overreaction in long-term growth forecasts – show promise in explaining core puzzles about asset prices, there is an ongoing effort to understand their source. One approach is to break the problem into two parts: to first identify the source of the underreaction, then the source of the overreaction, and to then combine the two mechanisms in a single framework.

Three main mechanisms have been proposed for underreaction:

Conservatism in belief updating. Laboratory studies of belief updating, first conducted decades ago and recently reviewed by Benjamin (2019), find that people often underreact when updating their beliefs in response to new information: if the new information would lead a rational person to update their assessment of the probability of an event from 0.5 to 0.9, the typical individual may update only from 0.5 to 0.7, say. One approach to understanding this is based on “cognitive uncertainty,” the subjective feeling that it is hard to figure out the answer to the updating problem. Enke & Graeber (2023) argue, with evidence, that due to this feeling of uncertainty, the individual shrinks her answer toward a prior belief, and hence underreacts.

Inattention. If an individual does not pay full attention to a piece of information, it is more likely that she will underreact to it. Many studies document inattention to news in financial markets. These include DellaVigna & Pollet (2007, 2009), Hirshleifer, Lim & Teoh (2009), and Huang (2015); see Section 9.1 of Barberis (2018) for a more comprehensive list. Bouchaud et al. (2019) provide evidence that inattention plays a role specifically in the

underreaction in near-term earnings forecasts.

Herding. If a financial analyst prefers to state a belief that is not too different from the median belief of other analysts, this may lead her to underreact to her own private information about a firm. Banerjee, Gondhi & Zhao (2026) and Bastianello (2026a) argue that such herding among analysts is one source of the underreaction in their earnings forecasts.

To explain the overreaction in long-term forecasts, a number of other mechanisms have been explored:

Base-rate neglect / representativeness. These closely-related mechanisms posit that, when updating their beliefs, people put too much weight on the likelihood term in Bayes' rule and too little weight on the base rate or prior term. In a financial context, following good news about a firm, investors over-estimate the firm's long-run earnings growth because they neglect the base rate – the fact that it is rare for a firm to have a high long-term growth rate. Representativeness, in particular, is often modeled through the “diagnostic expectations” approach of Bordalo, Gennaioli & Shleifer (2018).

An incorrect belief in a law of small numbers. Tversky and Kahneman (1971) argue that people incorrectly think that even a small sample of data will reflect the properties of the process that generated the data. Rabin (2002) and Rabin & Vayanos (2010) analyze a key implication of this, namely that people will overinfer from small samples. For example, following good news about a firm, they will become too optimistic about its future prospects.

Overconfidence. If investors are too confident in their assessment of a firm's future growth, they will be too optimistic following good news and too pessimistic following bad news (Daniel, Hirshleifer & Subrahmanyam 1998).

Memory. One memory-based mechanism of overreaction relies on the associative property of memory – the fact that, after we are prompted with a cue, we tend to recall things that share the features of the cue. For example, after hearing good news about a firm, we are more likely to recall other pieces of good news about it; this leads good news to be over-represented in our thinking, leading to overreaction (Enke, Schwerter & Zimmermann 2024).

One approach, then, is to argue for one mechanism in the underreaction category and for another in the overreaction category; to combine them in a single framework; and to develop testable implications. For example, an early study of under- and over-reaction by Barberis, Shleifer & Vishny (1998) combines conservatism and representativeness and shows that the resulting framework can account for several prominent anomalies in the cross-section of stock returns. More recently, Ba, Bohren & Imas (2026) study the interaction of representativeness and cognitive uncertainty, while Banerjee, Gondhi & Zhao (2026) argue for a combination of herding and overconfidence.

Rather than combining two different mechanisms, we can instead look for a single mechanism that generates both the underreaction and the overreaction. Augenblick, Lazarus &

Thaler (2025) propose and study a promising candidate: that people overreact to weak information and underreact to strong information. Their data nest the evidence of conservatism in belief updating mentioned above: conservatism is primarily observed when people are presented with strong evidence; when faced with weak evidence, they more typically overreact. Such behavior can be explained using the cognitive uncertainty framework. Suppose that, when an investor sees a piece of news about a firm, she is unsure how informative the news is. She therefore shrinks her estimate of its informativeness toward a moderate prior: she treats a weakly informative piece of news as moderately informative, and hence overreacts to it, but also treats a strongly informative piece of news as, again, moderately informative, and hence underreacts to it.

This approach can explain both the underreaction in near-term earnings forecasts and the overreaction in long-term growth forecasts. Precisely because near-term earnings are “near,” we have strong signals about them; the above framework then says that we will underreact to these signals. And precisely because long-term earnings growth involves the long term, the signals we have about it are weaker; we will therefore overreact to these signals. Direct tests of this idea in the context of stock prices would be welcome.

2.2 Open questions

The work described above leaves many open questions, most obviously the first-order question of why there is underreaction in near-term forecasts but overreaction in long-term forecasts. Below, I discuss two more.

Which types of news do investors misreact to, and why? One way to make progress on the under- vs. over-reaction question is to see if investors react differently to different types of news.

Kwon & Tang (2026) take a database of corporate news and classify the news into 24 categories – for example, news about a CEO change, a lawsuit, or an M&A transaction. They find that the stock market underreacts to about half the news types and overreacts to the rest. They also show that there is more overreaction for types of news that occasionally lead to extreme and hence memorable outcomes, and that thereby mislead investors into reacting too much. For example, if a technology firm announces a product launch, many investors will instinctively recall the original iPhone launch and the very high subsequent returns on Apple stock. This leads them to move the firm’s stock price a lot – but this is an overreaction because they are neglecting the many product launches which had a much smaller impact.

Another dimension on which news varies is its degree of quantitateness. Ke (2026) uses LLMs to summarize the factual information in over a million equity analyst reports and to categorize it into 29 different topics. He finds that analysts tend to underreact to

quantitative information when forecasting future earnings and to overreact to qualitative information. Since analysts pay relatively more attention to quantitative information when forecasting near-term earnings and relatively more attention to qualitative information when forecasting long-term earnings, this helps explain the underreaction in near-term forecasts and overreaction in long-term forecasts. Using the same database as Kwon and Tang (2026), Hong (2026) provides complementary evidence from stock prices: he finds more stock market underreaction to news that is more quantitative. Didisheim et al. (2026) obtain a similar result using a different methodology.

