**RESEARCH OVERVIEW**
**Nicholas Barberis, Yale University**
**July 2010**[1]

This note describes the research agenda my co-authors and I have developed over the past 15 years, and explains how our papers fit into that agenda. The intended audience is academic researchers, but I would like the note to be accessible to non-academic readers as well, so I've tried to avoid using too much technical jargon (although there is still some). I welcome your comments. If you find this write-up useful, please let me know. And if you find it confusing, please let me know that as well!

Most of my research is in the field of **behavioral finance**, and I focus on that research below. I have written papers on other topics – on asset allocation and on privatization, for example – but I don't discuss those papers here. And while I'm going to discuss most of my behavioral finance papers, I'm not going to discuss all of them.

What *is* behavioral finance? It is a large and active subfield of finance which investigates whether some financial phenomena might be the result of less than fully rational behavior. To understand what "less than fully rational" means, we should first clarify what "rational" means. In the field of finance, the label "rational" typically means two things. First, that people have sensible *preferences* -- technically, that they evaluate risk according to the Expected Utility framework (with a utility function that is defined over consumption in a dynamic model or over wealth in a single period model). And second, "rational" means that people have rational *beliefs*, in the sense that they update their beliefs according to Bayes' rule when they receive new information.

So now it is easier to explain what we study in behavioral finance. We study *non-standard preferences*, i.e. situations where people use non-Expected Utility models to evaluate risk, or where they do use Expected Utility but with a utility function that is defined over things other than consumption or wealth. And we study *non-standard beliefs*, i.e. situations where people deviate from Bayes' rule in forming their beliefs. To these two categories (non-standard preferences and non-standard beliefs), I will add a third category, namely *non-standard decision-making*. This third category covers situations where people have sensible preferences and beliefs, but use a non-standard *process* when they make a decision.

My work focuses on behavioral finance *theory*. I build models of financial markets in which people have non-standard preferences or non-standard beliefs, or in which they use a non-standard decision-making process. A natural question is: How can we figure out *how* people depart from full rationality, i.e. exactly what kinds of non-standard preferences or beliefs they might have? My preferred approach to answering this question is to study the psychology literature – both cognitive psychology and social psychology. Research in these fields documents a number of ways in which people systematically depart from full rationality. A lot of the important work in this area has been done by two

[1] This is a preliminary draft. Please do not quote or cite.

very famous psychologists, Daniel Kahneman and Amos Tversky. Much of my research takes the departures from full rationality that these two people have documented and incorporates them into models of financial markets. Once I have written a model down, I study its implications to see if it can help us make sense of puzzling financial phenomena. And I try to tease new predictions out of the model – predictions that I hope empiricists will test.

Below I describe my research in more detail. I structure my discussion in terms of the three categories described above: First I discuss papers on non-standard preferences; then some papers on non-standard beliefs; and finally some papers on non-standard decision-making. Within each category, I organize the papers according to the key departure from rationality that is driving the results. Here is a roadmap:

Non-standard preferences:

- loss aversion
- narrow framing
- probability weighting (static setting)
- probability weighting (dynamic setting)
- realization utility

Non-standard beliefs:

- representativeness, conservatism

Non-standard decision-making:

- categorization

The goal in the next few pages is to explain, intuitively, the ideas behind a number of my research papers. To keep the discussion brief, I will sometimes gloss over important technical issues. For full details, please see the original research papers.

Finally, I emphasize that the research agenda I describe here is not mine alone – it is shared with a sizeable group of other researchers in behavioral finance. Most of my papers are co-authored with people in this group. Both I, and the papers, have benefited enormously from these collaborations.

# NON-STANDARD PREFERENCES

## Loss aversion

Relevant papers:

(1) "Prospect Theory and Asset Prices," with Ming Huang and Tano Santos, *Quarterly Journal of Economics*, 2001.
(2) "The Loss Aversion / Narrow Framing Approach to the Equity Premium Puzzle," with Ming Huang, *Handbook of the Equity Premium*, 2007.

A famous idea of Kahneman and Tversky is that people are *loss averse*: they are more sensitive to losses – even small losses – than to gains of the same magnitude. My co-authors and I were interested in exploring the implications of loss aversion for financial markets. In particular, we were interested in the idea that, if investors are loss averse over annual fluctuations in the value of their stock market holdings, then that might help us understand the puzzlingly high historical equity premium (the high historical average return on the U.S. stock market relative to the return on Treasury Bills). The very simple logic is that, if people are loss averse, they will perceive the stock market to be very risky. They will think to themselves: "If the stock market goes up next year, that will feel good; but if it goes down, that will feel *really* bad." As a result of this kind of thinking, they require a high average return on equity in order to hold the market supply.

