Behavioral Finance

Nicholas Barberis
Yale School of Management
October 2016

Overview

- from the 1950's to the 1990's, finance research was dominated by the **rational agent** framework
  - assumes that all market participants are rational
- starting in the 1990's, a second framework has emerged
  - the **behavioral finance** framework
- tries to make sense of financial phenomena using models where some people *are not fully rational*
- more broadly, using models that are *psychologically more realistic*

Overview

- on several dimensions, behavioral finance has been successful
  - explains observed facts in simple, intuitive ways; makes testable predictions
  - strong interest among academics, practitioners, and policy makers
  - citations, prizes
- but it still has a long way to go
  - one goal: that all active finance researchers are familiar with the core ideas in the field, and apply them as appropriate

Overview

- in this talk, I pick out three ideas from behavioral finance that appear to be particularly useful
  - over-extrapolation of the past
  - overconfidence
  - gain-loss utility with prospect theory
- show that these three ideas explain many of the central facts in asset pricing
- end with some broad remarks about progress in the field

Over-extrapolation

- the idea that, when forming beliefs about the future, people put too much weight on the recent past
- this can be applied to forecasts about either *fundamentals* or *returns*
  - if recent fundamentals have been good (poor), people over- (under-) estimate future fundamentals
  - if recent returns have been good (poor), people over- (under-) estimate future returns

Over-extrapolation

- here, focus on extrapolation of *returns*
  - models where some investors' beliefs about future price changes are a weighted average of past price changes

$$E_t^e(P_{t+1} - P_t) =$$
$$(1-\theta)((P_{t-1} - P_{t-2}) + \theta(P_{t-2} - P_{t-3})$$
$$+ \theta^2(P_{t-3} - P_{t-4}) + \ldots)$$

- return extrapolation is an old idea
  - can be found in qualitative accounts going back decades
  - first wave of academic research started in the 1990's (Cutler, Poterba, Summers 1990, De Long et al. 1990, Hong and Stein 1999)
- recently, there has been a second wave of research, spurred by *survey* data

Over-extrapolation

- several surveys ask investors, individual and institutional, for their forecasts of future stock market returns
- these data provide evidence of extrapolation
  - after good (poor) past returns, investors expect continued good (poor) returns
- but also of *over*-extrapolation
  - investors' beliefs are incorrect
  - Greenwood and Shleifer (2014)

The Role of Past Stock Market Returns in Explaining Survey Expectations

The solid line denotes the 12-month rolling nominal return on the CRSP VW stock index. The line marked with circles denotes the percentage of investors who are bullish in the Gallup survey.