The results in Ke (2026), Hong (2026), and Didisheim et al. (2026) raise the question of why there is more overreaction to qualitative information. One possibility is that it is easier to be overconfident in one’s interpretation of qualitative information. Another is that qualitative, story-like information is more memorable and therefore more likely to have a persistent impact on beliefs (Graeber, Roth, & Zimmermann 2024).

The findings on qualitative and quantitative news dovetail with earlier analyses of tangible and intangible information, where tangible information refers to accounting measures of past performance. Daniel and Titman (2006) argue that there is little evidence of overreaction to tangible information in the cross-section of stock returns – past growth in accounting fundamentals does not predict future stock returns with a negative sign – and that overreaction is therefore more likely a response to intangible information.²

Do beliefs move prices or do prices move beliefs? Many behavioral finance papers on asset prices argue that, due to psychological mechanisms, analysts and investors form biased beliefs about future earnings, and these beliefs then affect prices. An alternative view is that price movements are driven by something else – by changing preferences, say – and that the beliefs analysts state about future earnings are not necessarily their true beliefs but rather the beliefs that fit current prices. For example, if the stock market attains a high P/E ratio because of a decline in investor risk aversion, analysts may, at this moment, state a belief that long-term earnings growth will be high, as this fits the high P/E ratio. Under this view, prices move beliefs rather than the other way around. Chaudhry (2025) presents evidence consistent with this view, while Bordalo et al. (2024, 2025) argue against it.

3 Biased beliefs about returns

When forecasting the future return of the stock market, many investors extrapolate past returns: if the stock market performed well in the past year, they think it will keep performing

²This is not to say that there is no evidence of overreaction to tangible information. Nagel & Xu (2022) find that a weighted average of past aggregate fundamentals predicts the stock market’s future return with a negative sign, suggesting overreaction to these tangible fundamentals. Bordalo et al. (2024) find that the component of analysts’ revisions in their long-term growth forecasts that is driven by earnings surprises predicts the market’s return with a negative sign, suggesting overreaction to these tangible earnings surprises.

well over the next year; and if it performed poorly over the past year, they think it will keep performing poorly. Such beliefs are not rational: stock market returns are not as positively autocorrelated as these investors think. As such, these investors are *over*-extrapolating past returns.

An influential study on stock market return expectations is by Greenwood & Shleifer (2014), who analyze data from six surveys of investors. Four of these surveys are of individual investors; one is of CFOs; and one summarizes the views of financial market newsletters. In all six surveys, the typical respondent extrapolates past returns when forming beliefs about future returns.

Motivated by these results, researchers have worked to understand how asset prices will behave when some investors forecast future returns by extrapolating past returns. This line of research is promising: it shows that models of return extrapolation can explain a number of important facts about asset prices.

To see the implications of return extrapolation, take a financial market with two assets – a risk-free asset and a risky asset – and in which some investors, “extrapolators,” form beliefs according to

$$E_t(P_{t+1} - P_t) = X_t \equiv (1 - \theta) \sum_{k=1}^{t-1} \theta^{k-1} (P_{t-k} - P_{t-k-1}) + \theta^{t-1} X_1, \quad (1)$$

where P_t is the price of the risky asset at time t and X_1 measures initial extrapolator enthusiasm at time 1. In words, an extrapolator’s expectation about the future price change is a weighted average of past price changes, with more weight on more recent price changes. The parameter θ lies between 0 and 1; the closer θ is to 0, the more relative weight these investors put on price changes in the recent past. If, at time t , an extrapolator makes his asset allocation decision using an exponential utility function defined over wealth at time $t + 1$, then, under some additional simplifying assumptions, his share demand for the risky asset at time t is given by

$$N_t = \frac{X_t}{\gamma \sigma^2}, \quad (2)$$

where γ is risk aversion and σ^2 is a measure of risk. We can close the model with a second group of more rational investors – for example, “fundamental traders” whose demand depends on the difference between the price and a sensible estimate of future cash flows; or, fully rational traders who know how extrapolators form beliefs.

Models of this type show that the extrapolators amplify movement in asset prices: following good news, the price of the risky asset rises too high and then corrects downward; following bad news, it falls too low and then corrects upward. In more detail: if, at time t , there is good news about the risky asset’s future cash flows that pushes its price up, then, at time $t + 1$, the extrapolators, seeing the positive price change at time t , become more

optimistic about future returns and increase their demand, pushing the price further up; a similar effect occurs at time $t + 2$. However, the price increases created by extrapolator demand at time $t + 1$ and time $t + 2$ are typically smaller than the initial price increase at time t . At time $t + 3$, then, extrapolators become less optimistic about future returns, reduce their demand, and cause a drop in price. The negative price change at time $t + 3$ leads to an even bigger drop in extrapolator enthusiasm, and hence prices, at time $t + 4$.³

This cycle of amplification and correction allows models of return extrapolation to account for a number of puzzling facts about asset prices. If we take the risky asset to be the aggregate stock market, then a model with return extrapolation can help explain the excess volatility and time-series predictability of stock market returns (Barberis et al. 2015, Adam, Marcet & Beutel 2017, Jin & Sui 2022). If we take the risky asset to be an individual stock, the model can account for momentum, long-term reversals, and the value premium in the cross-section of stock returns (Hong and Stein 1999). And regardless of the identity of the risky asset, if the initial news about its future cash flows is sufficiently positive, and if there are enough extrapolators in the market, the asset’s price will first rise far above the rational level and then sharply decline – a price pattern that resembles a bubble (Kindleberger 1978, Barberis et al. 2018). Return extrapolation thereby offers a framework for thinking about bubbles that complements the other prominent behavioral model of bubbles based on investor disagreement (Harrison & Kreps 1978, Scheinkman & Xiong 2003).