The idea that the equity premium might have something to do with loss aversion is not mine. The link was first made in a famous paper by Shlomo Benartzi and Richard Thaler (Benartzi and Thaler, 1995). What my co-authors and I do is to develop the argument in a couple of ways.

First, in [1], we embed the Benartzi-Thaler argument in a more formal model. Benartzi and Thaler analyze a model in which investors derive utility only from annual fluctuations in financial wealth. Virtually every other model in the field of asset pricing, however, assumes that investors derive utility only from consumption. We suspect that the truth lies somewhere in between. We therefore develop an asset pricing model in which investors derive utility both from consumption *and* from annual fluctuations in the value of their stock market holdings. (This turns out to be very hard to do, from a technical perspective, but we eventually figured out a way to do it). We find that the equity premium in our economy is high, but not as high as in Benartzi and Thaler (1995): in our framework, loss aversion is not as central a feature of investor preferences as it is in Benartzi and Thaler (1995).

The second contribution in [1] is to argue that experimental evidence on how loss aversion changes over time may also be useful for understanding financial markets. Specifically, there is experimental evidence that, after a prior loss, people become *more* loss averse than before (perhaps because, after getting one piece of bad news, they can't bear the thought of more bad news). There is also evidence that, after a prior gain, people

become *less* loss averse than before (perhaps because, with some good news as a cushion, they aren't as scared of potential future bad news).

Huang, Santos, and I build this evidence on changing loss aversion into our framework and argue that it might help us understand the puzzlingly high historical *volatility* of the stock market. The idea is that changing loss aversion amplifies fluctuations in the stock market caused by economic fundamentals. If bad fundamental news pushes the stock market down, this creates a loss for people holding the stock market and makes them more scared of further drops in the market. Their greater risk aversion pushes the stock market even further down, amplifying the initial shock.

[2] is a survey paper prepared for the *Handbook of the Equity Premium*. It reviews a number of models that build loss aversion into investor preferences and uses them to think about the equity premium. The theme, once again, is that, in the presence of loss aversion, the high historical equity premium isn't so puzzling.

**Narrow framing**

Relevant papers:

(3) "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing," with Ming Huang and Richard Thaler, *American Economic Review*, 2006.
(4) "Preferences with Frames: A New Utility Specification that Allows for the Framing of Risks," with Ming Huang, *Journal of Economic Dynamics and Control*, 2009.

If you are thinking about taking a financial risk of some kind, the rational way to evaluate it is to mix the new risk with other risks you are *already* facing and to check whether the resulting wealth distribution is an improvement or not. In experimental settings, however, people often fail to do this. Instead, they often seem to evaluate new risks *in isolation*, separately from other risks they are already facing. This is called "narrow framing".

In [3], my co-authors and I argue that narrow framing might be more common than previously realized. We take a simple empirical fact: the fact that most people turn down a 50:50 bet to win $110 or lose $100. It has long been thought that loss aversion (see above) is enough to explain this fact. We show that loss aversion is probably *not* enough. The argument is this. Most people who are offered the 50:50 bet are already facing *other* risks – stock market risk, housing risk, salary risk, and so on. So if they are loss averse but do *not* exhibit narrow framing, they will mix the 50:50 bet with these pre-existing risks and then check if the resulting wealth distribution is an improvement. And if they do that, they will find that they *should* take the 110/100 bet, because, roughly speaking, it diversifies their pre-existing risks.

In that case, how *can* we explain why people turn the 50:50 bet down? Our hypothesis is that people are not only loss averse, but that they also evaluate the 50:50 bet *in isolation* of their other risks – in other words, they frame it narrowly. And that, together with their loss aversion, explains why they turn it down (they think to themselves: "If I win, I get $110, which will feel good; but if I lose, I'll be down $100, which will feel really bad. So I'll pass on this bet").

Huang, Thaler, and I also argue that narrow framing may play an important role in financial markets. In other words, when people think about whether they should invest in the stock market, they may think about the stock market in isolation, separately from their other risks. And if they are thinking about whether to invest in a specific stock, they may think about that stock in isolation from their other risks. We point out that narrow framing of this kind may shed light on a number of puzzling phenomena: the fact that, historically, many households did not participate in the stock market; the fact that many households bias their portfolios heavily towards domestic, rather than international equity; and the fact that many households have a surprisingly large amount of money concentrated in just a few stocks.

