Stocks as Lotteries: The Implications of Probability Weighting for Security Prices

Nicholas Barberis and Ming Huang
Yale University and Cornell University

February 2007*

Abstract

We study the asset pricing implications of Tversky and Kahneman's (1992) cumulative prospect theory, with particular focus on its probability weighting component. Our main result, derived from a novel equilibrium with non-unique global optima, is that, in contrast to the prediction of a standard expected utility model, a security's own skewness can be priced: a positively skewed security can be "overpriced," and can earn a negative average excess return. Our results offer a unifying way of thinking about a number of seemingly unrelated financial phenomena, such as the low average return on IPOs, private equity, and distressed stocks; the diversification discount; the low valuation of certain equity stubs; the pricing of out-of-the-money options; and the lack of diversification in many household portfolios.

JEL classification: D1, D8, G11, G12

Keywords: prospect theory, asset pricing, skewness, under-diversification

*We are grateful to Alon Brav, Michael Brennan, Markus Brunnermeier, John Campbell, Bing Han, Harrison Hong, Jon Ingersoll, Bjorn Johnson, Mungo Wilson, Hongjun Yan, and seminar participants at the AFA meetings, Columbia University, Cornell University, Dartmouth University, Duke University, Hong Kong University of Science and Technology, the NBER, New York University's 5-star conference, Ohio State University, the Stockholm Institute for Financial Research, the University of Illinois, and the University of Maryland for helpful comments. Comments are welcome at nick.barberis@yale.edu and mh375@cornell.edu.

1 Introduction

Over the past few decades, researchers have accumulated a large body of experimental evidence on attitudes to risk. This evidence reveals that, when people evaluate risk, they often depart from the predictions of expected utility. In an effort to capture the experimental data more accurately, economists have developed so-called non-expected utility models. Perhaps the most prominent of these is Tversky and Kahneman's (1992) "cumulative prospect theory."

In this paper, we study the pricing of financial securities when investors make decisions according to cumulative prospect theory. Our goal is to see if a model like cumulative prospect theory, which captures attitudes to risk in experimental settings very effectively, can also help us understand investor behavior in financial markets. Of course, there is no guarantee that this will be the case. Nonetheless, given the difficulties the expected utility framework has encountered in addressing a number of financial phenomena, it may be useful to document the pricing predictions of non-expected models and to see if these predictions shed any light on puzzling aspects of the data.

Cumulative prospect theory is a modified version of "prospect theory" (Kahneman and Tversky, 1979). Under this theory, people evaluate risk using a value function that is defined over gains and losses, that is concave over gains and convex over losses, and that is kinked at the origin; and using *transformed* rather than objective probabilities, where the transformed probabilities are obtained from objective probabilities by applying a weighting function. The main effect of the weighting function is to overweight the tails of the distribution it is applied to. The overweighting of tails does not represent a bias in beliefs; it is simply a modeling device for capturing the common preference for a lottery-like, or positively skewed, wealth distribution.

Previous research on the pricing implications of prospect theory has focused mainly on the implications of the kink in the value function (Benartzi and Thaler, 1995; Barberis, Huang, and Santos, 2001). Here, we turn our attention to other, less-studied aspects of cumulative prospect theory, and, in particular, to the probability weighting function.

First, we show that, in a one-period equilibrium setting with Normally distributed security payoffs and homogeneous investors, the CAPM can hold even when investors evaluate risk according to cumulative prospect theory. Under the assumption of Normality, then, the pricing implications of cumulative prospect theory are no different from those of expected utility.

Our second and principal result is that, as soon as we relax the assumption of Normality, cumulative prospect theory can have novel pricing predictions. We demonstrate this using the most parsimonious model possible, one with the minimum amount of additional structure.

Specifically, we introduce a small, independent, positively skewed security into the economy. In a standard concave expected utility model, this security would earn an average return of zero in excess of the risk-free rate. We show that, in an economy with cumulative prospect theory investors, the skewed security can become overpriced, relative to the prediction of the expected utility model, and can earn a *negative* average excess return. To be clear, an investor who overweights the tails of a portfolio return distribution will, of course, value a positively skewed *portfolio* highly; what is surprising is that he also values a positively skewed *security* highly, even if that security is small and independent of other risks.

Our unusual result emerges from an equilibrium structure which, to our knowledge, is new to the finance literature. In an economy with cumulative prospect theory investors and a skewed security, there are non-unique global optima, so that even though investors have *homogeneous* preferences, they can hold *different* portfolios. In particular, some investors take a large, undiversified position in the skewed security, because by doing so, they make the distribution of their overall wealth more lottery-like, which, as people who overweight tails, they find highly desirable. The skewed security is therefore very useful to these investors; as a result, they are willing to pay a high price for it and to accept a negative average excess return on it. We show that this effect persists even if there are several skewed securities in the economy. We also argue that the effect cannot easily be arbitraged away: while arbitrageurs can try to exploit the overpricing by taking short positions in skewed securities, there are significant risks and costs to doing so, and this limits the impact of their trading.

Our results suggest a unifying way of thinking about a number of seemingly unrelated empirical facts. Consider, for example, the low long-term average return on IPO stocks (Ritter, 1991). IPOs have positively skewed returns, probably because they are issued by young firms, a large fraction of whose value is in the form of growth options. Our analysis implies that, in an economy with cumulative prospect theory investors, IPOs can become overpriced and earn low average returns. Under cumulative prospect theory, then, the poor historical performance of IPOs may not be so puzzling. We discuss several other applications, including the low average return on private equity and on distressed stocks, the diversification discount, the low valuations of certain equity stubs, the pricing of out-of-the-money options, and the under-diversification in many household portfolios.

Through the probability weighting function, cumulative prospect theory investors exhibit a preference for skewness. There are already a number of papers that analyze the implications of skewness-loving preferences. We note, however, that the pricing effects we demonstrate here are new to the skewness literature. Earlier papers have shown that a security's coskewness with the market portfolio can be priced (Kraus and Litzenberger, 1976). We show that it is not just coskewness with the market that can be priced, but also a security's own skewness. For example, in our economy, a skewed security can earn a negative average excess return even if it is small and independent of other risks; in other words, even if its coskewness with the market is zero.

How do we obtain this new effect? In our model, the pricing of idiosyncratic skewness traces back to the fact that, in equilibrium, some investors hold an undiversified position in a skewed security. The earlier skewness literature considers economies in which investors have concave expected utility preferences. Such investors always hold diversified portfolios and, as a result, only coskewness with the market is priced; idiosyncratic skewness is not.

Our first result – that, under cumulative prospect theory, the CAPM can still hold – was originally proved by De Giorgi, Hens, and Levy (2003). We include this result here for two reasons. First, it provides a very useful springboard for our main contribution, namely the analysis of how skewed securities are priced. Second, we are able to offer a different proof of the CAPM result, one that is much shorter. As part of our proof, we show that, within certain classes of distributions, cumulative prospect theory preferences satisfy *second*-order stochastic dominance – a result that is interesting in its own right and that is new to the literature.

In Section 2, we discuss cumulative prospect theory and its probability weighting feature in more detail. In Section 3, we present our assumptions on investor preferences. We then examine how cumulative prospect theory investors price Normally distributed securities (Section 4) and positively skewed securities (Section 5). Section 6 considers applications of our results and Section 7 concludes.

2 Cumulative Prospect Theory and Probability Weighting

Tversky and Kahneman's (1992) cumulative prospect theory is one of the best-known models of decision-making under risk. We introduce it by first reviewing the original version of prospect theory, laid out by Kahneman and Tversky (1979), on which it is based.

Consider the gamble

$$(x, p; y, q),$$

to be read as "get x with probability p and y with probability q, independent of other risks," where $x \leq 0 \leq y$ or $y \leq 0 \leq x$, and where $p + q = 1$. In the expected utility framework, an agent with utility function $U(\cdot)$ evaluates this risk by computing

$$pU(W + x) + qU(W + y), \tag{1}$$

where W is his current wealth. In the original version of prospect theory, the agent assigns the gamble the value

$$\pi(p)v(x) + \pi(q)v(y), \tag{2}$$

where $v(\cdot)$ and $\pi(\cdot)$ are known as the value function and the probability weighting function, respectively. Figure 1 shows the forms of $v(\cdot)$ and $\pi(\cdot)$ suggested by Kahneman and Tversky (1979). These functions satisfy $v(0) = 0$, $\pi(0) = 0$, and $\pi(1) = 1$.

