The Behavioral Foundations of Asset Prices: A Pedagogical Survey of Investor Psychology, Limits to Arbitrage, and Market Anomalies*

Agustín Shehadi Candela
QUAFI Research
research@quafi.io

April 2026

Abstract

Classical finance assumes investors are rational optimizers and that competition drives prices to their fundamental values. Decades of evidence — from the excess volatility of Shiller (1981) and the equity-premium puzzle of Mehra and Prescott (1985) to the long-horizon overreaction of De Bondt and Thaler (1985) — strain that assumption. *Behavioral finance* responds by grounding asset pricing in two empirically disciplined ideas: that real investors are subject to systematic psychological biases (Tversky and Kahneman, 1974; Kahneman and Tversky, 1979), and that the arbitrage which is supposed to correct the resulting mispricings is itself *limited* (De Long et al., 1990; Shleifer and Vishny, 1997). This didactic survey, written for a numerate reader with no specialist training, develops both pillars from first principles and shows how they combine into testable theories of market behavior. We catalog the core biases — heuristics, prospect-theoretic loss aversion, mental accounting, overconfidence, and the disposition effect (Kahneman and Tversky, 1979; Thaler, 1985; Odean, 1998; Shefrin and Statman, 1985) — and trace how each maps to a market phenomenon. We then present the three canonical behavioral models that reconcile short-run momentum with long-run reversal (Barberis et al., 1998; Daniel et al., 1998; Hong and Stein, 1999), and the sentiment-based view of the cross-section (Baker and Wurgler, 2006). Crucially, we give the efficient-markets rebuttal its full weight: Fama (1998) argues that anomalies split evenly between over- and underreaction and are fragile to methodology, a caution reinforced by the multiple-testing and replication evidence of Harvey et al. (2016) and Hou et al. (2020). We close with the adaptive-markets synthesis (Lo, 2004) and an agenda for further research — including the practical question, central to QUAFI's mission, of how a disciplined, quantitative process can help ordinary investors act against their own biases.

*QUAFI Working Paper Series, No. 2026-02. This is a survey and educational working paper of QUAFI Research. It is preliminary, circulated to inform and to invite comment; the views are the author's own. *Nothing herein is investment advice.* All references are to published, peer-reviewed sources; any errors are mine.

Keywords: behavioral finance; prospect theory; loss aversion; limits to arbitrage; investor sentiment; overreaction; underreaction; disposition effect; overconfidence; market anomalies; adaptive markets.
JEL: G02, G11, G12, G40, G41, D81, D91.

Contents

1 Introduction: two ways to read a market

There are two ways to explain why an asset trades where it does. The first, the *rational* tradition, says the price equals the market's best estimate of the asset's fundamental value: investors are rational, they process information efficiently, and any temporary mispricing is competed away by arbitrageurs (Fama, 1970). The second, the *behavioral* tradition, says prices are set by human beings who are predictably imperfect — who fear losses more than they value equivalent gains, who extrapolate recent trends, who are overconfident in their own judgment — and that the forces meant to correct their mistakes are weaker than the theory assumes.

These are not rival religions; they are complementary lenses, and the most useful practitioners switch between them deliberately. The purpose of this survey is to make the behavioral lens precise enough to use. Following the organizing framework of Barberis and Thaler (2003), behavioral finance rests on *two* pillars, and a phenomenon is only convincingly "behavioral" when both are present:

1. **Psychology** — a documented, systematic way in which investor beliefs or preferences depart from the rational ideal (§3);

2. **Limits to arbitrage** — a reason the "smart money" cannot simply trade the resulting mispricing away (§4).

We proceed as follows. Section 2 states the rational benchmark and the aggregate puzzles that first cracked it. Section 3 surveys the psychology of the individual investor. Section 4 explains why arbitrage is limited. Section 5 assembles these ingredients into the canonical behavioral asset-pricing models. Section 6 confronts the hard question of *measurement* — how one tells a real behavioral effect from a statistical mirage, and how seriously to take the efficient-markets rebuttal. Section 7 sets out open problems, and Section 8 concludes. A recurring practical thread — how disciplined process can protect investors from their own biases — connects the academic material to QUAFI's purpose.

2 The rational benchmark and the puzzles that strain it

The rational benchmark has two components. First, investors maximize *expected utility*: they weigh outcomes by their objective probabilities and care about final wealth. Second, markets are *efficient*: prices reflect available information, so they should respond only to genuine news and should not be predictable from public data (Fama, 1970). Together these imply that prices should be no more volatile than the fundamentals they track, and that broad, persistent mispricings should not exist.