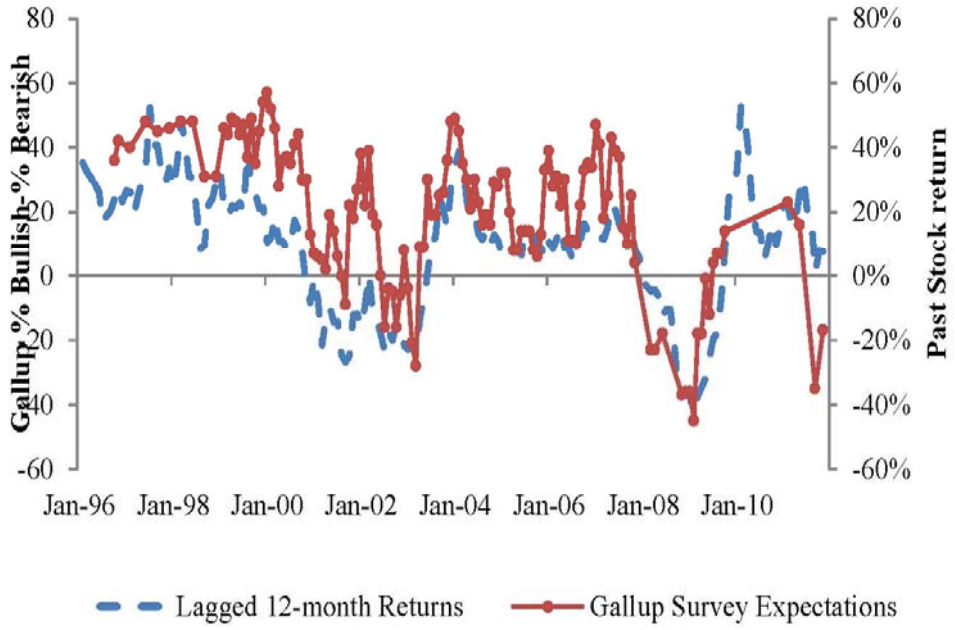

Over-extrapolation

- models in which some investors extrapolate past returns can explain several of the most important facts in asset pricing
  - momentum and reversals
  - time-series predictability
  - bubbles

Momentum and reversals

- in the cross-section of stocks, and also in other asset classes, there is medium-term momentum
  - stocks with high past 6-month returns have *higher* subsequent returns, on average, than stocks with low past 6-month returns
- but also long-term reversals
  - stocks with high past 3-year returns have *lower* subsequent returns, on average, than stocks with low past 3-year returns

Over-extrapolation

Time-series predictability

- ratios of price to fundamentals, e.g. P/E or P/D, negatively predict subsequent returns in the time series, in aggregate asset classes
  - e.g. the P/E ratio of the stock market negatively predicts the market's subsequent return

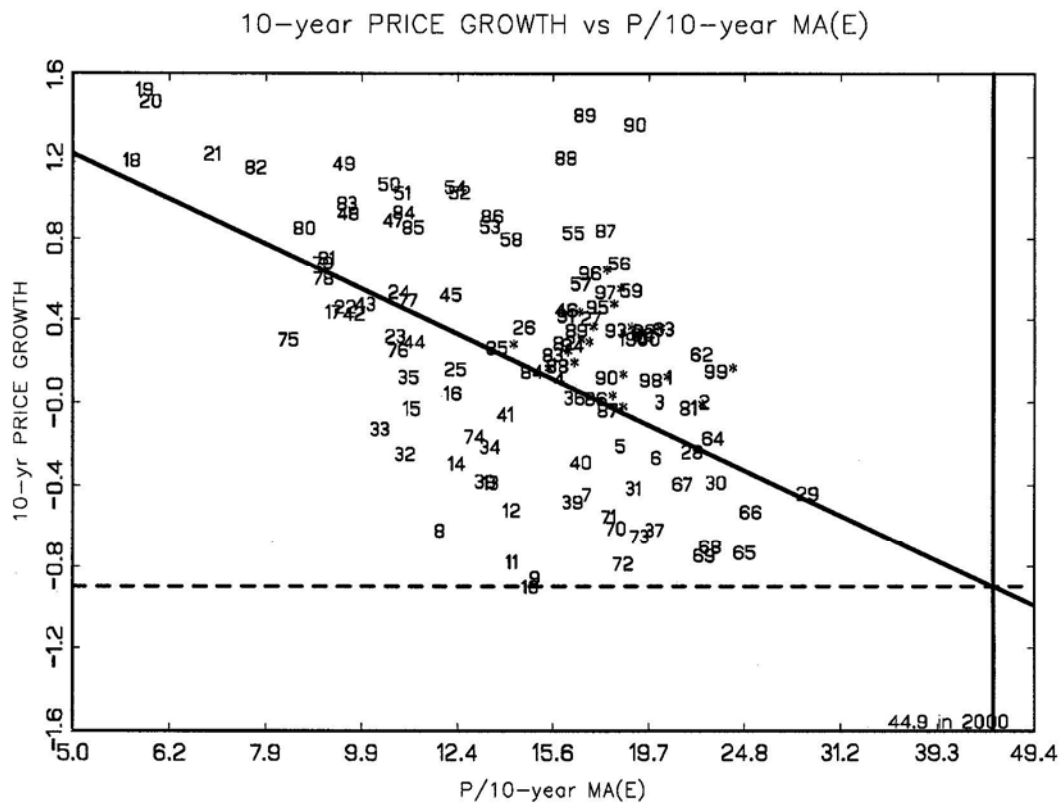

- source: Campbell and Shiller (2001)