There are four important differences between (1) and (2). First, the carriers of value in prospect theory are gains and losses, not final wealth levels: the argument of $v(\cdot)$ in (2) is x, not $W + x$. This is motivated in part by experimental evidence, but is also consistent with the way in which our perceptual apparatus is more attuned to a *change* in the level of an attribute – brightness, loudness, or temperature, say – than to the level itself.

Second, the value function $v(\cdot)$ is concave over gains, but convex over losses. Kahneman and Tversky (1979) infer this from subjects' preference for a certain gain of \$500 over[1]

$$(\$1000, \frac{1}{2}),$$

and from their preference for

$$(-\$1000, \frac{1}{2})$$

over a certain loss of \$500. In short, people are risk averse over moderate-probability gains, but risk-seeking over moderate-probability losses.

Third, the value function is kinked at the origin, so that the agent is more sensitive to losses – even small losses – than to gains of the same magnitude. This element of prospect theory is known as loss aversion. Kahneman and Tversky (1979) infer the kink from the widespread aversion to bets of the form

$$(\$110, \frac{1}{2}; -\$100, \frac{1}{2}).$$

Such aversion is hard to explain with differentiable utility functions, whether expected utility or non-expected utility, because the very high local risk aversion required to do so typically predicts implausibly high aversion to large-scale gambles (Epstein and Zin, 1990; Rabin, 2000; Barberis, Huang, and Thaler, 2006).

Finally, under prospect theory, the agent does not use objective probabilities when evaluating the gamble, but rather, transformed probabilities obtained from objective probabilities via the probability weighting function $\pi(\cdot)$. This function has two salient features. First, low probabilities are overweighted: in the lower panel of Figure 1, the solid line lies above the dotted line for low p. Given the concavity (convexity) of the value function in the region of gains (losses), this is inferred from people's preference for

$$(\$5000, 0.001)$$

[1]We abbreviate $(x, p; 0, q)$ to (x, p).

over a certain \$5, and from their preference for a certain loss of \$5 over

$$(-\$5000, 0.001);$$

in other words, it is inferred from their simultaneous demand for both lotteries and insurance. Spelling this out in more detail, note that

$$\begin{aligned}(\$5, 1) &\prec (\$5000, 0.001)\\ &\Rightarrow v(5)\pi(1) < v(5000)\pi(0.001) < 1000\ v(5)\pi(0.001)\\ &\Rightarrow \pi(0.001) > 0.001,\end{aligned}$$

so that low probabilities are overweighted. A similar calculation in the case of the $(-\$5000, 0.001)$ gamble, using the fact that $v(\cdot)$ is convex over losses, produces the same result.

The other main feature of the probability weighting function is a greater sensitivity to differences in probability at higher probability levels: in the lower panel of Figure 1, the solid line is flatter for low p than for high p. For example, subjects tend to prefer a certain \$3000 to $(\$4000, 0.8)$, but also prefer $(\$4000, 0.2)$ to $(\$3000, 0.25)$. This pair of choices violates expected utility, but, under prospect theory, implies

$$\frac{\pi(0.25)}{\pi(0.2)} < \frac{\pi(1)}{\pi(0.8)}. \tag{3}$$

The intuition is that the 20 percent jump in probability from 0.8 to 1 is more striking to people than the 20 percent jump from 0.2 to 0.25. In particular, people place much more weight on outcomes that are certain relative to outcomes that are merely probable, a feature sometimes known as the certainty effect.

The transformed probabilities $\pi(p)$ and $\pi(q)$ should not be thought of as beliefs, but as decision weights that help capture evidence on individual risk attitudes. In Kahneman and Tversky's (1979) framework, an agent evaluating the lottery-like $(\$5000, 0.001)$ gamble understands that he will only receive the \$5000 will probability 0.001. The overweighting of 0.001 introduced by prospect theory is simply a modeling device for capturing the agent's preference for the lottery over a certain \$5.

In this paper, we do not work with the original version of prospect theory, but with a modified version, cumulative prospect theory, proposed by Tversky and Kahneman (1992). In this version, Tversky and Kahneman (1992) suggest explicit functional forms for $v(\cdot)$ and $\pi(\cdot)$. Moreover, they apply the probability weighting function to the *cumulative* probability distribution, not to the probability density function. This ensures that the preferences do not violate first-order stochastic dominance – a weakness of the original 1979 version of prospect theory – and also that they can be applied to gambles with any number of outcomes, not just two. Finally, Tversky and Kahneman (1992) allow the probability weighting functions for gains and losses to differ.

Formally, cumulative prospect theory says that the agent evaluates a gamble

$$(x_{-m}, p_{-m}; \ldots; x_{-1}, p_{-1}; x_0, p_0; x_1, p_1; \ldots; x_n, p_n),$$

where $x_i < x_j$ for $i < j$ and $x_0 = 0$, by assigning it the value

$$\sum_{i=-m}^{n} \pi_i v(x_i), \tag{4}$$

where

$$\pi_i = \begin{cases} w^+(p_i + \ldots + p_n) - w^+(p_{i+1} + \ldots + p_n) \\ w^-(p_{-m} + \ldots + p_i) - w^-(p_{-m} + \ldots + p_{i-1}) \end{cases} \text{for} \begin{array}{c} 0 \leq i \leq n \\ -m \leq i < 0 \end{array}, \tag{5}$$

and where $w^+(\cdot)$ and $w^-(\cdot)$ are the probability weighting functions for gains and losses, respectively.

Equation (5) emphasizes that, under cumulative prospect theory, the weighting function is applied to the cumulative probability distribution. If it were instead applied to the probability density function, as in the original prospect theory, the probability weight π_i, for $i < 0$ say, would be $w^-(p_i)$. Instead, equation (5) shows that, under cumulative prospect theory, the probability weight π_i is obtained by taking the total probability of all outcomes equal to or worse than x_i, namely $p_{-m} + \ldots + p_i$, the total probability of all outcomes strictly worse than x_i, namely $p_{-m} + \ldots + p_{i-1}$, applying the weighting function to each, and subtracting one from the other.

The effect of applying the probability weighting function to a *cumulative* probability distribution is to make the agent overweight the *tails* of that distribution. In equations (4)-(5), the most extreme outcomes, x_{-m} and x_n, are assigned the probability weights $w^-(p_{-m})$ and $w^+(p_n)$, respectively. Since they are probability weighting functions, $w^-(\cdot)$ and $w^+(\cdot)$ overweight low probabilities, so that if p_{-m} and p_n are small, $w^-(p_{-m}) > p_{-m}$ and $w^+(p_n) > p_n$. The most extreme outcomes – the outcomes in the tails – are therefore overweighted. Just as in the original prospect theory, then, a cumulative prospect theory agent likes positively skewed, or lottery-like, wealth distributions. This will play an important role in our analysis.

Tversky and Kahneman (1992) propose the functional forms

$$v(x) = \begin{cases} x^\alpha \\ -\lambda(-x)^\alpha \end{cases} \text{for} \begin{array}{c} x \geq 0 \\ x < 0 \end{array}, \tag{6}$$

and

$$w^+(P) = w^-(P) = w(P) = \frac{P^\delta}{(P^\delta + (1-P)^\delta)^{1/\delta}}. \tag{7}$$

For $0 < \alpha < 1$ and $\lambda > 1$, $v(\cdot)$ captures the features of the value function highlighted earlier: it is concave over gains, convex over losses, and exhibits a greater sensitivity to losses than to gains. The degree of sensitivity to losses is determined by λ, which is known as the coefficient of loss aversion. For $0 < \delta < 1$, $w(\cdot)$ captures the features of the weighting function described earlier: it overweights low probabilities, so that $w(P) > P$ for low P, and is flatter for low P than for high P.

Using experimental data, Tversky and Kahneman (1992) estimate $\alpha = 0.88$, $\lambda = 2.25$, and $\delta = 0.65$. Figure 2 shows the form of the probability weighting function $w(\cdot)$ for $\delta = 0.65$ (the dashed line), for $\delta = 0.4$ (the dash-dot line), and for $\delta = 1$, which corresponds to no probability weighting at all (the solid line). The overweighting of low probabilities and the greater sensitivity to changes in probability at higher probability levels are both clearly visible for $\delta < 1$.

In our subsequent analysis, we work with the specification of cumulative prospect theory laid out in equations (4)-(5) and (6)-(7), adjusted only to allow for continuous probability distributions.