Three early findings put that benchmark under strain.

- **Excess volatility.** Shiller (1981) showed that stock prices fluctuate far more than the present value of subsequent dividends can justify — prices "move too much" relative to the fundamentals they are supposed to forecast. Markets look more emotional than the rational model permits.

- **The equity-premium puzzle.** Mehra and Prescott (1985) found that the historical excess return of stocks over bonds is far too large to be explained by plausible risk aversion in the standard model. Either investors are implausibly risk-averse, or something other than fundamental risk is at work.

- **Long-horizon overreaction.** De Bondt and Thaler (1985) found that portfolios of past long-term "losers" subsequently outperform past "winners" — as if the market overreacts to extended runs of good or bad news and then corrects. This is hard to square with prices being unpredictable.

None of these *proves* irrationality — each has prompted ingenious rational explanations — but collectively they motivated a search for a model of the investor that is psychologically realistic.

3 The psychology of the individual investor

3.1 Bounded rationality and heuristics

Long before behavioral finance, Simon (1955) argued that human cognition is *boundedly rational*: people economize on scarce attention by using *heuristics* — mental shortcuts that are usually serviceable but systematically biased. Tversky and Kahneman (1974) catalogued the most important ones:

- **Representativeness** — judging probability by similarity to a stereotype, which leads investors to over-extrapolate (a few good earnings reports become "a great company," a few bad ones "a failing one").

- **Availability** — overweighting events that come easily to mind, such as recent or vivid market moves.

- **Anchoring** — clinging to an initial reference number when forming estimates. A concrete market example: George and Hwang (2004) show that a stock's distance from its 52-week high carries much of momentum's predictive power, as if investors anchor on that salient threshold.

Kahneman (2011) later organized these under the now-familiar metaphor of two mental systems — a fast, intuitive, error-prone "System 1" and a slow, deliberate "System 2" — a useful teaching device for why biases persist even among intelligent people.

3.2 Prospect theory: how people actually value risk

The cornerstone of behavioral finance is *prospect theory* (Kahneman and Tversky, 1979), later refined as cumulative prospect theory (Tversky and Kahneman, 1992). It replaces expected-utility theory with three psychologically grounded features.

1. **Reference dependence.** People evaluate outcomes as *gains and losses* relative to a reference point (often the purchase price), not as levels of final wealth.

2. **Loss aversion.** Losses hurt roughly twice as much as equivalent gains feel good — the value function is steeper below the reference point. This single fact has enormous consequences for how investors hold and sell.

3. **Probability weighting.** People overweight small probabilities (hence lotteries and crash insurance) and underweight moderate-to-large ones, distorting risky choice in predictable ways.

The value function is concave over gains and convex over losses — so people are *risk-averse* when ahead but *risk-seeking* when behind, willing to gamble to avoid realizing a loss. We will see this asymmetry resurface directly in trading behavior.

3.3 Mental accounting, framing, and the endowment effect

Thaler (1985, 1999) introduced *mental accounting*: people sort money into separate psychological "accounts" (a retirement account, a "play-money" account) and treat them non-fungibly, violating the rational principle that a dollar is a dollar. Closely related, *narrow framing* leads investors to evaluate each gamble in isolation rather than as part of a portfolio. The *endowment effect* — people demand more to give up an object than they would pay to acquire it (Kahneman et al., 1990, 1991) — is loss aversion applied to ownership, and helps explain inertia and the reluctance to rebalance.

A direct asset-pricing payoff is *myopic loss aversion* (Benartzi and Thaler, 1995): if loss-averse investors evaluate their portfolios too frequently, they experience the short-run pain of stock volatility acutely and demand a large premium to bear it — a behavioral resolution of the equity-premium puzzle of Mehra and Prescott (1985).

3.4 Overconfidence and the disposition effect in real trading

Two biases are visible directly in investors' brokerage records.

- **Overconfidence.** Investors overestimate the precision of their own information and trade too much as a result. Odean (1999) documents excessive trading; Barber and Odean (2000) show that the most active individual traders *underperform* precisely because trading costs erode returns; and Barber and Odean (2001) find men trade more than women and earn lower net returns, consistent with overconfidence. Overconfidence is also the engine of the behavioral model of Daniel et al. (1998) below.