When we take the risky asset in Equations 1 and 2 to be an individual stock, we are assuming that a large group of investors – most likely, retail investors – has extrapolative beliefs about the returns on individual stocks, just as they do for aggregate market returns. There is, as yet, little direct evidence of this: we do not have a large-scale dataset on the beliefs of retail investors about the future returns of individual stocks. The construction of such a dataset would be a significant contribution.

Research on “experience effects” shows that, when making decisions, people put heavy weight on data they have personally lived through. In the current context, this suggests that people may base their stock market allocations on a weighted average of the past market returns they have personally experienced – not simply, as in Equation 2, a weighted average of *all* past returns. Malmendier & Nagel (2011) show that a framework of this type is helpful not only for explaining households’ allocations to the stock market, but stock market fluctuations as well.

In a model of aggregate stock market prices such as Barberis et al. (2015) that features both extrapolators and rational investors, when the P/E ratio of the stock market is high, the extrapolators forecast high future market returns. Importantly, though, at this very moment, rational investors forecast *low* returns: they recognize the overvaluation created by the extrapolators and sensibly forecast a correction. In other words, even a model of return

³The precise timing of the cycle depends on the model parameters. See Section 4.1 of Barberis (2018) for a more detailed explanation.

extrapolation predicts that *some* investors will have countercyclical expectations. As such, it is a puzzle that none of the six surveys studied by Greenwood and Shleifer (2014) reveals such countercyclical forecasts. More recent work is resolving the puzzle: studies of the beliefs of *institutional* investors find clearer evidence of countercyclical expectations (Dahlquist & Ibert 2024, 2026; Bastianello 2026b; Coutts, Goncalves & Loudis 2026; Hillenbrand & McCarthy 2026; Thesmar & Verner 2026).

3.1 Sources of return extrapolation

Why would people have extrapolative beliefs about returns in the first place? I now summarize some of the research on this. Based on the work to date, it is likely that return extrapolation has multiple sources.

There are two types of foundations for extrapolative beliefs: psychological mechanisms which lead investors to process past returns incorrectly; and mechanisms where investors use returns to make inferences about underlying fundamentals, but in a boundedly-rational way.

The psychological mechanisms are similar to those used to understand overreaction in forecasts of long-term earnings growth. This is not surprising: return extrapolation is a type of overreaction.

Base-rate neglect / representativeness. After seeing good returns on a risky asset, investors over-estimate the asset’s true mean return because they neglect the base rate – the fact that it is relatively rare for an asset to have a high true mean return.

Incorrect belief in a law of small numbers. If investors overinfer from small samples, then, after seeing a small sample of high returns on a risky asset, they will over-estimate the asset’s true mean return. Jin & Peng (2025) explore the connection between this bias and return extrapolation, and use it to make sense of retail investor trading.

Memory. Due to the Law of Recency, a fundamental feature of human memory whereby we recall recent instances more readily, an individual who is trying to recall an asset’s past returns is more likely to recall its recent past returns. These recent returns will then have an outsize influence on his forecast of future returns. Using data on both investors’ memories and their beliefs, Jiang et al. (2025) present direct evidence for this mental process.

Mechanisms based on boundedly-rational inference about fundamentals include:

Information diffusion. In the influential model of Hong and Stein (1999), some investors, seeing high past returns on an asset, infer that good news about its future cash flows must be diffusing through the investor population. They then reason that, over the next few months, this information will keep diffusing through the population and that the asset will therefore continue to earn high returns. However, this reasoning leads to *over*-extrapolation of past returns because, in some cases, it is incorrect: unbeknownst to the investors, the good news

about cash flows will sometimes have already fully diffused through the population.

Partial equilibrium thinking. Bastianello & Fontanier (2025) propose that, when uninformed investors observe high returns on an asset, they correctly infer that there must be some good news about the asset’s future cash flows; however, they forget that part of the high return is probably due to buying by other uninformed investors who, like themselves, inferred some good news. As a result, they become too optimistic about the asset’s future prospects.⁴

3.2 Open questions

There are several open questions related to return extrapolation. I discuss some of them below.

What value of the parameter θ in Equation 1 are investors using when forecasting future returns, and does this value of θ fit the facts about aggregate and individual stock prices? To see the kind of challenge that arises here, notice that, in the six surveys that Greenwood & Shleifer (2014) study, investors effectively have a low value of θ : their beliefs about the future return depend primarily on the past year or two of stock market returns. The problem is that, in an asset pricing model that incorporates such low- θ investors, the predicted price-to-dividend (P/D) ratio of the stock market will be much less persistent than in actual data: if investors’ beliefs about future returns depend heavily on the past year’s return, these investors will switch from bullish to bearish at a high frequency, leading to overly rapid oscillations in the P/D ratio. This disconnect has not yet been resolved.

Why does θ vary over time? Cassella and Gulen (2018) show that, for the average investor, θ varies over time: during market booms, it goes down, meaning that investors base their forecasts even more than usual on recent returns.

There is no consensus on why θ varies in this way, but there are two broad possibilities. It could be that, during market booms, each investor puts more weight on more recent returns. Alternatively, it could be that, during market booms, there is an influx of new investors who weight recent returns heavily, thereby leading the average θ across investors to fall. Chincio (2022) and Bastianello & Fontanier (2026) offer more specific theories of why θ varies over time.

Are models of return extrapolation consistent with facts about investor portfolios? Models of return extrapolation can explain facts about asset prices, but we also want them to be consistent with facts about investor portfolios. Are they?

At a qualitative level, the answer is yes: following good stock market returns, we see

⁴See Glaeser & Nathanson (2017) and Andre, Schirmer & Wohlfart (2026) for other foundations of return extrapolation based on boundedly-rational thinking.

positive net flows to equity mutual funds, with many households choosing to buy more stocks (Greenwood & Shleifer 2014, Gabaix et al. 2026).

At a quantitative level, however, there is a challenge that has not been resolved. A traditional model of portfolio choice, calibrated with reasonable parameter values, says that an investor who increases his estimate of the equity premium by 1%, from 3% to 4%, say, should increase his allocation to the stock market by about 10% – from 50% to 60%, say. However, Giglio et al. (2021) find that the empirical sensitivity is much lower: the typical investor increases his allocation only by about 1% – from 50% to 51%, say. This challenge applies to return extrapolation models too: in Barberis et al. (2015), for example, investors’ allocations are more sensitive to their beliefs than in Giglio et al.’s data.