3 Investor Preferences

In Sections 4 and 5, we study security prices in economies where investors evaluate risk using cumulative prospect theory, paying particular attention to the implications of the probability weighting function. In this section, we lay the groundwork for that analysis by specifying investor preferences in more detail.

Suppose that an investor uses cumulative prospect theory to evaluate risk, and that his beginning-of-period wealth and end-of-period wealth are W_0 and $\widetilde{W} = W_0\widetilde{R}$, respectively. In prospect theory, utility is defined over gains and losses, which we interpret as final wealth $\widetilde{W}$ minus a reference wealth level W_z. In symbols, the gain or loss in wealth, $\widehat{W}$, is

$$\widehat{W} = \widetilde{W} - W_z. \tag{8}$$

One possible reference level is initial wealth W_0. In this paper, we use another reference level, namely W_0R_f, where R_f is the gross risk-free rate, so that

$$\widehat{W} = \widetilde{W} - W_0R_f. \tag{9}$$

This specification is more tractable, and potentially more plausible: the agent thinks of the change in his wealth as a gain only if it exceeds what he would have achieved by investing at the risk-free rate. We also assume:

Assumption 1: $|E(\hat{W})| < \infty$, and $\mathrm{Var}(\hat{W}) < \infty$.

In the economies we study later, each investor has the goal function:

$$U(\widetilde{W}) \equiv V(\hat{W}) = V(\hat{W}^+) + V(\hat{W}^-), \tag{10}$$

where $\hat{W}^+ = \max(\widehat{W}, 0)$, $\hat{W}^- = \min(\widehat{W}, 0)$, and

$$\begin{aligned} V(\hat{W}^+) &= -\int_0^\infty v(W)\, dw^+(1 - P(W)) & (11)\\ V(\hat{W}^-) &= \int_{-\infty}^0 v(W)\, dw^-(P(W)), & (12) \end{aligned}$$

and where $P(\cdot)$ is the cumulative probability distribution function of $\widehat{W}$. Equations (10)-(12) are equivalent to equations (4)-(5), modified to allow for continuous probability distributions. As before, $w^+(\cdot)$ and $w^-(\cdot)$ are the probability weighting functions for gains and losses, respectively. We assume:

Assumption 2: $w^+(\cdot) = w^-(\cdot) \equiv w(\cdot)$.

Assumption 3: $w(\cdot)$ takes the form proposed by Tversky and Kahneman (1992),

$$w(P) = \frac{P^\delta}{(P^\delta + (1-P)^\delta)^{1/\delta}}, \tag{13}$$

where $\delta \in (0, 1)$. As mentioned above, experimental evidence suggests $\delta \approx 0.65$.

Assumption 4: $v(\cdot)$ takes the form proposed by Tversky and Kahneman (1992),

$$v(x) = \begin{cases} x^\alpha & \text{for } x \geq 0 \\ -\lambda(-x)^\alpha & \text{for } x < 0, \end{cases} \tag{14}$$

where $\lambda > 1$ and $\alpha \in (0, 1)$. As noted earlier, experimental evidence suggests $\alpha \approx 0.88$ and $\lambda \approx 2.25$.

Assumption 5: $\alpha < 2\delta$.

Taken together with Assumptions 1-4, Assumption 5 ensures that the goal function $V(\cdot)$ in (10) is well-behaved at $\pm\infty$, and therefore well-defined. The values of α and δ estimated by Tversky and Kahneman (1992) satisfy this condition. In the case of Normal or Lognormal distributions, Assumption 5 is not needed.

Before embarking on our analysis, we present a useful lemma. In informal terms, the lemma shows that we can reverse the order of $v(\cdot)$ and $w(\cdot)$ in equations (11)-(12).

Lemma 1: *Under Assumptions 1-5,*

$$\begin{aligned} V(\hat{W}^+) &= \int_0^\infty w(1 - P(x))dv(x) & (15)\\ V(\hat{W}^-) &= -\int_{-\infty}^0 w(P(x))dv(x). & (16) \end{aligned}$$

Proof of Lemma 1: See the Appendix for a full derivation. In brief, the lemma follows by applying integration by parts to equations (11) and (12).

4 The Pricing of Normally Distributed Securities

We now study the implications of cumulative prospect theory for asset prices. We first show that, in a one-period equilibrium model with Normally distributed security payoffs, the CAPM can hold. Under the assumption of Normality, then, the pricing implications of cumulative prospect theory are no different from those of expected utility.

Our proof of the CAPM builds on the following two propositions, which show that cumulative prospect theory preferences satisfy first-order stochastic dominance and, within certain classes of distributions, second-order stochastic dominance as well. That they satisfy first-order stochastic dominance is not surprising: Tversky and Kahneman (1992) themselves point this out. The result that, under certain conditions, they can also satisfy second-order stochastic dominance *is* surprising and is new to the literature.

Proposition 1: *Under Assumptions 1-5, the preferences in (10)–(12) satisfy the first-order stochastic dominance property. That is, if $\hat{W}_1$ first-order stochastically dominates $\hat{W}_2$, then $V(\hat{W}_1) \geq V(\hat{W}_2)$. Moreover, if $\hat{W}_1$ strictly first-order stochastically dominates $\hat{W}_2$, then $V(\hat{W}_1) > V(\hat{W}_2)$.*

Proof of Proposition 1: Since $\hat{W}_1$ first-order stochastically dominates $\hat{W}_2$, $P_1(x) \leq P_2(x)$ for all $x \in \Re$, where $P_i(\cdot)$ is the cumulative distribution function for $\hat{W}_i$. Equations (15) and (16) imply $V(\hat{W}_1^+) \geq V(\hat{W}_2^+)$ and $V(\hat{W}_1^-) \geq V(\hat{W}_2^-)$, and therefore that $V(\hat{W}_1) \geq V(\hat{W}_2)$. If, moreover, $\hat{W}_1$ strictly first-order stochastically dominates $\hat{W}_2$, then $P_1(x) < P_2(x)$ for some $x \in \Re$. Given that cumulative distribution functions are right continuous, we have $V(\hat{W}_1) > V(\hat{W}_2)$.[2]

Proposition 2: *Suppose that Assumptions 1-5 hold. Take two distributions, $\hat{W}_1$ and $\hat{W}_2$, and suppose that:*

(i) $E(\hat{W}_1) = E(\hat{W}_2) \geq 0$

(ii) $\hat{W}_1$ *and* $\hat{W}_2$ *are both symmetrically distributed*

(iii) $\hat{W}_1$ *and* $\hat{W}_2$ *satisfy a single-crossing property, so that if* $P_i(\cdot)$ *is the cumulative*

[2]Not all of Assumptions 1-5 are needed for this result. Assumptions 2-4 can be replaced by "$w^+(\cdot)$, $w^-(\cdot)$, and $v(\cdot)$ are strictly increasing and continuous," and Assumption 5 can be replaced by "the integrals in equations (11) and (12) are well-defined and finite."

distribution function for $\hat{W}_i$, there exists z such that $P_1(x) \leq P_2(x)$ for $x < z$ and $P_1(x) \geq P_2(x)$ for $x > z$.

Then, $V(\hat{W}_1) \geq V(\hat{W}_2)$. If, furthermore, the inequalities in condition (iii) hold strictly for some x, then $V(\hat{W}_1) > V(\hat{W}_2)$.

Proof of Proposition 2: See the Appendix.

Proposition 2 immediately implies that, for certain classes of distributions – specifically, for any set of symmetric distributions that have the same non-negative mean and that, pairwise, satisfy a single-crossing property – cumulative prospect theory preferences satisfy second-order stochastic dominance. To see this, take any two distributions in the set, $\widehat{W}_1$ and $\widehat{W}_2$, say. The single-crossing property means that we can rank $\widehat{W}_1$ and $\widehat{W}_2$ according to the second-order stochastic dominance criterion: we can obtain one distribution from the other by adding a mean-preserving spread. Proposition 2 then shows that the distribution which dominates is preferred by a cumulative prospect theory investor.

Given that the probability weighting function $w(\cdot)$ and the convexity of the value function $v(\cdot)$ in the region of losses induce risk-seeking, cumulative prospect theory preferences do not, in general, satisfy second-order stochastic dominance: an agent with these preferences is not necessarily averse to a mean-preserving spread. The intuition for why, within certain classes of distributions, he *is* averse to such a spread, is discussed in detail in our proof of Proposition 2. In brief, the idea is that, within these classes, a mean-preserving spread fattens the right tail of the wealth distribution – an attractive feature for a cumulative prospect theory investor – but also fattens the *left* tail of the distribution, which is unattractive. Since the investor is loss averse, he is more sensitive to changes in the left tail, and so, on balance, is averse to the mean-preserving spread.