- **The disposition effect.** Shefrin and Statman (1985) predicted, and Odean (1998) confirmed in trading data, that investors *sell winners too early and ride losers too long* — exactly the risk-seeking-in-losses pattern of prospect theory. Frazzini (2006) shows this same reluctance to realize losses causes prices to *underreact* to news, linking an individual bias to a market-level anomaly.

4 Why mispricing survives: the limits to arbitrage

Documenting biases is not enough. The efficient-markets reply is that even if some investors are irrational, rational arbitrageurs will trade against them and erase any mispricing. Behavioral finance's second pillar shows why this correction is incomplete.

De Long et al. (1990) formalize *noise-trader risk*: irrational traders' sentiment is itself unpredictable, so an arbitrageur who bets against mispricing faces the risk that sentiment grows *more* extreme before it corrects — mispricing can widen and bankrupt the arbitrageur first. Shleifer and Vishny (1997) sharpen the argument: real arbitrage is conducted by specialized, capital-constrained professionals managing other people's money. When a position moves against them, their investors withdraw funds at exactly the wrong moment, forcing liquidation when the opportunity is best. Arbitrage is therefore *risky, costly, and capital-constrained* — strongest exactly where it is most needed. Add real-world frictions (short-sale costs, the difficulty of finding a horizon over which mispricing reliably closes) and the conclusion is that psychological mispricings can persist for long periods.

This is why the *two pillars* must hold together: a bias creates the mispricing, and limited arbitrage lets it survive long enough to show up in prices.

5 From psychology to prices: behavioral asset-pricing models

The signature empirical challenge is that markets display *both* short-to-medium-run **momentum** (winners keep winning over 3–12 months, Jegadeesh and Titman, 1993) *and* long-run **reversal** (extreme movers revert over 3–5 years, De Bondt and Thaler, 1985; Jegadeesh, 1990). Any serious behavioral theory must generate this twin pattern. Three canonical models, all published in 1998–1999, do so from different psychological primitives.

- **Barberis, Shleifer, and Vishny (1998)** build on *representativeness* and *conservatism*. A representative investor oscillates between believing earnings are mean-reverting (causing *underreaction* to single news items, hence momentum) and believing they trend (causing *overreaction* to long streaks, hence reversal).

- **Daniel, Hirshleifer, and Subrahmanyam (1998)** build on *overconfidence* and *biased self-attribution.* Investors overweight their private signals, and when public news confirms them they grow *more* confident — driving prices past fundamentals (momentum) before the eventual correction (reversal).
- **Hong and Stein (1999)** need no individual irrationality of beliefs, only *bounded rationality and gradual information diffusion.* "Newswatchers" react to fundamentals that spread slowly (causing underreaction/momentum); "momentum traders" then chase the trend and push it too far (causing overreaction/reversal). The interaction of two limited groups produces both patterns.

That three different psychological stories all reproduce the momentum–reversal signature is both a strength (the pattern is robust to the exact mechanism) and a caution (the data underdetermine the mechanism) — a point we take up in §6.

At the aggregate level, Baker and Wurgler (2006) construct a measure of *investor sentiment* and show it predicts the cross-section: when sentiment is high, speculative, hard-to-value, hard-to-arbitrage stocks (small, young, volatile, non-dividend-paying) subsequently underperform, and vice versa — precisely the stocks where the two pillars bite hardest.

6 How to measure a behavioral effect — and how skeptical to be

A pedagogical survey would be irresponsible if it presented only the behavioral case. The discipline of the field lies in how hard it tries to *disprove* its own claims.

6.1 Evidence from the laboratory and from real accounts

Behavioral finance draws strength from two evidentiary sources the rational tradition lacks. *Controlled experiments* isolate a bias under conditions where the rational prediction is unambiguous — the endowment-effect tests of Kahneman et al. (1990) are the model. *Field evidence from trading records* (Odean, 1998; Barber and Odean, 2000) shows the same biases operating with real money and real consequences. When a bias appears in both the lab and the brokerage account, the case is strong.

6.2 The efficient-markets rebuttal, taken seriously

The most important counterweight is Fama (1998). His critique has three prongs that every behavioral claim must survive:

1. **Over- vs. underreaction split.** Across the anomaly literature, apparent overreaction is about as common as underreaction — consistent with *chance* deviations around efficiency

rather than a systematic behavioral bias in one direction.