In [4], Huang and I make a methodological contribution. While it seems plausible that, in the real world, people often engage in narrow framing, i.e. often evaluate risks in isolation, economists do not have a way of incorporating narrow framing into the formal models they use to study financial markets. Huang and I therefore develop a new model of investor preferences that *does* allow the researcher to incorporate narrow framing, and show how the model can be used to analyze both portfolio choice and asset pricing. We hope that this will be a useful tool for other researchers who want to study narrow framing.

**Probability weighting (static setting)**

Relevant paper:

(5) "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices," with Ming Huang, *American Economic Review*, 2006.

Another famous idea of Kahneman and Tversky is that the brain weights probabilities in a non-linear way – in particular, it overweights extreme, low-probability events. In [5], Ming Huang and I investigate the implications of this idea for financial markets. In short, we show that in a financial market in which investors overweight extreme events, assets with *positively* skewed returns are going to be *overpriced* and will therefore earn *low* returns, on average. And assets with negatively skewed returns will be underpriced and will earn high returns, on average.

Huang and I argue that this idea has many potential applications in finance. For example, it can help us understand why IPO stocks have a poor long-term average return. An important clue is that the returns of IPO stocks are highly positively skewed – most

stocks don't perform very well after their IPO, but some – Microsoft, say, or Google – do extraordinarily well. Our paper says that in a world where investors weight probabilities nonlinearly, IPO stocks *should* have a low average return. The intuition, roughly speaking, is that by taking a significant position in an IPO stock, you are giving yourself a chance – a small chance, admittedly – of becoming wealthy. According to Kahneman and Tversky, the brain overweights this extreme, low probability outcome. The IPO stock is therefore very attractive, and you are willing to pay a high price for it initially, and to accept a low average return on it.

In the years since we wrote this paper, a number of researchers have tested its core prediction: that assets with more positively skewed returns should have lower average returns. While more research is needed, the initial empirical results are broadly supportive of our prediction.

[A technical aside. The key theoretical prediction of our paper is that, if investors overweight extreme portfolio returns, they will overpay for positively skewed assets. This is more subtle than it seems. What is obvious is that an investor who overweights extreme portfolio returns will find a positively skewed *portfolio* desirable. What is much more subtle is to show that this investor will also find a positively skewed *stock* attractive. After all, if he only takes a *small* position in the stock, that won't add any skewness to his portfolio: even if the stock does well, he still won't become wealthy. So, at first sight, it's not clear that the investor will find a positively skewed stock appealing. The key to our prediction is to show that the investor optimally chooses a *large* position in the positively skewed stock. In this case, if the stock does well, the investor *does* become wealthy, and he really values this outcome. As a result, he is willing to pay a lot for the positively skewed stock and to accept a low average return on it].

**Probability weighting (dynamic setting)**

Relevant paper:

(6) "A Model of Casino Gambling," Working paper, 2010.

In the previous section, I discussed Kahneman and Tversky's idea that people overweight extreme, low probability events. The paper I referenced in that section, [5], studies the implications of this idea in a *static*, i.e. one-period model. This was appropriate for the financial applications I had in mind. But I then became interested in understanding the implications of probability weighting in a *dynamic*, i.e. multi-period model. This is the topic of [6].

It turns out that, in a dynamic context, probability weighting leads to something quite interesting, namely a *time-inconsistency*; in other words, it predicts that how people act in specific states of the world will often be different from how they initially *planned* to act in those states of the world. One place where I think this is very relevant is in casinos. Specifically, I show in [6] that, in a casino setting, the nonlinear probability weighting

postulated by Kahneman and Tversky predicts the following behavior: That people will enter the casino with a plan to stop gambling should they lose a certain amount of money – X dollars, say. But that if they *actually* lose X dollars, they will *continue* gambling, contrary to their initial plan. I like this prediction because I think it has a good chance of describing actual behavior. Anecdotally, at least, many people do seem to enter casinos with the intention of stopping if they lose a specific amount of money; but if they actually lose that amount, they keep going, contrary to their initial plan.