We now use Propositions 1 and 2 to derive a CAPM result. We make the following assumptions:

Assumption 6: We study a one-period economy with two dates, $t = 0$ and $t = 1$.

Assumption 7: Asset supply. The economy contains a risk-free asset, which is in perfectly elastic supply, and has a gross return of R_f. There are also J risky assets. Risky asset j has $n_j > 0$ shares outstanding, a per-share payoff of $\tilde{X}_j$ at time 1, and a gross return of $\widetilde{R}_j$. The random payoffs $\{\tilde{X}_1, \cdots, \tilde{X}_J\}$ have a positive-definite variance-covariance matrix, so that no linear combination of the J payoffs is a constant.

Assumption 8: Distribution of payoffs. The joint distribution of the time 1 payoffs on the J risky assets is multivariate Normal.

Assumption 9: Investor preferences. The economy contains a large number of price-taking investors who derive utility from the time 1 gain or loss in wealth, $\widehat{W}$, defined in (9). All investors have the *same* preferences, namely those described in equations (10)–(12) and Assumptions 2-5. In particular, the parameters α, δ, and λ are the same for all investors.

Assumption 10: Investor beliefs. All investors assign the same probability distribution to future payoffs and security returns.

Assumption 11: Investor endowments. Each investor is endowed with positive net wealth in the form of traded securities.

Assumption 12: There are no trading frictions or constraints.

We can now prove:

Proposition 3: *Under Assumptions 6-12, there exists an equilibrium in which the CAPM holds, so that*

$$E(\tilde{R}_j) - R_f = \beta_j(E(\tilde{R}_M) - R_f), \quad j = 1, \cdots, J, \tag{17}$$

where

$$\beta_j \equiv \frac{\mathrm{Cov}(\tilde{R}_j, \tilde{R}_M)}{\mathrm{Var}(\tilde{R}_M)}, \tag{18}$$

and where $\widetilde{R}_M$ is the market return. The excess market return, $\hat{R}_M \equiv \tilde{R}_M - R_f$, satisfies

$$V(\hat{R}_M) \equiv -\int_{-\infty}^{0} w(P(\hat{R}_M))dv(\hat{R}_M) + \int_{0}^{\infty} w(1 - P(\hat{R}_M))dv(\hat{R}_M) = 0, \tag{19}$$

and the market risk premium is positive:

$$E(\widehat{R}_M) > 0. \tag{20}$$

Proof of Proposition 3: See the Appendix.

The intuition behind the proposition is straightforward. When security payoffs are Normally distributed, the goal function in (10)-(12) becomes a function of the mean and standard deviation of the distribution of wealth. Since these preferences satisfy first-order stochastic dominance, all investors choose a portfolio on the mean-variance efficient frontier, in other words, a portfolio that combines the risk-free asset and the tangency portfolio. Market clearing means that the tangency portfolio is the market portfolio, and the CAPM then follows in the usual way.

The previous paragraph explains why, if there *is* an equilibrium, that equilibrium must be a CAPM equilibrium. In our proof of Proposition 3, we also show that a CAPM equilibrium satisfying conditions (17), (19), and (20) does indeed exist. It is in this part of the argument that we make use of the second-order stochastic dominance result in Proposition 2.

The result that, under cumulative prospect theory, the CAPM can still hold, also appears in De Giorgi, Hens, and Levy (2003). For two reasons, we include this result here as well. First, it provides a very useful springboard for our main contribution, namely the analysis in Section 5 of how skewed securities are priced. Second, thanks to our new result on second-order stochastic dominance in Proposition 2, we are able to offer a different proof of the CAPM; specifically, one that is much shorter.[3]

5 The Pricing of Skewed Securities

Under the assumption of Normality, then, the pricing implications of cumulative prospect theory are identical to those of expected utility. We now show that, as soon as we relax the assumption of Normality, cumulative prospect theory can have novel pricing predictions. We demonstrate this using the most parsimonious model possible, one with the minimum amount of additional structure. Specifically, we study an economy in which Assumptions 6-12 still apply, but which, in addition to the risk-free asset and the J Normally distributed risky assets, also contains a positively skewed security. We make the following simplifying assumptions:

Assumption 13: Independence. The return on the skewed security is independent of the returns on the J original risky securities.

Assumption 14: Small supply. The payoff of the skewed security is infinitesimal relative to the total payoff of the J original risky securities.[4]

In a representative agent economy with concave, expected utility preferences, a small, independent, skewed security earns an average return infinitesimally above the risk-free rate; in other words, an average *excess* return infinitesimally above zero. We now show that, when investors have the cumulative prospect theory preferences in (10)-(12), we obtain a very

[3]De Giorgi, Hens, and Levy (2003) also point out that, if investors have cumulative prospect theory preferences with *heterogeneous* preference parameters, some of them may want to take an infinite position in the market portfolio and a CAPM equilibrium may therefore not exist. This problem can be avoided by imposing the condition that each investor's terminal wealth be non-negative; by adding a second term to investors' utility, namely a concave utility of consumption term; or, as De Giorgi, Hens, and Levy (2003) themselves suggest, by slightly modifying Tversky and Kahneman's (1992) specification.

[4]We assume an "infinitesimal" payoff for technical convenience. In practice, the payoff of the skewed security simply needs to be small, relative to the total payoff of the J original risky securities. Just how small it needs to be will become clearer when we present quantitative examples of equilibria.

different prediction: a small, independent, skewed security earns a *negative* average excess return. For simplicity, our initial analysis imposes short-sale constraints. Our results are not driven by these constraints, however: in Section 5.7, we show that our main conclusions are valid even when investors can sell short.

We first note that, in any equilibrium, all investors must hold portfolios that are some combination of the risk-free asset, the skewed security, and the tangency portfolio T formed in the mean / standard deviation plane from the J original risky assets. To see this, suppose that an investor allocates a fraction $1-\theta$ of his wealth to a portfolio P which is some combination of the risk-free asset and the J original risky assets, and a fraction θ of his wealth to the skewed security. If the gross returns of portfolio P and of the skewed security are $\widetilde{R}_p$ and $\widetilde{R}_n$, respectively, the expected return E and variance V of the overall allocation strategy are

$$\begin{aligned} E &= (1-\theta)E(\widetilde{R}_p) + \theta E(\widetilde{R}_n) & (21) \\ V &= (1-\theta)^2 \mathrm{Var}(\widetilde{R}_p) + \theta^2 \mathrm{Var}(\widetilde{R}_n). & (22) \end{aligned}$$

Now, recall that cumulative prospect theory satisfies first-order stochastic dominance. The investor is therefore interested in portfolios which, for given variance in (22), maximize expected return in (21). For a fixed position in the skewed security, these are portfolios that maximize $E(\widetilde{R}_p)$ for given $\mathrm{Var}(\widetilde{R}_p)$, in other words, as claimed above, portfolios that combine the risk-free asset with the tangency portfolio T formed in the mean / standard deviation plane from the J original risky assets. Market clearing further implies that the tangency portfolio T must be the market portfolio formed from the J original risky assets alone, excluding the skewed asset. If we call the latter portfolio the "J-market portfolio," for short, we conclude that all investors hold portfolios that are some combination of the risk-free asset, the J-market portfolio, and the skewed security.

The simplest candidate equilibrium is a homogeneous holdings equilibrium: an equilibrium in which all investors hold the same portfolio. In Section 5.1, however, we show that, for a wide range of parameter values, no such equilibrium exists. We therefore consider the next simplest candidate equilibrium: a heterogeneous holdings equilibrium with two groups of investors, where all investors in the same group hold the same portfolio. Specifically, we conjecture an equilibrium with the following structure: all investors in the first group hold a portfolio that combines the risk-free asset and the J-market portfolio, but takes no position at all in the skewed security; and all investors in the second group hold a portfolio that combines the risk-free asset, the J-market portfolio, *and* a long position in the skewed security.

The heterogeneous holdings in our conjectured equilibrium do not stem from heterogeneous preferences: as specified in Assumption 9, all investors have identical preferences. Rather, they stem from the existence of non-unique optimal portfolios. By assigning each investor to one of the two proposed optimal portfolios, we can clear the market in the skewed security, even though that security is in small supply.