Is the extrapolative demand in Equation 2 really about extrapolative beliefs? Over the past decade, perhaps due to the surge of interest in beliefs, it has become common to motivate the extrapolative *demand* in Equation 2 by saying that investors have the extrapolative *beliefs* in Equation 1. However, some of the evidence described above – the low sensitivity of allocations to beliefs, and the disconnect between the value of θ estimated from surveys and the persistence of the P/D ratio – raises the question of whether extrapolative demand is primarily, or even partly, due to extrapolative beliefs. After all, to account for excess volatility, predictability, momentum, and reversals in prices, it is enough to assume extrapolative demand; the additional assumption of extrapolative beliefs is not needed.

An alternative possibility is that extrapolative demand is driven by a *non-belief* mechanism. If investors are unable to discern, through introspection, what this mechanism is, they may mistakenly attribute it to beliefs: if, after good returns, they increase their allocation to the stock market for non-belief reasons, they may mistakenly attribute their action to greater optimism about the stock market’s future return. Under this view, the return expectations some investors report to survey compilers are beliefs that the investors reverse-engineered from actions taken for non-belief reasons, and not true beliefs.

One possible non-belief mechanism among many is “observational learning,” an algorithm used by the human brain according to which, if I observe another person being rewarded for taking some action, I raise the value I attach to this action and am more likely to take it myself (Burke et al. 2010). For example, if I observe the stock market performing well, I infer that many people are being rewarded for investing in the stock market. I therefore perceive stock market investing to be a valuable activity and am more likely to engage in it myself. This mechanism does not involve beliefs – it does not assign probabilities to future outcomes – but it generates extrapolative demand.

4 Gain-loss utility and prospect theory

The vast majority of models of investor behavior posit that these investors make decisions according to the Expected Utility (EU) framework. However, a large body of laboratory evidence shows that people depart from EU in systematic ways when making choices under risk. This has spurred the development of non-EU theories that capture people’s behavior more accurately. The most influential of these, by far, is prospect theory (Kahneman & Tversky 1979, Tversky & Kahneman 1992).

Prospect theory has four elements. *Reference dependence*: people derive utility not only from final wealth levels, but also from gains and losses. *Loss aversion*: people are more sensitive to potential losses than to potential gains. *Diminishing sensitivity*: people are risk averse over moderate-probability gains – they typically prefer \$500 to a 50% chance of \$1,000 – but risk-seeking over moderate-probability losses, typically preferring a 50% chance of losing \$1,000 to losing \$500 for sure. And *probability weighting*: people overweight low-probability outcomes, leading to a strong preference for positively-skewed payoff distributions and a strong aversion to negatively-skewed distributions.

A sustained research effort over the past 30 years has sought to use prospect theory to make sense of asset prices. Here, I briefly summarize some of this work. One theme of this research is that, when making investment decisions, people think about the future gains and losses that their decisions could give rise to, and about the utility impact of these gains and losses. This contrasts with traditional models of asset prices where people maximize utility functions defined only over consumption outcomes.

Benartzi & Thaler (1995) argue that prospect theory, and loss aversion in particular, can help explain the equity premium puzzle – the fact that the U.S. stock market has historically earned an average return in excess of T-Bills that is much higher than predicted by a garden-variety rational model. Inspired by their work, and by Kahneman and Tversky’s notion of reference dependence more generally, Barberis, Huang & Santos (2001) build a model of stock market prices in which investors derive utility not only from consumption, but also from annual gains and losses in their financial wealth. The investors are loss averse, in other words, more sensitive to losses in financial wealth than to gains. The model generates a large equity premium in line with the historical one. The intuition echoes that in Benartzi & Thaler (1995): an investor who allocates a sizable fraction of his wealth to the stock market knows that this will lead to large annual fluctuations in the value of his financial wealth; due to loss aversion, this will be painful on average, and so to compensate, he requires a high equity premium. De Giorgi & Legg (2012) point out that, if we also take account of probability weighting, then, since the distribution of stock market returns is negatively-skewed, this will lead to an even larger equity premium. Dimmock & Kouwenberg (2010) and Dimmock et al. (2021) provide portfolio-level evidence for these mechanisms: they find that individuals with higher loss aversion and stronger probability weighting are less likely

to participate in the stock market.

In an effort to provide a behavioral alternative to the traditional CAPM, Barberis, Jin & Wang (2021) analyze prospect theory’s implications for the cross-section of stock returns. They find that, when investors derive prospect theory utility from gains and losses in their holdings of individual stocks, the average return of a stock is determined not only by its beta – the only factor that matters in the CAPM – but also by three additional characteristics: *return volatility* (due to loss aversion, investors dislike volatility); *return skewness* (due to probability weighting, investors like skewness); and *capital gain overhang*, the average paper gain or loss across the investors holding the stock (as first argued by Grinblatt & Han (2005), due to diminishing sensitivity, investors accept a lower average return on their stocks that are trading at a loss in exchange for the opportunity to break even on these positions). Barberis, Jin & Wang (2021) show that this framework is helpful for thinking about a majority of 23 prominent stock market anomalies, including the momentum, profitability, distress, and volatility anomalies.

A key prediction of the above model, due to probability weighting, is that assets whose returns have higher skewness will be more overpriced and have a lower average return, even if the skewness is fully idiosyncratic.⁵ This prediction offers a way of explaining the low average return on many types of assets – volatile stocks and distressed stocks, but also IPO stocks and out-of-the-money options: the returns on these assets are highly positively-skewed. Chen, Hwang & Peng (2025) provide empirical support for this mechanism: they show that investors’ written rationales for buying the types of stocks that are known to have low average returns often emphasize the skewness in these stocks’ future returns.

These findings raise several questions, most obviously: Is prospect theory the right explanation for the applications discussed above? Are there other applications of these mechanisms that we have neglected? And can we attain a deeper understanding of these mechanisms by studying the cognitive roots of the elements of prospect theory? I return to this last question in Section 5.1. Below, I discuss two additional research directions.