Over-extrapolation

Bubbles

- episodes where:
  - the price of an asset rises dramatically and then collapses
  - and during the price rise, there is much talk of possible overvaluation and high volume

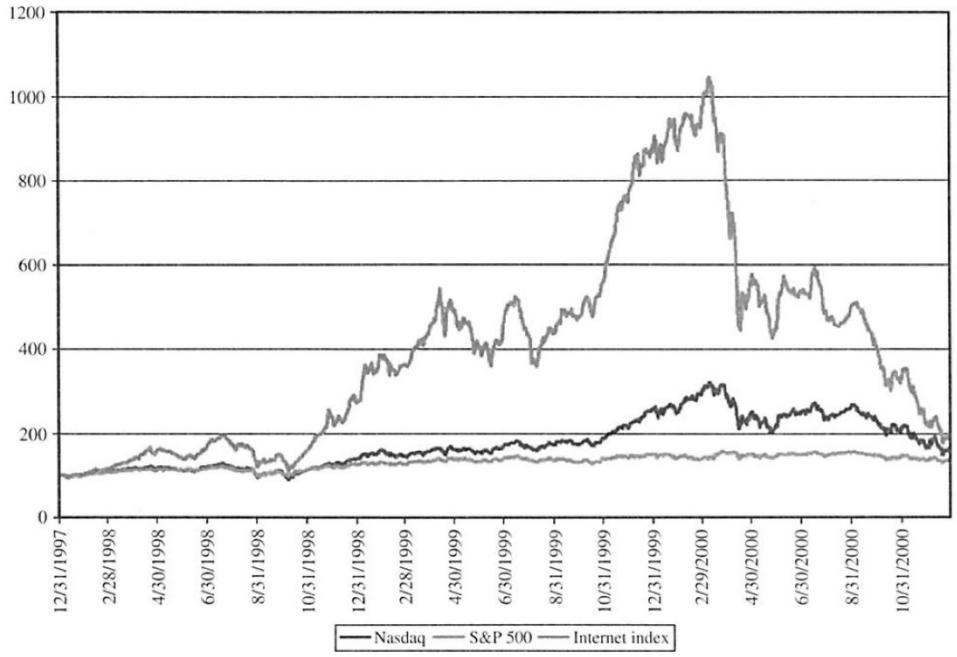

- source: Ofek and Richardson (2003)

Over-extrapolation

- models where some investors extrapolate past returns explain these facts
  - Barberis, Greenwood, Jin, Shleifer (2015, 2016)
- the graph below shows the price path of a risky asset in such an economy, following good cash-flow news

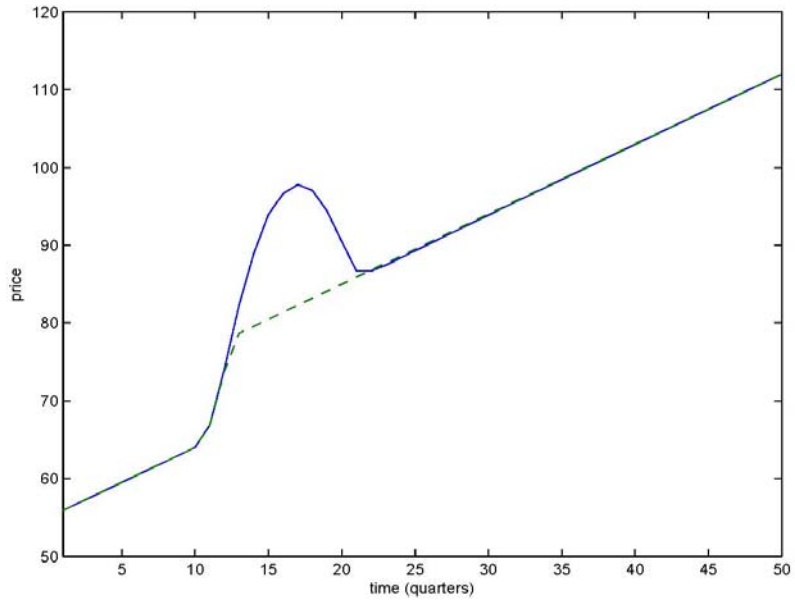

- we can (informally) see all three facts in this picture
- note: extrapolators earn profits at some points in the cycle
  - but more often, they perform poorly