Let $\widehat{R}_M$ and $\widehat{R}_n \equiv \widetilde{R}_n - R_f$ be the excess returns of the J-market portfolio and of the skewed security, respectively. The conditions for our proposed equilibrium are then:

$$V(\widehat{R}_M) = V(\widehat{R}_M + x^*\widehat{R}_n) = 0 \tag{23}$$

$$V(\widehat{R}_M + x\widehat{R}_n) < 0 \text{ for } 0 < x \neq x^*, \tag{24}$$

where

$$V(\hat{R}_M + x\hat{R}_n) = -\int_{-\infty}^{0} w(P_x(R))dv(R) + \int_{0}^{\infty} w(1 - P_x(R))dv(R) \tag{25}$$

and

$$P_x(R) = \Pr(\widehat{R}_M + x\widehat{R}_n \leq R). \tag{26}$$

Here, x^* is the fraction of wealth allocated to the skewed security relative to the fraction allocated to the J-market portfolio, for those investors who allocate a positive amount to the skewed security.

Why are these the appropriate equilibrium conditions? First, recall that, in the conjectured equilibrium, each investor in the first group holds a portfolio with return $(1-\theta)R_f + \theta\widetilde{R}_M$, with $\theta > 0$. Since

$$U(W_0((1-\theta)R_f + \theta\widetilde{R}_M)) = V(W_0\theta\widehat{R}_M) = W_0^\alpha\theta^\alpha V(\widehat{R}_M), \tag{27}$$

an investor will only choose a finite and positive θ if $V(\widehat{R}_M) = 0$. Each investor in the second group holds a portfolio with return $(1-\phi_1-\phi_2)R_f + \phi_1\widetilde{R}_M + \phi_2\widetilde{R}_n$, with $\phi_1 > 0$ and $\phi_2 > 0$. If this portfolio is to be a second global optimum, it must also offer a utility level of zero, so that, if $x^* = \phi_2/\phi_1$,

$$U(W_0((1-\phi_1-\phi_2)R_f + \phi_1\widetilde{R}_M + \phi_2\widetilde{R}_n)) = W_0^\alpha\phi_1^\alpha V(\widehat{R}_M + x^*\widehat{R}_n) = 0. \tag{28}$$

This explains condition (23). Condition (24) ensures that these two portfolios are the *only* global optima.

In general, when a new security is introduced into an economy, the prices of existing securities are affected. A useful feature of our conjectured equilibrium, which we derive formally in the Appendix, is that the prices of the J original risky assets are *not* affected by the introduction of the skewed security: their prices in the heterogeneous holdings equilibrium are identical to what they were in the economy of Section 4, where there was no skewed security.

5.1 An example

We now show that an equilibrium satisfying conditions (23)-(24) actually exists. To do this, we construct an explicit example.

From Assumption 8, the J-market return – the return on the market portfolio excluding the skewed security – is Normally distributed:

$$\widehat{R}_M \sim N(\mu_M, \sigma_M). \tag{29}$$

We model the skewed security in the simplest possible way, using a binomial distribution: with some low probability q, the security pays out a large "jackpot" L, and with probability $1 - q$, it pays out nothing. Using our earlier notation, the payoff distribution is therefore

$$(L, q; 0, 1 - q). \tag{30}$$

For *very* large L and *very* low q, this resembles the payoff distribution of a lottery ticket. If the price of this security is p_n, its gross return $\widetilde{R}_n$ and excess return $\widehat{R}_n = \widetilde{R}_n - R_f$ are distributed as:

$$\widetilde{R}_n \sim (\frac{L}{p_n}, q; 0, 1 - q) \tag{31}$$

$$\widehat{R}_n \sim (\frac{L}{p_n} - R_f, q; -R_f, 1 - q). \tag{32}$$

We now specify the preference parameters $(\alpha, \delta\,, \lambda)$, the skewed security payoff parameters (L, q), the risk-free rate R_f, and the standard deviation of the J-market return σ_M, and search for a mean excess return on the J-market, μ_M, and a price p_n for the skewed security, such that conditions (23)-(24) hold. Specifically, we take the unit of time to be a year and set the annual stock market standard deviation to $\sigma_M = 0.15$ and the annual risk-free rate to $R_f = 1.02$. We set $L = 10$ and $q = 0.09$, which imply substantial positive skewness in the new security's payoff. Finally, we set

$$\alpha = 0.88,\ \delta = 0.65,\ \lambda = 2.25,$$

the values estimated by Tversky and Kahneman (1992).

For these parameter values, the condition $V(\widehat{R}_M) = 0$ in (23) implies $\mu_M = 7.5\%$. This is consistent with Benartzi and Thaler (1995) and Barberis, Huang, and Santos (2001), who show that, in an economy with prospect theory investors who derive utility from annual fluctuations in the value of their stock market holdings, the equity premium can be very substantial. The intuition is that, under prospect theory, investors are much more sensitive to stock market losses than to stock market gains. They therefore perceive the stock market to be very risky, and charge a high average return as compensation.

We now search for a price p_n of the skewed security such that conditions (23)-(24) hold. To do this, we need to compute $P_x(R)$, defined in (26). Given our assumptions about the distribution of security returns, it is given by

$$P_x(R) = \Pr(\widehat{R}_M + x\widehat{R}_n \le R)$$

$$= \Pr(\widehat{R}_n = \frac{L}{p_n} - R_f)\Pr(\widehat{R}_M \leq R - x(\frac{L}{p_n} - R_f)) + \Pr(\widehat{R}_n = -R_f)\Pr(\widehat{R}_M \leq R + xR_f)$$

$$= qN(\frac{R - x(\frac{L}{p_n} - R_f) - \mu_M}{\sigma_M}) + (1-q)N(\frac{R + xR_f - \mu_M}{\sigma_M}), \tag{33}$$

where $N(\cdot)$ is the cumulative Normal distribution.

We find that the price level $p_n = 0.925$ satisfies conditions (23)-(24). Figure 3 provides a graphical illustration. For this value of p_n, the solid line in the figure plots the goal function $V(\widehat{R}_M + x\widehat{R}_n)$ for a range of values of x, where x is the amount allocated to the skewed security relative to the amount allocated to the J-market portfolio. The two global optima are clearly visible: one at $x = 0$ and one at $x = 0.086$. So long as the skewed security is in small supply – specifically, so long as its value is less than 8.6% of the value of the J-market portfolio – we can clear the market for it by assigning each investor to one of the two global optima. Given the return distribution in (32), the equilibrium average excess return on the skewed security is

$$E(\hat{R}_n) = \frac{qL}{p_n} - R_f = \frac{(0.09)(10)}{0.925} - 1.02 = -0.047, \tag{34}$$

so that the average net return is $E(\widetilde{R}_n) - 1 = E(\widehat{R}_n) + R_f - 1 = -0.027$.

The shape of the solid line can be understood as follows. Adding a small position in the skewed security to an existing position in the J-market portfolio initially lowers utility because of the security's negative average excess return and because of the lack of diversification the strategy entails. As we increase x further, however, the security starts to add skewness to the return on the investor's portfolio. Since the investor overweights the tails of his wealth distribution, he values this highly and his utility increases. At a price level of $p_n = 0.925$, the skewness effect offsets the diversification and negative excess return effects in a way that produces two global optima at $x = 0$ and $x = 0.086$. As x increases beyond 0.086, utility falls again: at this point, a higher value of x preserves the lottery-like structure of the investor's wealth but increases the size of the lottery jackpot. Since the prospect theory value function is concave over gains, the benefit of a larger jackpot is too small to compensate for the lack of diversification, and utility falls.

Figure 3 also explains *why* the skewed security earns a negative average excess return. By taking a large position in this security, some investors can add skewness to their portfolio return; they value this highly, and are therefore willing to accept a low average return on the security.

In summary, we have shown that, under cumulative prospect theory, a positively skewed security can become overpriced, relative to its price in a concave expected utility model, and can earn a low average return. We emphasize that this result is by no means an obvious one. An investor who overweights the tails of a portfolio return distribution will, of course, value a positively skewed *portfolio* highly; what is surprising is that he also values a skewed *security* highly, even if that security is small and independent of other risks.

It is natural to ask whether the parameter values in our example, namely $(\sigma_M, L, q) = (0.15, 10, 0.09)$, admit any equilibria other than the heterogeneous holdings equilibrium described above. While it is difficult to give a definitive answer, we can at least show that, for these parameter values, there is no *homogeneous* holdings equilibrium, in other words, no equilibrium in which all investors hold the same portfolio.