Beliefs vs. preferences. The prospect theory view of the low average return on IPO stocks, say, is that, while investors have the correct perception of the distribution of IPO returns, they overweight the right tail of the distribution due to probability weighting and therefore overvalue any given IPO stock. An alternative view is that investors have incorrect beliefs: they over-estimate the *likelihood* that a given IPO stock will post a right-tail performance. Distinguishing between these views is not just of intellectual interest; it also has practical implications. If the high prices of IPO stocks are driven by incorrect beliefs, this is a clear mistake that, in principle, can be corrected. If they instead reflect a preference, intervention may not be needed.

⁵In earlier work on probability weighting and asset prices, Barberis & Huang (2008) show that this prediction – that idiosyncratic skewness is priced – emerges even in a model where investors derive prospect theory utility from gains and losses in *total* wealth.

The role of narrow framing. In their explanation of the equity premium puzzle, Benartzi & Thaler (1995) and Barberis, Huang & Santos (2001) assume that investors engage in “narrow framing”: they derive utility from fluctuations in *financial* wealth. This assumption is important: if investors instead derive utility from fluctuations in *total* wealth, the predicted equity premium is much smaller. The reason is that shocks to human capital have a low correlation with stock market returns; an investor who focuses on total wealth fluctuations that account for human capital will therefore be less scared of the stock market and will require a lower equity premium. The justification for assuming that investors focus on financial wealth is that financial wealth is what brokerage firms report to us. But this raises the question: Is it really the case that providing investors with information on narrowly-framed gains and losses leads them to take less risk? Given that, if true, this could have a first-order impact on the cost of capital, more experimental and field evidence on this question would be welcome.

5 New directions

5.1 Cognitive foundations

The 1970s and 1980s saw important advances in our understanding of human belief biases and preferences over risk. For much of the past 30 years, behavioral finance has focused on incorporating these advances, which are particularly associated with Daniel Kahneman and Amos Tversky, into financial economics. This effort continues today.

Over the past decade, a new area of research has begun to study the deeper cognitive foundations of these beliefs and preferences (Enke 2026). This work has several potential benefits: it can offer a deeper understanding of financial phenomena; it can generate new predictions about finance; and it can help unify the large number of scattered findings about beliefs and preferences.

There are several lines of research on cognitive foundations. These include: (i) cognitive noise and efficient coding (Woodford 2020, Frydman & Jin 2022, Enke & Graeber 2023, Oprea & Woodford 2026); (ii) memory (Bordalo et al. 2023, Wachter & Kahana 2024, Bordalo et al. 2025, Charles 2025, Jiang et al. 2025); (iii) attention (Bordalo, Gennaioli & Shleifer 2012, Gabaix 2019, Bohren et al. 2025, Loewenstein & Wojtowicz 2025); (iv) learning algorithms, including reinforcement learning (Barberis & Jin 2026); and (v) complexity (Puri 2025, Enke and Shubatt 2026, Shubatt and Yang 2026). Bordalo et al. (2026) unify some of these strands.

To give one example: In Section 4, we noted that probability weighting is helpful for understanding several facts about asset prices. Recent papers in behavioral economics have studied the deeper foundations of probability weighting: Where does it come from? One

answer builds on the psychology of attention: low-probability outcomes are often extreme and therefore attract attention, leading to overweighting (Bordalo, Gennaioli & Shleifer 2012). Another is cognitive uncertainty: if an individual has a hard time valuing a gamble offering a low-probability prize, he may shade his valuation toward that of a more typical, moderate-probability prize, leading to overvaluation of the initial gamble (Enke & Graeber 2023). Can these, or other, interpretations of probability weighting lead to useful new predictions about asset prices or investor behavior?

5.2 The decision-making process

Almost all models of investor behavior, whether rational or behavioral, use the following template: they specify the investor’s beliefs – either rational or biased – and then couple these beliefs with a preference specification, either Expected Utility or non-EU, to determine the person’s course of action. But what if people make decisions in part using an entirely different process? Decades of research in psychology suggest that this may indeed be the case.

One example of an alternative decision-making process is reinforcement learning. Suppose that an individual is considering a range of possible actions. Over time, he updates his estimate of the value of an action by trying it and seeing what happens: if he tries the action and the outcome is better than anticipated, he increases the value he assigns to that action; if the outcome is worse than anticipated, he lowers his estimate of the action’s value. At each moment of time, he is more likely to take an action, the higher its estimated value. Reinforcement learning is a prominent idea in both psychology and neuroscience because of extensive neural evidence that the human brain implements a type of reinforcement learning.

While reinforcement learning has a smaller footprint in economics, this may be a good time to apply it more widely. The reason is that cognitive scientists now have well-developed computational models of how human reinforcement learning works, and have found a way to combine these models with the more traditional framework in economics where people explicitly form beliefs about future outcomes. Barberis and Jin (2026) implement this “hybrid” framework in a financial context and show that it can shed light on a number of facts about investor behavior, such as experience effects in portfolio choice and the low sensitivity of allocations to beliefs.

5.3 Artificial intelligence (AI)

AI tools are sure to have a significant impact on behavioral finance, just as they will in every area of economics. Most obviously, they allow researchers to summarize and digest large troves of data. Bastianello, Decaire & Guenzel (2025), Ke (2026), and Sarkar (2026) are just three studies that illustrate the potential of this approach.

However, there are many AI-related research questions that are more specific to behavioral finance. Below, I mention three of them.

Can LLMs help households make better financial decisions? Akuzawa et al. (2026) argue that the answer is “yes”. They find that LLMs give sensible advice – to participate in the stock market, to be diversified, to invest in index funds. However, the authors also find that the quality of the advice depends on the quality of the prompt. This may limit the usefulness of LLMs: the households who need the advice the most may be the least able to craft the right prompt.