Over-extrapolation

- a common question:
  - "How can extrapolation be a mistake, when 'momentum trading', which seems very similar, is thought to be a smart strategy?"
- answer: extrapolation and momentum trading are *not* the same
  - there is a crucial difference in timing
- if the price of the risky asset rises from time *t*-1 to *t*, momentum traders buy immediately at time *t*
  - but extrapolators buy one period later, at time *t*+1

$$E_t^e(P_{t+1} - P_t) =$$
$$(1-\theta)((P_{t-1} - P_{t-2}) + \theta(P_{t-2} - P_{t-3})$$
$$+ \theta^2(P_{t-3} - P_{t-4}) + ...)$$

- in this framework, momentum trading is profitable because it "front runs" the extrapolators

Over-extrapolation

- Cassella and Gulen (2016) use the extrapolation framework to uncover some striking results
- the relative weight θ that extrapolators put on recent vs. distant past returns varies over time
- more important, conditioning on θ improves our ability to forecast future returns
  - if the stock market is overvalued *and* θ is low, a short-term correction is much more likely

Panel A: D/P

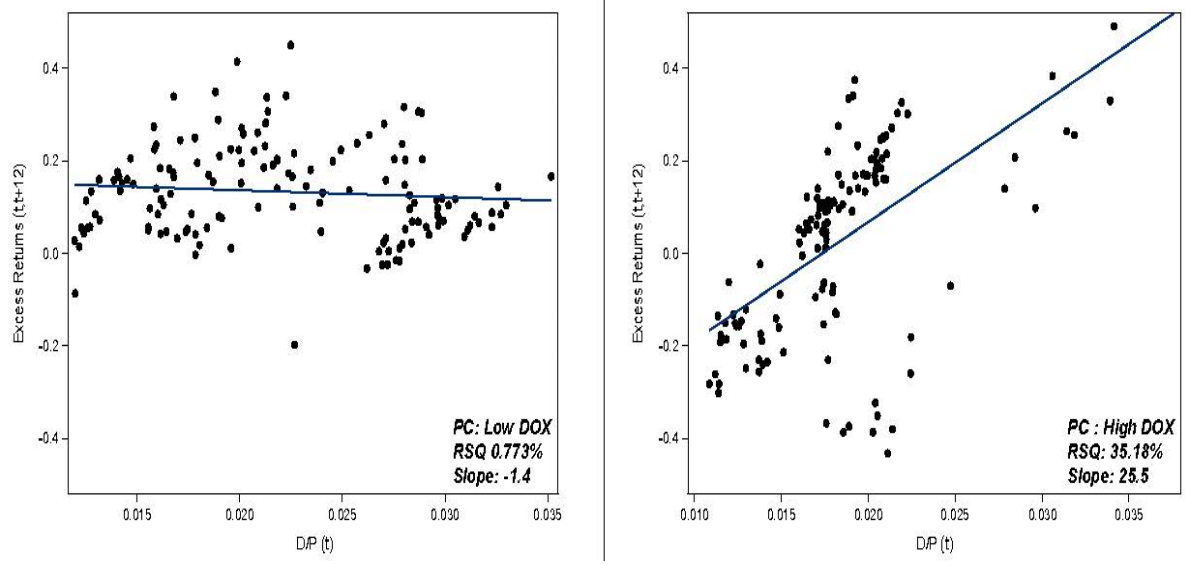

Over-extrapolation

- an important open question is:
  - *why* do investors extrapolate the past when forming beliefs?
- one hypothesis is that it stems from the "representativeness heuristic" (Kahneman and Tversky, 1974)
  - if the data reflect the essential characteristics of some model, people are too quick to embrace the model
  - they neglect the "base rate", i.e. the unconditional likelihood of the model
- extrapolation may also reflect a deep-seated "reward-seeking" behavior

Plan for the talk

Three core ideas:

- over-extrapolation of the past
- overconfidence
- gain-loss utility with prospect theory

Overconfidence

Type 1: “overplacement”

- people have overly rosy views of their abilities relative to other people
  - in surveys, over 80% believe themselves to be above the median on various dimensions

Type 2: “overprecision”

- people are too confident in the accuracy of their beliefs
  - 90% confidence intervals contain the correct answer approximately 50% of the time

Overconfidence

- the principal motivation for invoking overconfidence is the high trading volume in financial markets
  - non-speculative motives are unlikely to explain much of it
  - most trading is likely speculative, i.e. based on beliefs about future price changes
- key point: it is hard to generate a large amount of speculative trading in an economy with *rational* investors
  - each investor infers others' information from prices or from their willingness to trade
  - this reduces her own willingness to trade

Overconfidence

- overconfidence is an appealing way to break this logjam
- under this view, each investor overestimates the precision of her analysis, and underestimates the precision of others' analyses
  - heavy trading follows (Odean, 1998)

Overconfidence

- there is now direct evidence linking overconfidence to trading volume
- Grinblatt and Keloharju (2009) use military records to estimate overconfidence for a large sample of individuals in Finland
  - measure is self-reported confidence minus “appropriate” confidence inferred from aptitude tests
  - find a significant link between overconfidence and trading in subsequent years
- Glaser and Weber (2007) correlate trading frequency to measures of overplacement and overprecision among clients of an online brokerage
  - find a significant correlation for overplacement

Overconfidence

Key message:

- when we put on a trade, we should ask:
  - what makes us think we are on the right side of the trade?
  - that we are better informed than other market participants?
- overconfidence can be reduced through
  - a “why not?” approach
  - by explaining one’s reasoning in public
- other applications of overconfidence:
  - over- and under-valuation (Daniel, Hirshleifer, Subrahmanyam, 1998)
  - the popularity of active management, despite its lower average return than indices
  - firm acquisition activity (Malmendier and Tate, 2008)

Plan for the talk

Three core ideas:

- over-extrapolation of the past
- overconfidence
- gain-loss utility with prospect theory

Gain-loss utility / prospect theory

- so far, we have focused on people's beliefs
- we now turn to "preferences"
  - given people's beliefs about the potential future outcomes of an investment decision, how do they evaluate these outcomes?
- the vast majority of finance models assume that investors evaluate risk using "expected utility"
  - for any course of action, write down the potential future wealth outcomes
  - compute the utility of each outcome
  - multiply the utility of each outcome by the outcome's probability
  - sum up across outcomes
- the problem: experimental evidence suggests that expected utility is not an accurate description of decision-making under risk
  - and that "prospect theory" is much more accurate

Gain-loss utility / prospect theory

- prospect theory, due to Kahneman and Tversky (1979), is viewed by many psychologists as the best available summary of individual risk attitudes

Main features:

Reference dependence

- people think in terms of potential *gains* and *losses*

Loss aversion

- people are much more sensitive to potential losses

Gain-loss utility / prospect theory

Diminishing sensitivity

- people are risk averse in the domain of gains
  - but risk-seeking in the domain of losses

Probability weighting

- people process probabilities in a non-linear way
  - in particular, overweight low-probability outcomes
  - captures the simultaneous preference for both lottery tickets *and* insurance

Gain-loss utility / prospect theory

Probability weighting, ctd.