In any homogeneous holdings equilibrium, each investor would need to be happy to hold an infinitesimal amount ε^* of the skewed security. The equilibrium conditions are therefore

$$V(\widehat{R}_M + \varepsilon^* \widehat{R}_n) = 0 \tag{35}$$

$$V(\widehat{R}_M + \varepsilon \widehat{R}_n) < 0, \qquad 0 \leq \varepsilon \neq \varepsilon^*. \tag{36}$$

Using the same reasoning as for condition (23), we need condition (35) to ensure that investors will optimally choose positive but *finite* allocations to the J-market portfolio and the skewed security. Condition (36) ensures that an allocation ε^* to the skewed security is a global optimum. Since this global optimum is also a local optimum, a necessary condition for equilibrium is $V'(\widehat{R}_M + \varepsilon^* \widehat{R}_n) = 0$.

If a homogeneous holdings equilibrium exists, we can approximate it by studying the limiting case as $\varepsilon^* \to 0$. We therefore search for a price p_n of the skewed security such that $V(\widehat{R}_M) = 0$ and $V'(\widehat{R}_M) = 0$. We find that these conditions are satisfied for $p_n = 0.882$. The dashed line in Figure 3 plots the goal function $V(\widehat{R}_M + x\widehat{R}_n)$ for this case. We immediately see that $p_n = 0.882$ does *not* represent an equilibrium, as it violates condition (36): all investors would prefer a substantial positive position in the skewed security to an infinitesimal one, making it impossible to clear the market. There is therefore no homogeneous holdings equilibrium for these preference and payoff parameters.[5]

5.2 How does expected return vary with skewness?

The skewness of the new security's excess return in (32) is primarily determined by q, the probability of the large payoff: a low value of q corresponds to a high degree of skewness. In this section, we examine how the equilibrium average excess return on this security changes as we vary its skewness; or, more precisely, as we vary q.[6]

[5]Given the scale of Figure 3, it is hard to tell whether the dashed line really does have a derivative of 0 at $x = 0$. Magnifying the left side of the graph confirms that the derivative *is* 0 at $x = 0$, although it quickly turns negative as x increases.

[6]It is straightforward to check that the skewness of the excess return in (32) is $(L/p_n)(1 - 2q)$. We can approximate the price of the new security by $p_n \approx qL/R_f$, its price in a representative agent economy with

Our main finding, obtained by searching across many different values of q, is that the results in Section 5.1 for the case of $q = 0.09$ are typical of those for all low values of q. Specifically, for all $q \leq 0.105$, a heterogeneous holdings equilibrium can be constructed, but a homogeneous holdings equilibrium cannot. For q in this range, the expected excess return on the new security is negative, and becomes more negative as q falls. The intuition is that, when q is low, the skewed security is highly skewed and can add a large amount of skewness to the investor's portfolio; as a result, it is more valuable to him and he lowers the expected return he requires on it.

For q above 0.105, however – in other words, for a skewed security that is only mildly skewed – the opposite is true: a homogeneous holdings equilibrium can be constructed, but a heterogeneous holdings equilibrium cannot. The reason the heterogeneous holdings equilibrium breaks down for higher values of q is that, if the new security is not sufficiently skewed, no position in it, however large, adds enough skewness to the investor's portfolio to compensate for the lack of diversification the position entails.

To see this last point, suppose that, as before, $\sigma_M = 0.15$ and $L = 10$, so that, once again, $\mu_M = 0.075$, but that q is set to 0.2 rather than to 0.09. Figure 4 plots the goal function $V(\widehat{R}_M + x\widehat{R}_n)$ for various values of p_n, namely $p_n = 2.5$ (dashed line), $p_n = 1.96$ (solid line), and $p_n = 1.35$ (dash-dot line). While these lines correspond to only three values of p_n, they hint at the outcome of a more extensive search, namely that *no* value of p_n can deliver two global optima. In other words, no value of p_n can satisfy conditions (23)-(24) for a heterogeneous holdings equilibrium.

For the parameter values $(\sigma_M, L, q) = (0.15, 10, 0.2)$, we can only obtain a homogeneous holdings equilibrium, one which satisfies conditions (35)-(36). As before, we study this equilibrium in the limiting case of $\varepsilon^* \to 0$ by searching for a price p_n of the skewed security that satisfies $V(\widehat{R}_M) = 0$ and $V'(\widehat{R}_M) = 0$. We find that $p_n = 1.96$ satisfies these conditions. The solid line in Figure 4 plots the goal function for this value of p_n. The graph shows that $x = \varepsilon^*$ is not only a local optimum, but also a global optimum. We have therefore identified a homogeneous holdings equilibrium. In this equilibrium, the expected excess return of the skewed security is

$$E(\widehat{R}_n) = \frac{qL}{p_n} - R_f = \frac{(0.2)(10)}{1.96} - 1.02 = 0.$$

It is no coincidence that the skewed security earns an expected excess return of zero. The following proposition shows that, whenever a homogeneous holdings equilibrium exists, the expected excess return on the skewed security is *always* zero, or, more precisely, infinitesimally greater than zero. In other words, in a homogeneous holdings setting, cumulative prospect theory assigns the skewed security the *same* average return that a concave

concave expected utility, where it earns an average excess return of zero. A rough estimate of the skewness of the new security is therefore $R_f(1/q - 2)$, so that skewness is primarily determined by q.

expected utility specification would.

Proposition 4: *Consider an agent with the preferences in (10)-(12) who holds a portfolio with return* $\tilde{R} \equiv \hat{R} + R_f$. *Suppose that he adds a small amount of an independent security with excess return* $\widehat{R}_n$ *to his portfolio; and that he finances this by borrowing, so that his excess portfolio return becomes* $\hat{R}+\varepsilon\hat{R}_n$. *If* $\hat{R}$ *has a probability density function that satisfies* $\sigma(\hat{R}) > 0$, *then*

$$\lim_{\varepsilon \to 0} \frac{V(\hat{R} + \varepsilon \hat{R}_n) - V(\hat{R})}{\varepsilon} = E(\hat{R}_n) V'(\hat{R}), \tag{37}$$

where, with some abuse of notation, $V'(\cdot)$ *is defined as*

$$\begin{aligned} & V'(\hat{R}) \\ \equiv\ & \lim_{x \to 0} \frac{V(\hat{R}+x) - V(\hat{R})}{x} \\ =\ & \int_{-\infty}^{0} w'(P(R)) P'(R) dv(R) + \int_{0}^{\infty} w'(1 - P(R)) P'(R) dv(R) > 0, \end{aligned} \tag{38}$$

with $P(R) \equiv \mathrm{Prob}(\hat{R} \leq R)$.

Proof of Proposition 4: See the Appendix.

In any homogeneous holdings equilibrium, we need $V(\widehat{R}_M + x\widehat{R}_n)$ to have a local optimum at $x = \varepsilon^*$, for infinitesimal ε^*. From the proposition, this implies that $E(\widehat{R}_n) \approx 0$.

Figure 5 summarizes the findings of this section by plotting the expected return of the skewed security as a function of q when $(\sigma_M, L) = (0.15, 10)$. For $q \leq 0.105$, we obtain heterogeneous holdings equilibria in which the expected excess return is negative and falls as q falls. For $q > 0.105$, a heterogenous holdings equilibrium can no longer be constructed, but a homogeneous holdings equilibrium can, and here, the skewed security earns an expected excess return of zero.

Figure 5 emphasizes that, while cumulative prospect theory predicts a relationship between a security's skewness and its average return, the predicted relationship is highly nonlinear. Only securities with a high degree of skewness earn a negative expected excess return. Those with merely moderate skewness have an expected excess return of zero.

5.3 Relation to other research on skewness

Our analysis of economies with cumulative prospect theory investors has led us to a prediction that is new to the asset pricing literature, namely that idiosyncratic skewness is priced. Our main motivation for working with cumulative prospect theory is that, given its status as a leading model of how investors evaluate risk, it is interesting to document its implications for security pricing. At the same time, it is reasonable to ask whether the pricing of idiosyncratic skewness can also be derived in an expected utility framework with skewness-loving preferences.