Can AI agents lead to more rational financial decision-making in markets? A first question is whether AI agents can learn to make sensible trading decisions, and in particular, to avoid common but subtle pitfalls. Can they learn not to chase high past returns or be drawn in by positively-skewed returns? A second question is: Even if his AI agent makes sensible trading recommendations, will the investor heed them, or instead give in to his own ill-conceived impulses? An AI agent *is* a commitment device, but only a soft one – there is no guarantee that the investor will stick to its suggestions. A third question is: What happens when AI agents tasked with investing their principals’ capital start interacting in large numbers in financial markets? At this point, the answer to this question is far from clear, and extensive research will be needed to grasp it.

Can LLMs simulate human beliefs? Sections 2 and 3 of this essay showed how survey data on economic expectations is helpful for understanding investor behavior. However, this survey data is limited in the asset classes, investor groups, and time periods that it covers. This raises the question: Can LLMs simulate human expectations about economic variables? If so, this would allow us to quickly generate a new type of survey data that could further advance our understanding of investor behavior. Studies such as Bybee (2025), Artrip et al. (2026), and Graham, Harvey & Jha (2026) have begun to investigate this possibility.

6 Conclusion

Behavioral finance has three main areas of application: asset prices, household finance, and corporate finance. This essay has primarily discussed applications to asset prices. And among the five behavioral asset pricing frameworks listed in the Introduction – limits to arbitrage; disagreement; biases in earnings forecasts; biases in return forecasts; and prospect theory – I have focused on the last three.

The behavioral finance research that has been painstakingly built up over the past 40 years by dozens of researchers now forms a large repository of knowledge that can be quickly deployed to make sense of ongoing developments. For example, in response to the commonly-asked question of 2025 and 2026 – “Is there a bubble in AI-related stocks?” – all five of the above frameworks have something to say. Given the disagreement about the prospects

of AI companies, the positive qualitative news about them, and the positive skewness that many perceive in them, the conditions are ripe, according to research conducted prior to the current boom and discussed in this article, for at least some overpricing. The high returns that accompany the overpricing are likely to be extrapolated by some investors, further amplifying the overvaluation. And limits to arbitrage mean that any overpricing will be difficult to correct.

Behavioral asset pricing has been vibrant over the past decade because of the strong interest in survey data on expectations. Behavioral ideas are also heavily used in household finance, and behavioral corporate finance is moving beyond its early classic contributions and exploring multiple new avenues. All of this makes me optimistic about the next 40 years of behavioral finance, and along with many others, I look forward to seeing the form it will take.

7 References

- Adam K, Marcet A, Beutel J. 2017. Stock price booms and expected capital gains. *Am. Econ. Rev.* 107:2352-408.
- Akuzawa M, Choukhmane T, de Silva T, Lin W. 2026. AI finance advice: Supply, demand, and life cycle implications. Working paper.
- Andre P, Schirmer P, Wohlfart J. 2026. Mental models of the stock market. *Q. J. Econ.* forthcoming.
- Artrip A, Horton J, Kazinnik S, Puzzello D, Zarifhonarvar A. 2026. Simulating the survey of professional forecasters. Working paper.
- Augenblick N, Lazarus E, Thaler M. 2025. Overinference from weak signals and underinference from strong signals. *Q. J. Econ.* 140:335-401.
- Ba C, Bohren JA, Imas A. 2026. Over- and underreaction to information: Belief updating with cognitive constraints. *Q. J. Econ.* forthcoming.
- Baker M, Ruback R, Wurgler J. 2007. Behavioral corporate finance: A survey. In *Handbook of Empirical Corporate Finance*, ed. E Eckbo, pp.145-86. Amsterdam: North Holland
- Baker M, Wurgler J. 2000. The equity share in new issues and aggregate stock returns. *J. Finance* 55:2219-57.
- Baker M, Wurgler J. 2002. Market timing and capital structure. *J. Finance* 57:1-32.
- Banerjee S, Gondhi N, Zhao X. 2026. Analyst forecast biases and the cost of wishful thinking. Working paper.
- Barber B, Odean T. 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *J. Finance* 55:773-806.
- Barberis N. 2018. Psychology-based models of asset prices and trading volume. In *Handbook of Behavioral Economics*, Vol. 1, ed. D Bernheim, S DellaVigna, D Laibson, pp.79-175. Amsterdam: North Holland
- Barberis N, Greenwood R, Jin L, Shleifer A. 2015. X-CAPM: An extrapolative capital asset pricing model. *J. Financ. Econ.* 115:1-24.
- Barberis N, Greenwood R, Jin L, Shleifer A. 2018. Extrapolation and bubbles. *J. Financ. Econ.* 129:203-27.
- Barberis N, Huang M. 2008. Stocks as lotteries: The implications of probability weighting for security prices. *Am. Econ. Rev.* 98:2066-100.

- Barberis N, Huang M, Santos T. 2001. Prospect theory and asset prices. *Q. J. Econ.* 116:1-53.
- Barberis N, Jin L. 2026. Model-free and model-based learning as joint drivers of investor behavior. *J. Finance* forthcoming.
- Barberis N, Jin L, Wang B. 2021. Prospect theory and stock market anomalies. *J. Finance* 76:2639-87.
- Barberis N, Shleifer A, Vishny R. 1998. A model of investor sentiment. *J. Financ. Econ.* 49:307-45.
- Barberis N, Thaler R. 2003. A survey of behavioral finance. In *Handbook of the Economics of Finance*, Vol. 1, ed. G Constantinides, M Harris, R Stulz, pp.1053-128. Amsterdam: North Holland
- Bastianello F. 2026a. Bias and predictability in analysts' beliefs. Working paper.
- Bastianello F. 2026b. Time-series and cross-section of risk premia expectations: Evidence from financial analysts. Working paper.
- Bastianello F, Decaire P, Guenzel M. 2025. Mental models and financial forecasts. Working paper.
- Bastianello F, Fontanier P. 2025. Expectations and learning from prices. *Rev. Econ. Stud.* 92:1341-74.
- Bastianello F, Fontanier P. 2026. Partial equilibrium thinking, extrapolation, and bubbles. *Rev. Financ. Stud.* forthcoming.
- Benartzi S, Thaler R. 1995. Myopic loss aversion and the equity premium puzzle. *Q. J. Econ.* 110:75-92.
- Benjamin D. 2019. Errors in probabilistic reasoning and judgment biases. In *Handbook of Behavioral Economics*, Vol. 2, ed. D Bernheim, S DellaVigna, D Laibson, pp.69-186. Amsterdam: North Holland
- Beshears J, Choi J, Laibson D, Madrian B. 2018. Behavioral household finance. In *Handbook of Behavioral Economics*, Vol. 1, ed. D Bernheim, S DellaVigna, D Laibson, pp.177-276. Amsterdam: North Holland
- Bohren JA, Hascher J, Imas A, Ungeheuer M, Weber M. 2025. Mental models of information: The role of attention and memory. Working paper.
- Bordalo P, Burro G, Coffman K, Gennaioli N, Shleifer A. 2025. Imagining the future: Memory, simulation, and beliefs. *Rev. Econ. Stud.* 92:1532-63.