Notice that this framework also predicts that, if people are *aware* of the time-inconsistency, they will look for commitment devices that can help them overcome the inconsistency. For example, if your plan is to stop gambling once you lose $100, and you are aware that you may find it hard to stick to this plan, you might use the following commitment device: you might bring $100 with you to the casino and also leave your ATM card at home. That way, if you lose $100, you will really want to continue gambling, but you won't be able to, because you won't have your ATM card with you. Anecdotally, at least, some people do use commitment devices of this kind.

Incidentally, the model in [6] is intended to explain not only how people behave in casinos, but also why they go to casinos in the first place. I won't go into that part here – see the introduction to the paper if you're interested.

**Realization utility**

Relevant paper:

(7) "Realization Utility," with Wei Xiong, Working paper, 2010.

Wei Xiong and I have studied the idea that people might get a burst of utility at the moment that they *sell* an asset in their portfolio. For example, if you buy a stock at $40 and sell it at $60, you might get a *positive* burst of utility right then, at the moment of sale. And if you buy a stock at $20 and sell it at $10, you might get a *negative* burst of utility right then. We label this type of utility "realization utility," i.e. utility that derives from the act of *realizing* a gain or loss.

We think that the notion of realization utility is a plausible one. Surprisingly, though, there are very few papers in the finance literature that even mention it. When finance researchers model investor behavior, they almost always assume that investors derive utility only from wealth or consumption, not from the act of selling an asset. The idea of realization utility was first discussed by Shefrin and Statman (1985) but there has been very little mention of it in the literature since then.

In [7], Xiong and I analyze a model in which investors make decisions based on realization utility. We show that this simple idea can help us understand several different facts about trading behavior and asset prices. Perhaps the most obvious application is to something called the "disposition effect," the puzzling tendency of individual investors to

sell stocks that have risen in value, rather than fallen in value, since purchase. (Why is this behavior "puzzling"? Recall that, empirically, stocks exhibit *momentum*: stocks that have done well (poorly) in the past six months tend to keep doing well (poorly) in the next six months. Given this, the rational thing to do is to hold on to stocks that have been doing well and to get rid of stocks that have been doing poorly. But people tend to do exactly the opposite).

Why do the investors in our model exhibit a disposition effect? In other words, why do they have a greater propensity to sell stocks that have risen in value since purchase? Simply put, because it feels good to do so: When they sell a stock at a gain, investors get a burst of positive realization utility. (To some readers, this explanation may sound too simple. I agree that it is simple – but it may nonetheless be right! The only way to find out is to test the model's predictions; Xiong and I list a number of new predictions towards the end of the paper).

Another contribution of this paper, we hope, is to offer an explanation for why people might experience realization utility in the first place. In our view, the answer has to do with the way people think about their investing history. Under this view, people do not think about their investing history in terms of overall portfolio return; rather, they think about it as a series of investing "episodes", where each episode is characterized by three things: the name of the asset, the purchase price, and the sale price. "I bought IBM at $80 and sold it at $120" might be one such episode. "I bought my house for $280,000 and sold it for $320,000" might be another. So the reason an investor might get a positive burst of utility when he sells a stock at a gain is because, at the instant he sells the stock, he is creating a positive investing episode, one that is pleasant to look back on and to talk about. And the reason he gets a negative burst of utility when he sells a stock at a loss is because, at the instant he does so, he is creating a negative investing episode, one that is *not* pleasant to look back on or to talk about.

## NON-STANDARD BELIEFS

### Representativeness

Relevant paper:

(8) "A Model of Investor Sentiment," with Andrei Shleifer and Robert Vishny, *Journal of Financial Economics*, 1998.

In this paper, Shleifer, Vishny, and I study the implications of two important belief biases -- representativeness and conservatism – for asset prices.

Representativeness is another of Kahneman and Tversky's classic ideas. One of its many implications is something called "the law of small numbers", whereby people think that even a *small* sample should reflect the properties of the process that generated it. Put differently, it implies that people will try to infer the data-generating process from too

small a data sample. My co-authors and I point out that if investors in the stock market exhibit this bias, we should observe long-run mean reversion in stock prices, as well as a value premium. Representativeness is therefore one possible cause of the long-run mean reversion and value premium that we actually observe in the data.

The intuition behind our argument is straightforward. If a firm posts several periods of impressive earnings growth, investors who exhibit representativeness will be too quick to say that the firm's *true* earnings growth rate is high (their mistake is to forget that the sample is too short to allow for such a strong conclusion and that even a firm with average earnings growth can post a few periods of good earnings). As a result of this mistake, investors push the firm's stock price up too high. From that overvalued point, the stock must eventually decline. The stock therefore exhibits long-run mean-reversion: an increase in price followed by a decline.