In fact, models in which investors have expected utility preferences with concave, skewness-loving utility functions do *not* predict the pricing of idiosyncratic skewness. In such economies, only a security's coskewness with the market portfolio is priced; the security's own skewness is not (Kraus and Litzenberger, 1976; Harvey and Siddique, 2000). A small, independent, skewed security therefore earns a zero risk premium: its coskewness with the market is zero. It does not earn the negative risk premium we observe under cumulative prospect theory.

One way to think about this point is to note that, in our model, the pricing of idiosyncratic skewness traces back to the undiversified positions some investors hold in the skewed security. By contrast, investors with concave, skewness-loving, expected utility preferences always hold diversified portfolios. As a result, only coskewness with the market is priced; idiosyncratic skewness is not.

Can idiosyncratic skewness be priced when investors have expected utility preferences with *convex*, skewness-loving utility functions, such as cubic utility functions? It is hard to give a definitive answer, because the pricing implications of these preferences have not yet been fully analyzed. One well-known difficulty with such preferences, however, is that, given a skewed security as an investment option, the optimal portfolio may involve an infinite position in that security, a phenomenon known as "plunging" (Kane, 1982; Polkovnichenko, 2005).[7]

One framework that does predict the pricing of idiosyncratic skewness is the optimal expectations model of Brunnermeier and Parker (2005), in which investors choose their beliefs in order to maximize the discounted value of expected future utility flows. Brunnermeier, Gollier, and Parker (2007) show that, in this framework, all investors allocate a significant fraction of their wealth to positively skewed assets, which, in equilibrium, earn low average returns.

[7] In order to provide a theoretical framework for some empirical results on skewness, Mitton and Vorkink (2007) consider an expected utility model in which some investors have convex, skewness-loving preferences. A potential pitfall of this model, however, is that the global optimum for these investors may indeed involve an infinite position in the skewed security; the authors focus on a finite local optimum, but do not prove that it is also a global optimum. We suggest later that their empirical results may be more easily interpreted using the cumulative prospect theory model we present here.

5.4 How does expected return vary with the preference parameters?

In Section 5.2, we saw how the expected return on the skewed security changes as we vary the probability q of winning the large payoff. Throughout that analysis, we kept the preference parameters fixed at the values estimated by Tversky and Kahneman (1992), namely $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$. In this section, we fix q, and examine the effect of varying α, δ, and λ. The three panels in Figure 6 plot the expected excess return of the skewed security against each of these parameters in turn, holding the other two constant.

The top left panel shows that, as δ increases, the expected return of the skewed security also rises. A low value of δ means that the investor weights the tails of a probability distribution particularly heavily and therefore that he is strongly interested in a positively skewed portfolio. Since the skewed security offers a way of constructing such a portfolio, it is very valuable, and the investor is willing to hold it in exchange for a very low average return. Once δ rises above 0.68, however, no heterogeneous holdings equilibrium is possible: by this point, the investor does not care enough about having a positively skewed portfolio for him to want to take on an undiversified position in the skewed security. Only a homogeneous holdings equilibrium is possible, and, in such an equilibrium, the skewed security earns an average excess return of zero.

The top right panel shows that, as λ increases, the expected return on the skewed security also goes up. The parameter λ governs the investor's aversion to losses. In order to add skewness to his portfolio, the investor needs to hold a large, undiversified position in the skewed security. As λ increases, he finds it harder to tolerate the high volatility of this undiversified position and is therefore only willing to hold the skewed security if it offers a high expected return. Once λ rises above 2.48, no position in the skewed security, however large, contributes sufficient skewness to offset the painful swings in the value of the overall portfolio. In this range, only a homogeneous holdings equilibrium is possible.

Finally, the lower left panel shows that, as α falls, the expected return on the skewed security goes up. A lower α means that the value function in the region of gains, depicted in Figure 1, is more concave. This means that the investor derives less utility from a positively skewed portfolio. The skewed security is therefore less useful to him, and he is only willing to hold it in exchange for a higher average return.

5.5 Additional skewed securities

We now show that our main result – that, under cumulative prospect theory, a positively skewed security can earn a negative average excess return – continues to hold even when we introduce additional skewed securities into the economy.

Suppose that Assumptions 6-12 apply as before, but that, in addition to the risk-free asset and J Normally distributed securities described there, the economy now contains *two* skewed securities, each in infinitesimal supply, independent of other securities, and with the payoff distribution in (30). The excess return and price of the j'th skewed security are $\widehat{R}_{n,j}$ and $p_{n,j}$, respectively.

As before, optimizing investors hold portfolios that combine the risk-free asset, the J-market portfolio, and the skewed securities. Each investor's goal function is therefore

$$V(\widehat{R}_M + x_1\widehat{R}_{n,1} + x_2\widehat{R}_{n,2}) = -\int_{-\infty}^{0} w(P_{x_1,x_2}(R))dv(R) + \int_{0}^{\infty} w(1 - P_{x_1,x_2}(R))dv(R) \quad (39)$$

where

$$P_{x_1,x_2}(R) = \Pr(\widehat{R}_M + x_1\widehat{R}_{n,1} + x_2\widehat{R}_{n,2} \leq R), \quad (40)$$

and where x_j is the amount allocated to skewed security j relative to the amount allocated to the J-market portfolio.

Now that there are two skewed securities, we conjecture an equilibrium with *three* global optima: a portfolio that combines the risk-free asset and the J-market portfolio with a large, undiversified position $x^* > 0$ in just the first skewed security; a portfolio that combines the risk-free asset and the J-market portfolio with a large, undiversified position x^* in just the second skewed security; and a portfolio that holds only the risk-free asset and the J-market portfolio and takes no position at all in either of the skewed securities. In other words, the three conjectured optima are

$$(x_1, x_2) = (x_*, 0), \ (0, x_*), \text{ and } (0, 0). \quad (41)$$

By assigning each investor to one of these three optima, we can clear markets in all securities.

Earlier, we saw that if $(x_1, x_2) = (0, 0)$ is an optimum – in other words, if combining the risk-free asset with a positive, finite position in the J-market portfolio is an optimum – we need $V(\widehat{R}_M) = 0$. The goal function must therefore take the value 0 at all three optima. This leads to the equilibrium conditions:

$$V(\widehat{R}_M) = V(\widehat{R}_M + x^*\widehat{R}_{n,j}) = 0, \quad j = 1, 2 \quad (42)$$

$$V(\widehat{R}_M + x_1\widehat{R}_{n,1} + x_2\widehat{R}_{n,2}) < 0, \quad 0 < x_1 \neq x^*, \ 0 < x_2 \neq x^*. \quad (43)$$

We now check that this conjectured equilibrium exists. Suppose that, as in the example in Section 5.1, $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$, $(\sigma_M, R_f) = (0.15, 1.02)$, and that, for each of the two skewed securities, $(L, q) = (10, 0.09)$. The condition $V(\widehat{R}_M) = 0$ again implies $\mu_M = 0.075$. All that remains is to find $p_{n,1}$ and $p_{n,2}$ that satisfy conditions (42)-(43).

We find that the values $p_{n,1} = p_{n,2} = 0.925$ satisfy these conditions with $x^* = 0.086$. In other words, for these prices, the goal function has three global optima, namely

$$(x_1, x_2) = (0.086, 0), (0, 0.086), \text{ and } (0, 0).$$

We note that the two skewed securities have the same price as the skewed security in the original one-skewed-security economy of Section 5.1. Their average excess return is therefore also the same, namely -4.7%.

An important feature of this heterogeneous holdings equilibrium is that investors prefer a large, undiversified position in just one skewed security to a diversified position in two of them. The intuition is that, by diversifying, an investor lowers the volatility of his overall portfolio – which is good – but also lowers its skewness – which is bad. Which of the two forces dominates depends on the distribution of security returns. For the binomial distribution in (31), skewness falls faster than volatility as the investor diversifies; diversification is therefore unattractive.

Our goal here is a modest one. We simply show that, when there are two skewed securities, we can construct an equilibrium in which these skewed securities earn a negative average excess return. We do not claim to characterize the full set of preference and payoff parameters for which this is possible; nor do we claim that this equilibrium is unique. At the same time, we have explored many sets of preference and payoff parameters, and have found that, whenever these parameters support a heterogenous holdings equilibrium in the case of one skewed security, a heterogeneous holdings equilibrium can also be constructed in the case of two skewed securities, and the average excess return of the skewed securities in this case is the same as in the original one-skewed-security economy.