- Bordalo P, Conlon J, Gennaioli N, Kwon S, Shleifer A. 2023. Memory and probability. *Q. J. Econ.* 138:265-311.
- Bordalo P, Gennaioli N, La Porta R, Shleifer A. 2019. Diagnostic expectations and stock returns. *J. Finance* 74:2839-74.
- Bordalo P, Gennaioli N, La Porta R, Shleifer A. 2024. Belief overreaction and stock market puzzles. *J. Political Econ.* 132:1450-84.
- Bordalo P, Gennaioli N, La Porta R, Shleifer A. 2025. Finance without exotic risk. *J. Financ. Econ.* 173. Article 104145.
- Bordalo P, Gennaioli N, Shleifer A. 2012. A salience theory of choice under risk. *Q. J. Econ.* 127:1243-85.
- Bordalo P, Gennaioli N, Shleifer A. 2018. Diagnostic expectations and credit cycles. *J. Finance* 73:199-227.
- Bordalo P, Gennaioli N, Shleifer A. 2026. Cognitive economics. *Annu. Rev. Econ.* forthcoming.
- Borgschulte M, Guenzel M, Liu C, Malmendier U. 2025. CEO stress, aging, and death. *J. Finance* 80:3401-42.
- Bouchaud J-P, Kruger P, Landier A, Thesmar D. 2019. Sticky expectations and the profitability anomaly. *J. Finance* 74:639-74.
- Burke C, Tobler P, Baddeley M, Schultz W. 2010. Neural mechanisms of observational learning. *PNAS* 107:14431-36.
- Bybee L. 2025. The ghost in the machine: Generating beliefs with large language models. Working paper.
- Cassella S, Gulen H. 2018. Extrapolation bias and the predictability of stock returns by price-scaled variables. *Rev. Financ. Stud.* 31:4345-97.
- Charles C. 2025. Memory moves markets. *Rev. Financ. Stud.* 38:1641-86.
- Chaudhry A. 2025. The impact of prices on analyst cash flow expectations: Reconciling subjective beliefs data with rational discount rate variation. *J. Financ. Econ.* 171. Article 104095.
- Chen H, Hwang B, Peng Z. 2025. Revisiting the cross-section of expected stock returns: Evidence from a textual analysis of buy recommendations. *Rev. Financ. Stud.* 38:3729-67.
- Chinco A. 2022. The ex ante likelihood of bubbles. *Manag. Sci.* 69:1222-44.

- Choi J. 2022. Popular personal financial advice versus the professors. *J. Econ. Perspect.* 36:167-92.
- Choukhmane T. 2025. Default options and retirement saving dynamics. *Am. Econ. Rev.* 115:3749-87.
- Couts S, Goncalves A, Loudis J. 2026. The subjective risk and return expectations of institutional investors. Working paper.
- Dahlquist M, Ibert M. 2024. Equity return expectations and portfolios: Evidence from large asset managers. *Rev. Financ. Stud.* 37:1887-928.
- Dahlquist M, Ibert M. 2026. Institutions' return expectations across assets and time. *J. Financ. Econ.* 175. Article 104188.
- Daniel K, Hirshleifer D, Subrahmanyam A. 1998. Investor psychology and security market under- and overreactions. *J. Finance* 53:1839-85.
- Daniel K, Titman S. 2006. Market reactions to tangible and intangible information. *J. Finance* 61:1605-43.
- De Bondt W, Thaler R. 1985. Does the stock market overreact? *J. Finance* 40:793-808.
- De Giorgi E, Legg S. 2012. Dynamic portfolio choice and asset pricing with narrow framing and probability weighting. *J. Econ. Dyn. Control* 36:951-72.
- De Long JB, Shleifer A, Summers L, Waldmann R. 1990. Noise trader risk in financial markets. *J. Political Econ.* 98:703-38.
- Delao R, Myers S. 2021. Subjective cash flow and discount rate expectations. *J. Finance* 76:1339-87.
- DellaVigna S, Pollet J. 2007. Demographics and industry returns. *Am. Econ. Rev.* 97:1667-702.
- DellaVigna S, Pollet J. 2009. Investor inattention and Friday earnings announcements. *J. Finance* 64:709-49.
- Didisheim A, Kelly B, Pourmohammadi M, Tian H. 2026. The inefficient pricing of news. Working paper.
- Dimmock S, Kouwenberg R. 2010. Loss-aversion and household portfolio choice. *J. Empir. Finance* 17:441-59.
- Dimmock S, Kouwenberg R, Mitchell O, Peijnenburg K. 2021. Household portfolio under-diversification and probability weighting: Evidence from the field. *Rev. Financ. Stud.* 34:4524-63.