2. **Methodological fragility.** Many long-horizon anomalies shrink or vanish under different but equally reasonable methods (e.g. value-weighting, alternative factor models). A result that depends on the technique is weak evidence.

3. **The joint-hypothesis problem.** Any claim of mispricing is jointly a claim about the risk model used to define "fair" return; a behavioral anomaly under one model can be a risk premium under a better one (Fama, 1970).

These cautions connect directly to the broader credibility crisis in empirical finance. Harvey et al. (2016) argue that, given how many factors have been tested, the bar for "discovery" must be far higher than the conventional one; Hou et al. (2020) re-test hundreds of anomalies under uniform methodology and find a large majority lose significance; and McLean and Pontiff (2016) show that documented predictors weaken markedly after publication — partly statistical artifact, partly the arbitrage the behavioral story itself predicts. The honest reader treats any single anomaly as provisional until it survives these filters.

6.3 A balanced verdict

The weight of evidence supports a measured conclusion. The *biases* are real and replicable at the individual level; their *aggregate price impact* is genuine but smaller, more episodic, and more fragile than enthusiastic early readings suggested. Lo (2004)'s *adaptive-markets hypothesis* offers a synthesis: rationality and bias coexist and compete, efficiency waxes and wanes with the population of strategies and the environment, and behavioral mispricings are real but *perishable*. Markets are neither perfectly efficient nor reliably foolish; they are a moving equilibrium between the two.

7 Open problems and directions for further research

1. **Aggregation.** The deepest open question is *when* individual biases survive aggregation into prices rather than cancelling out. The two-pillar framework (Barberis and Thaler, 2003) says limited arbitrage is the key, but a general, quantitative theory of which biases aggregate, and how strongly, is still missing.

2. **Measuring sentiment at scale.** Baker and Wurgler (2006) pioneered sentiment indices; modern text and alternative data (news, social media, search) promise far richer, real-time measures — but raise acute multiple-testing and overfitting risks (Harvey et al., 2016) that the field is only beginning to manage.

3. **Bias and machine learning.** Flexible models can detect behavioral footprints invisible to linear methods, yet a model without an economic mechanism is more likely a statistical artifact. Marrying behavioral theory with disciplined machine learning is an active frontier.

4. **Debiasing and choice architecture.** If investors are predictably biased, can products help them act against those biases? Default options, commitment devices, automatic rebalancing, and disciplined rules-based processes that pre-commit to selling winners and cutting losers — counteracting the disposition effect (Shefrin and Statman, 1985; Odean, 1998) and overtrading (Barber and Odean, 2000) — are exactly the design space in which a quantitative investment platform such as QUAFI operates. Rigorous field evidence on which designs actually improve outcomes is a valuable, under-explored agenda.

5. **Behavioral portfolio construction.** Mean-variance optimization assumes rational preferences; integrating loss aversion, mental accounting, and reference dependence (Kahneman and Tversky, 1979; Thaler, 1999) into *normative* portfolio advice that real investors will actually follow remains open.

6. **Perishability.** Under adaptive markets (Lo, 2004), the practical object is a bias's *remaining*, not historical, strength. Modeling how behavioral edges decay as they become known — the same arms race documented for quantitative anomalies (McLean and Pontiff, 2016; Hou et al., 2020) — is high-value and unsolved.

8 Conclusion

Behavioral finance does not overturn rational finance; it completes it. Real investors are boundedly rational and loss-averse (Simon, 1955; Kahneman and Tversky, 1979), they trade too much and sell their winners too soon (Odean, 1998; Barber and Odean, 2000), and because arbitrage is limited (Shleifer and Vishny, 1997) these tendencies leave fingerprints on prices — momentum, reversal, and sentiment-driven cross-sectional patterns (Barberis et al., 1998; Daniel et al., 1998; Hong and Stein, 1999; Baker and Wurgler, 2006). Yet the same discipline that documents these effects also warns, through Fama (1998) and the replication literature (Harvey et al., 2016; Hou et al., 2020), that their aggregate force is easy to overstate. The mature position is the adaptive one (Lo, 2004): biases and rationality coexist, and the edge they create is real but perishable.

For the individual investor the practical lesson is liberating. The biases are predictable, which means they are *manageable*. A disciplined, rules-based, cost-aware process — one that rebalances automatically, resists the urge to overtrade, and does not anchor on a purchase price — is, in effect, a technology for acting against one's own worst instincts. Building exactly that kind of disciplined process, accessible to non-specialists, is the purpose to which QUAFI's quantitative methodology is dedicated.