Representativeness suggests that people sometimes put too *much* weight on sample data, relative to their priors. There is also evidence, however, that people sometime put too *little* weight on sample data. This evidence is sometimes given the label "conservatism".

In the same paper, Shleifer, Vishny, and I show that if investors also exhibit conservatism, then that has its own rich predictions, such as post-earnings announcement drift and medium-term momentum. As such, conservatism may be a way of understanding the post-earnings announcement drift and momentum that we actually see in the data. The intuition is again straightforward. If a firm reports surprisingly good earnings, investors who exhibit conservatism will be too slow to update their beliefs about the firm's prospects (they put too little weight on this data point relative to their priors). The stock price will therefore jump up too little on the day of the announcement. From that undervalued point, the stock's subsequent return will be high, on average. This is precisely post-earnings announcement drift.

Representativeness and conservatism are opposite effects, in some ways. In one case, people put too *much* weight on sample evidence; in the other, they put too *little* weight on sample evidence. It is therefore important to understand when one effect or the other dominates. Psychologists have not yet resolved this issue. In our model, we propose our own reconciliation. The idea is that if the investor sees a good earnings announcement *in isolation* – i.e. one that is not part of a *sequence* of good earnings – then conservatism dominates and the investor reacts too little. But if he sees a sequence of good earnings announcements, then representativeness takes over and he reacts too much.

## NON-STANDARD DECISION-MAKING

### Categorization

Relevant papers:

(9) "Style Investing," with Andrei Shleifer, *Journal of Financial Economics*, 2003.

(10) "Comovement," with Andrei Shleifer and Jeffrey Wurgler, *Journal of Financial Economics*, 2005.

The investment problem people face when they allocate their money across individual stocks is dizzyingly complex: How should they split their wealth across the thousands of different stocks out there? In [9], Andrei Shleifer and I argue that, to simplify this decision problem, people often make decisions at the level of asset *categories*. In other words, they first put stocks into categories – small-cap, mid-cap, large-cap, value, growth, technology stocks, utility stocks, etc – and then allocate funds at the level of these asset categories. We couple this categorization assumption with one another assumption, namely that people move money into categories that have recently done well and out of categories that have recently done poorly. So if small-cap stocks have recently done well relative to large-cap stocks, we assume that investors move into small-caps and out of large-caps. (This performance-chasing assumption can be motivated in a number of ways – for example, as a consequence of the representativeness heuristic discussed in the previous section).

Shleifer and I investigate the model's predictions. Perhaps the most interesting thing that comes out of the model is a new theory of return comovement. The traditional explanation for why a group of stocks have correlated returns is that they have correlated *earnings news* (e.g. automobile stocks move together because their earnings are correlated). Our model leads to an alternative theory: a group of stocks may have correlated returns because the stocks in the group comprise a salient *category* for many investors, and as those investors move money in and out of the category, the demand pressure makes the stocks in the category move together over and above what would be expected based on earnings correlation alone.

We use this theory to shed light on several puzzling instances of comovement – for example, the fact that small stocks and value stocks co-move in ways that cannot be fully attributed to earnings news. The idea is that small-cap stocks are a salient category for many investors, so as those investors move money into small-cap stocks this month and out of small-cap stocks the next, they make small-cap stocks move together even if their fundamentals are largely uncorrelated.

In [10], Shleifer, Wurgler, and I test a natural prediction of this framework, namely that immediately after a stock is added to the S&P 500 index, it should start co-moving more with other stocks in the index. The idea is that the S&P 500 index is a salient category for many investors. So as soon as a stock is added to the S&P 500, it is buffeted by investors' flows in and out of the index. This demand pressure should increase the stock's comovement with other stocks in the index.

The data support our prediction. Note that it is not easy to explain this finding using the traditional view of comovement – Standard and Poors do not add stocks to the index because they think that their earnings are about to start co-moving with those of other stocks already in the index. Nor is there any obvious reason why inclusion in the index

would cause a firm's earnings to suddenly start moving with those of other stocks in the index.

## REFERENCES

Benartzi, S. and R. Thaler (1995), "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics* 110, 73-92.

Shefrin, H. and M. Statman (1985), "The Disposition to Sell Winners too Early and Ride Losers too Long: Theory and Evidence," *Journal of Finance* 40, 777-790.