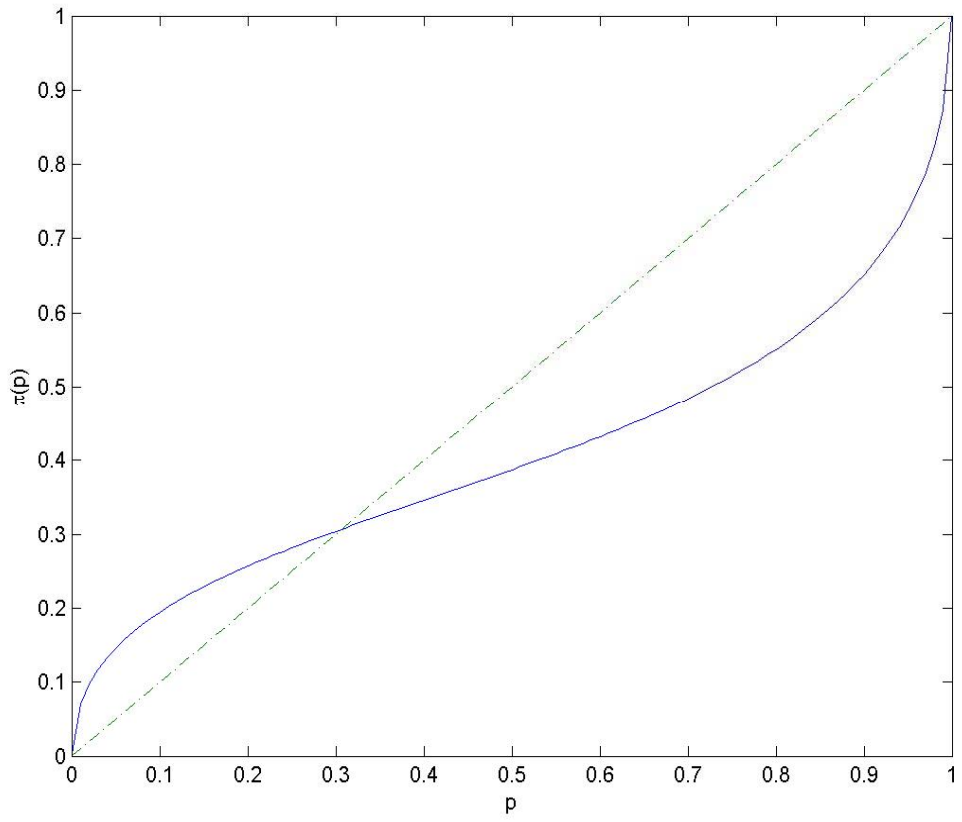

Gain-loss utility / prospect theory

- loss aversion is the best-known element of prospect theory
  - but it now appears that probability weighting and diminishing sensitivity have more applications in finance

Probability weighting

- probability weighting predicts that the *skewness* of an asset's returns will be priced
  - even idiosyncratic skewness
  - Barberis and Huang (2008)
- positively skewed assets will be overpriced and will earn low average returns
  - negatively skewed assets will be underpriced and will earn high average returns
- this prediction has many applications

Gain-loss utility / prospect theory

Application: *average returns, high and low*

- some average returns are puzzlingly *high*
  - the equity premium on the aggregate stock market
- other average returns are puzzlingly *low*
  - the average return on IPO stocks in the 5 years after issue