- Enke B. 2026. The cognitive turn in behavioral economics. Working paper.
- Enke B, Graeber T. 2023. Cognitive uncertainty. *Q. J. Econ.* 138:2021-67.
- Enke B, Schwerter F, Zimmermann F. 2024. Associative memory, beliefs, and market interactions. *J. Financ. Econ.* 157. Article 103853.
- Enke B, Shubatt C. 2026. Quantifying lottery choice complexity. Working paper.
- Frydman C, Jin L. 2022. Efficient coding and risky choice. *Q. J. Econ.* 137:161-213.
- Gabaix X. 2019. Behavioral inattention. In *Handbook of Behavioral Economics*, Vol. 2, ed. D Bernheim, S DellaVigna, D Laibson, pp.261-343. Amsterdam: North Holland
- Gabaix X, Koijen R, Mainardi F, Oh S, Yogo M. 2026. Asset demand of U.S. households. Working paper.
- Giglio S, Maggiori M, Stroebel J, Utkus S. 2021. Five facts about beliefs and portfolios. *Am. Econ. Rev.* 111:1481-522.
- Glaeser E, Nathanson C. 2017. An extrapolative model of house price dynamics. *J. Financ. Econ.* 126:147-70.
- Graeber T, Roth C, Zimmermann F. 2024. Stories, statistics, and memory. *Q. J. Econ.* 139:2181-225.
- Graham J, Harvey C, Jha M. 2026. CFOs meet LLMs. Working paper.
- Greenwood R, Shleifer A. 2014. Expectations of returns and expected returns. *Rev. Financ. Stud.* 27:714-46.
- Grinblatt M, Han B. 2005. Prospect theory, mental accounting, and momentum. *J. Financ. Econ.* 78:311-39.
- Guenzel M, Hamilton C, Malmendier U. 2025. Prosociality and layoffs. Working paper.
- Harrison M, Kreps D. 1978. Speculative investor behavior in a stock market with heterogeneous expectations. *Q. J. Econ.* 92:323-36.
- Hillenbrand S, McCarthy O. 2026. Heterogeneous beliefs and stock market fluctuations. Working paper.
- Hirshleifer D, Lim S, Teoh SH. 2009. Driven to distraction: Extraneous events and underreaction to earnings news. *J. Finance* 64:2289-325.
- Hong H, Stein J. 1999. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *J. Finance* 54:2143-84.

- Hong W. 2026. Selective recall and the story-statistics gap in stock market misreaction. Working paper.
- Huang X. 2015. Thinking outside the borders: Investors' underreaction to foreign operations information. *Rev. Financ. Stud.* 28:3109-52.
- Jiang Z, Liu H, Peng C, Yan H. 2025. Investor memory and biased beliefs: Evidence from the field. *Q. J. Econ.* 140:2749-804.
- Jin L, Peng C. 2025. The law of small numbers in financial markets: Theory and evidence. *Rev. Financ. Stud.* forthcoming.
- Jin L, Sui P. 2022. Asset pricing with return extrapolation. *J. Financ. Econ.* 145:273-95.
- Kahneman D, Tversky A. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47:263-91.
- Ke B. 2026. Dissecting misreactions in earnings expectations: Evidence from analyst reports. Working paper.
- Kindleberger C. 1978. *A history of manias, panics, and crashes*. New York: Basic Books
- Kwon S, Tang J. 2026. Extreme categories and overreaction to news. *Rev. Econ. Stud.* 93:1137-66.
- Lee S, Maxted P. 2026. Credit card borrowing in heterogeneous-agent models: Reconciling theory and data. Working paper.
- Loewenstein G, Wojtowicz Z. 2025. The economics of attention. *J. Econ. Lit.* 63:1038-89.
- Malmendier U. 2018. Behavioral corporate finance. In *Handbook of Behavioral Economics*, Vol. 1, ed. D Bernheim, S DellaVigna, D Laibson, pp.277-379. Amsterdam: North Holland
- Malmendier U, Nagel S. 2011. Depression babies: Do macroeconomic experiences affect risk-taking? *Q. J. Econ.* 126:373-416.
- Malmendier U, Tate G. 2005. CEO overconfidence and corporate investment. *J. Finance* 60:2661-700.
- Malmendier U, Tate G. 2008. Who makes acquisitions? CEO overconfidence and the market's reaction. *J. Financ. Econ.* 89:20-43.
- Nagel S, Xu Z. 2022. Asset pricing with fading memory. *Rev. Financ. Stud.* 35:2190-245.
- Odean T. 1998. Are investors reluctant to realize their losses? *J. Finance* 53:1775-98.
- Oprea R, Woodford M. 2026. Behavioral signatures of cognitive noise. *Annu. Rev. Econ.* forthcoming.

- Puri I. 2025. Simplicity and risk. *J. Finance* 80:1029-80.
- Rabin M. 2002. Inference by believers in the law of small numbers. *Q. J. Econ.* 117:775-816.
- Rabin M, Vayanos D. 2010. The gambler's and hot hand fallacies: Theory and applications. *Rev. Econ. Stud.* 77:730-78.
- Sarkar S. 2026. Economic representations. Working paper.
- Scheinkman J, Xiong W. 2003. Overconfidence and speculative bubbles. *J. Political Econ.* 111:1183-219.
- Shiller R. 1981. Do stock prices move too much to be justified by subsequent changes in dividends? *Am. Econ. Rev.* 71:421-36.
- Shleifer A, Vishny R. 1997. The limits of arbitrage. *J. Finance* 52:35-55.
- Shubatt C, Yang J. 2026. Tradeoffs and comparison complexity. Working paper.
- Thaler R, Sunstein C. 2009. *Nudge: Improving decisions about health, wealth, and happiness*. New Haven: Yale University Press
- Thesmar D, Verner E. 2026. Beliefs and stock market fluctuations: New evidence from the past seven decades. Working paper.
- Tversky A, Kahneman D. 1971. Belief in the law of small numbers. *Psychol. Bull.* 76:105-10.
- Tversky A, Kahneman D. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertain.* 5:297-323.
- Wachter J, Kahana M. 2024. A retrieved-context theory of financial decisions. *Q. J. Econ.* 139:1095-147.
- Woodford M. 2020. Modeling imprecision in perception, valuation, and choice. *Annu. Rev. Econ.* 12:579-601.

		APPLICATIONS		
		Asset prices	Household finance	Corporate finance
PSYCHOLOGY	Beliefs			
	Preferences			
	Decision process			

Figure 1. Most behavioral finance papers can be placed somewhere in this matrix. Each row of the matrix corresponds to an aspect of human psychology – one related to beliefs, preferences, or the decision-making process – and each column to an empirical fact about asset prices, household finance, or corporate finance.