References

Malcolm Baker and Jeffrey Wurgler. Investor sentiment and the cross-section of stock returns. *Journal of Finance*, 61(4):1645–1680, 2006.

Brad M. Barber and Terrance Odean. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance*, 55(2):773–806, 2000.

Brad M. Barber and Terrance Odean. Boys will be boys: Gender, overconfidence, and common stock investment. *Quarterly Journal of Economics*, 116(1):261–292, 2001.

Nicholas Barberis and Richard Thaler. A survey of behavioral finance. In George M. Constantinides, Milton Harris, and René M. Stulz, editors, *Handbook of the Economics of Finance*, volume 1, pages 1053–1128. Elsevier, 2003.

Nicholas Barberis, Andrei Shleifer, and Robert Vishny. A model of investor sentiment. *Journal of Financial Economics*, 49(3):307–343, 1998.

Shlomo Benartzi and Richard H. Thaler. Myopic loss aversion and the equity premium puzzle. *Quarterly Journal of Economics*, 110(1):73–92, 1995.

Kent Daniel, David Hirshleifer, and Avanidhar Subrahmanyam. Investor psychology and security market under- and overreactions. *Journal of Finance*, 53(6):1839–1885, 1998.

Werner F. M. De Bondt and Richard Thaler. Does the stock market overreact? *Journal of Finance*, 40(3):793–805, 1985.

J. Bradford De Long, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. Noise trader risk in financial markets. *Journal of Political Economy*, 98(4):703–738, 1990.

Eugene F. Fama. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2):383–417, 1970.

Eugene F. Fama. Market efficiency, long-term returns, and behavioral finance. *Journal of Financial Economics*, 49(3):283–306, 1998.

Andrea Frazzini. The disposition effect and underreaction to news. *Journal of Finance*, 61(4): 2017–2046, 2006.

Thomas J. George and Chuan-Yang Hwang. The 52-week high and momentum investing. *Journal of Finance*, 59(5):2145–2176, 2004.

Campbell R. Harvey, Yan Liu, and Heqing Zhu. . . . and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–68, 2016.

Harrison Hong and Jeremy C. Stein. A unified theory of underreaction, momentum trading, and overreaction in asset markets. *Journal of Finance*, 54(6):2143–2184, 1999.

Kewei Hou, Chen Xue, and Lu Zhang. Replicating anomalies. *Review of Financial Studies*, 33 (5):2019–2133, 2020.

Narasimhan Jegadeesh. Evidence of predictable behavior of security returns. *Journal of Finance*, 45(3):881–898, 1990.

Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1):65–91, 1993.

Daniel Kahneman. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, New York, 2011.

Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979.

Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. Experimental tests of the endowment effect and the coase theorem. *Journal of Political Economy*, 98(6):1325–1348, 1990.

Daniel Kahneman, Jack L. Knetsch, and Richard H. Thaler. Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, 5(1):193–206, 1991.

Andrew W. Lo. The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, 30(5):15–29, 2004.

R. David McLean and Jeffrey Pontiff. Does academic research destroy stock return predictability? *Journal of Finance*, 71(1):5–32, 2016.

Rajnish Mehra and Edward C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161, 1985.

Terrance Odean. Are investors reluctant to realize their losses? *Journal of Finance*, 53(5): 1775–1798, 1998.

Terrance Odean. Do investors trade too much? *American Economic Review*, 89(5):1279–1298, 1999.

Hersh Shefrin and Meir Statman. The disposition to sell winners too early and ride losers too long: Theory and evidence. *Journal of Finance*, 40(3):777–790, 1985.

Robert J. Shiller. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review*, 71(3):421–436, 1981.

Andrei Shleifer and Robert W. Vishny. The limits of arbitrage. *Journal of Finance*, 52(1):35–55, 1997.

Herbert A. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, 69 (1):99–118, 1955.

Richard Thaler. Mental accounting and consumer choice. *Marketing Science*, 4(3):199–214, 1985.

Richard H. Thaler. Mental accounting matters. *Journal of Behavioral Decision Making*, 12(3): 183–206, 1999.

Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157):1124–1131, 1974.

Amos Tversky and Daniel Kahneman. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5(4):297–323, 1992.