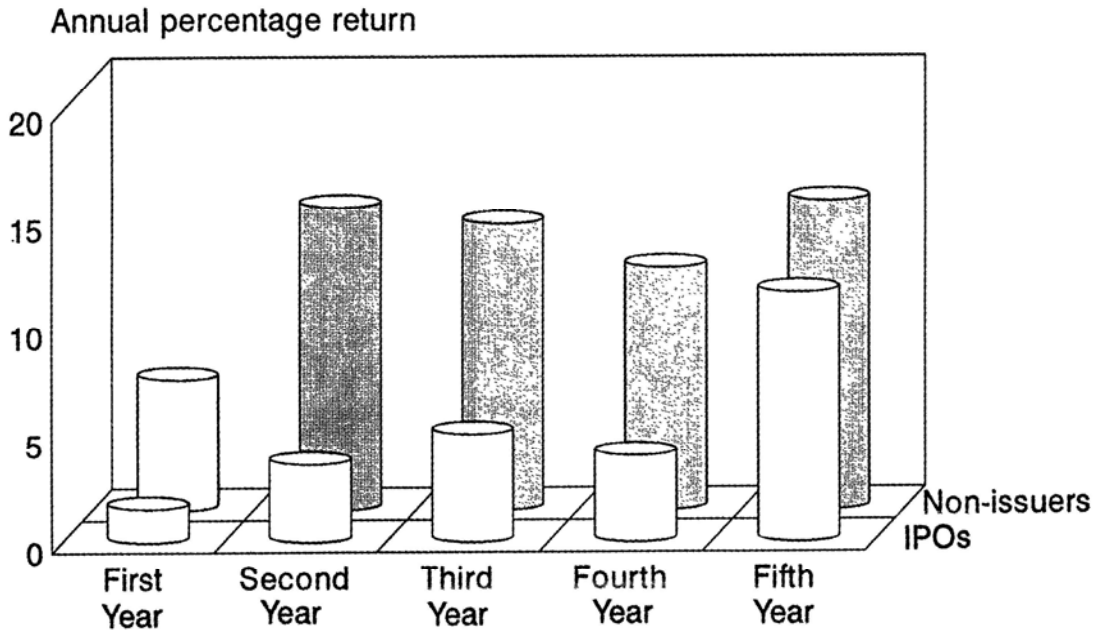

Gain-loss utility / prospect theory

Average returns, ctd.

- probability weighting provides a simple framework for understanding both facts
- the aggregate market has *negatively* skewed returns
  - under probability weighting, it should therefore have a high average return
- IPO stocks have *positively* skewed returns
  - under probability weighting, they should therefore have a low average return

Gain-loss utility / prospect theory

Average returns, ctd.

- the idea that positively skewed assets should have low average returns has many other applications
  - the low average returns of distressed stocks, bankrupt stocks, stocks traded in OTC markets
  - the low average return of out-of-the-money options
  - the low average return of stocks with high idiosyncratic volatility
  - Conrad, Kapadia, and Xing (2012), Boyer and Vorkink (2014), Eraker and Ready (2014)
- see also Ilmanen (2012) and Barberis (2013)

Gain-loss utility / prospect theory

Average returns, ctd.

- other studies have found support for the basic prediction: that more positively skewed assets will have lower average returns
  - Boyer, Mitton, and Vorkink (2010) use a regression model to predict skewness
  - Conrad, Dittmar, and Ghysels (2013) use an option-based measure of skewness

Gain-loss utility / prospect theory

Diminishing sensitivity

- diminishing sensitivity also has several applications
  - the disposition effect
  - momentum
- a striking recent application is to the risk-return relationship (Wang et al., 2016)
- the average raw return of volatile stocks is similar to that of less volatile stocks
  - the “beta” anomaly
- based on diminishing sensitivity, Wang et al. (2016) predict:
  - for stocks trading at a gain, there will be a *positive* relationship between volatility and returns
  - and a *negative* relationship between volatility and return for stocks trading at a loss
  - this would explain the flat overall relationship

Gain-loss utility / prospect theory

Diminishing sensitivity, ctd.

- Wang et al. (2016) show that these predictions hold in the data

	Proxy=IVOL		
P1	0.875	0.477	0.669
P3	0.233	0.492	0.797
P5	-1.050	0.072	0.989
P5-P1	-1.924	-0.405	0.320
t-stat	-6.00	-1.62	1.58
FF3-α	-2.093	-0.678	0.132
t-stat	-8.48	-3.59	0.73

- Source: Wang, Yan, and Yu (2016)

Discussion

- three ideas that appear particularly useful
  - over-extrapolation of the past
  - overconfidence
  - gain-loss utility with prospect theory
- these three ideas explain many of the central facts about asset prices:
  - average returns
  - time-series predictability
  - momentum and reversals
  - bubbles
  - trading volume
- and do so in simple, intuitive ways
- facts relating to volatility have been linked primarily to investor *beliefs*
  - facts about average returns have been linked mainly to *preferences*

Discussion

- in the 1990's, people worried about a "lack of discipline" in behavioral finance
- this concern has proven unfounded
  - the center of gravity of behavioral finance in the 1990's was in over-extrapolation, overconfidence, and gain-loss utility
  - today, the field's center of gravity remains in these three concepts
- extrapolation and gain-loss utility, in particular, are promising building blocks for an eventual "unified theory"

Conclusion

- behavioral finance has become successful not just by debating the rational side
  - but primarily by developing new models, making predictions, and conducting empirical tests
- this effort will continue with, hopefully, continued success for the field