What happens when there are more than two skewed securities? Suppose that there are N skewed securities in the economy, each in small supply, independent of other securities, and with the payoff distribution in (30). In this case, we conjecture an equilibrium with $N+1$ optima: a portfolio that combines the risk-free asset and the J-market portfolio with a large, undiversified position x^* in just the first skewed security; a portfolio that combines the risk-free asset and the J-market portfolio with a large, undiversified position x^* in just the second skewed security; and so on for each of the skewed securities; and, finally, a portfolio that holds only the risk-free asset and the J-market portfolio and takes no position at all in any of the skewed securities. By assigning each investor to one of these optima, we can clear markets for all securities.

We have analyzed economies with as many skewed securities as computational limits will allow: specifically, economies with $N \leq 10$. For N in this range, we find that we can construct a heterogeneous holdings equilibrium of the form conjectured; for example, the preference and payoff parameters considered above, namely $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$ and $(\sigma_M, L, q) = (0.15, 10, 0.09)$, support such an equilibrium. In this equilibrium, the average excess return on each of the skewed securities is the same as in the original one-skewed-security economy. Even as we add more skewed securities, then, investors show no interest in diversifying across them. Diversification lowers volatility, but also lowers skewness, and the latter effect dominates.

5.6 Relaxing the short-sales constraint

We now show that the novel pricing implications of cumulative prospect theory continue to hold even when short sales are allowed.

We start with the simplest case: that of one skewed security. Suppose that Assumptions 6-12 apply as before, but that, in addition to the risk-free asset and J Normally distributed securities described there, the economy also contains an independent, skewed security with the payoff distribution in (30), and with an excess return and price of $\widehat{R}_n$ and p_n, respectively. For the purposes of this section, we assume that the skewed security is in zero net supply; this will simplify our discussion.

We conjecture that, when shorting is allowed, there is an equilibrium with two optima: a portfolio which combines a position in the risk-free asset and the J-market portfolio with a positive position x^* in the skewed security; and a portfolio which combines a position in the risk-free asset and the J-market portfolio with a short position x^{**} in the skewed security. Since the skewed security is in zero net supply, we can clear markets by assigning an appropriate number of investors to each optimum.

The equilibrium we have just described requires that the goal function $V(\widehat{R}_M + x\widehat{R}_n)$ is maximized at $x = x^*$ and $x = -x^{**}$. As before, $V(\cdot)$ must take the value zero at both optima; otherwise, the investor would prefer to hold only the risk-free asset or to take an infinite position in the risky assets. The equilibrium conditions are therefore:

$$\begin{aligned} V(\widehat{R}_M + x^*\widehat{R}_n) &= V(\widehat{R}_M - x^{**}\widehat{R}_n) = 0 && (44) \\ V(\widehat{R}_M + x\widehat{R}_n) &< 0, \forall x \neq x^*, \ -x^{**}. && (45) \end{aligned}$$

We now check that this equilibrium exists. For the parameters values of Section 5.1, namely $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$, $(\sigma_M, R_f) = (0.15, 1.02)$, and $(L, q) = (10, 0.09)$, we search for μ_M and p_n such that conditions (44)-(45) hold.

We find that $\mu_M = 0.075$ and $p_n = 0.924$ satisfy these conditions. In this equilibrium, the skewed security earns an average excess return of

$$E(\widehat{R}_n) = \frac{qL}{p_n} - R_f = -0.046. \tag{46}$$

Our main result – that, under cumulative prospect theory, a positively skewed security can earn a negative average excess return – therefore continues to hold even when shorting is allowed.

The intuition behind the two optima is straightforward. Some investors take a long position in the skewed security. This gives them a tiny chance of a huge wealth payoff. Since they overweight the tails of the wealth distribution, they find this very attractive, and are willing to hold the skewed security even if it offers a low average return. Other investors take a short position in the skewed security. This exposes them to the possibility of a large drop in wealth. Since they overweight the tails of the wealth distribution, they find this unattractive, but, because they are compensated by a high average return, they are nonetheless willing to take the position. For $p_n = 0.924$, the two strategies are equally attractive. The expected excess return of the skewed security, -0.046, is slightly higher than in the no-short-sales economy of Section 5.1. Shorting allows some investors to exploit the "overpricing" of the skewed security: this exerts some downward pressure on the price, and some upward pressure on the expected return.

Now suppose that the economy contains $N > 1$ skewed securities, each in zero net supply, independent of other securities, and with the payoff distribution in (30). The excess return and price of the j'th skewed security are $\widehat{R}_{n,j}$ and $p_{n,j}$, respectively. What form does the equilibrium take now?

We conjecture a heterogeneous holdings equilibrium with $N+1$ optima: the N portfolios which combine a position in the J-market portfolio with a long position x^* in any one of the N skewed securities and a short position $x^{**}/(N-1)$ in the remaining $N-1$ skewed securities, where $0 < x^{**} < x^*$; and a portfolio which combines a position in the J-market portfolio with a short position x^{***} in each of the N skewed securities. Mathematically, our conjecture is that the goal function

$$V(\widehat{R}_M + x_1\widehat{R}_{n,1} + \ldots + x_M\widehat{R}_{n,M}), \tag{47}$$

where x_j is the allocation to the j'th skewed security relative to the allocation to the J-market portfolio, is maximized at the $N+1$ points

$$(x_1, x_2, \ldots, x_N) = (x_*, -\frac{x_{**}}{N-1}, \ldots, -\frac{x_{**}}{N-1}), \ldots, (-\frac{x_{**}}{N-1}, -\frac{x_{**}}{N-1}, \ldots, x_*) \tag{48}$$

and

$$(x_1, x_2, \ldots, x_M) = (-x_{***}, -x_{***}, \ldots, -x_{***}), \tag{49}$$

where

$$x_*,\ x_{**},\ x_{***} > 0 \text{ and } x_{**} < x_*. \tag{50}$$

Since $x^{**} < x^*$, we can clear markets for all securities by assigning each investor to one of these optima.

Why do we conjecture that the optimal portfolios involve a long position in one skewed security but a short position in many of them? As noted in Section 5.5, investors prefer to go long in just one skewed security because, while diversifying across several long positions reduces volatility, it also reduces desirable positive skewness; since the latter effect dominates, investors prefer that their long positions remain undiversified. On the short side, however, diversification across skewed securities is very valuable: it reduces both volatility and undesirable negative skewness.

Following the usual reasoning, the equilibrium conditions are that the goal function takes the value zero at each of the optima in (48) and (49), and that it takes a value less than zero for all other values of $(x_1, \ldots, x_M)$. For the parameter values of Section 5.1, namely $(\alpha, \delta, \lambda) = (0.88, 0.65, 2.25)$, $(L, q) = (10, 0.09)$, and $(\sigma_M, R_f) = (0.15, 1.02)$, we search for μ_M and $\{p_{n,j}\}_{j=1}^{N}$ that satisfy these conditions.

We have again analyzed economies with as many skewed securities as computational limits will allow: specifically, economies with $N \leq 10$. For N in this range, we find that the conjectured heterogeneous holdings equilibrium does indeed exist. Its properties are described in Table 1. For each N, we report the equilibrium values of μ_M, $E(\widehat{R}_{n,j}) = qL/p_{n,j} - R_f$, and, for comparison, $E(\widehat{R}_n^{NS})$, the average excess return of a skewed security in an economy with N skewed securities when short sales are *not* allowed; as discussed in Section 5.5, this is -0.047 for N in the range $1 \leq N \leq 10$.

The table shows that, for fixed N, the average excess return on a skewed security is higher when short sales are allowed. Intuitively, short sales make it easier to exploit the overpricing of the skewed security, thereby bringing the expected return slightly closer to zero. However, the table also shows that, even when there are ten skewed securities in the economy, each skewed security continues to earn a significantly negative excess return. In other words, even though, in equilibrium, all investors take a short position in many skewed securities, the shorting activity does not remove the overpricing. A strategy that shorts ten positively skewed stocks has significant negative skewness – ten stocks are not enough to diversify the negative skewness away. The strategy is therefore risky. In light of this risk, investors limit the size of their short positions. This, in turn, means that significant overpricing remains.[8]

Of course, if there were *many* skewed securities in the economy – if N were as high as 200, say – the overpricing of the skewed securities would be significantly reduced: a strategy which shorts 200 positively skewed stocks does not exhibit much negative skewness. As a result, investors would be willing to take more substantial short positions, and this would attenuate the overpricing.

[8] The table also shows that, when $N = 10$, the market risk premium μ_M falls slightly when short sales are allowed. The reason is that, when shorting is allowed, all investors take *some* position in the skewed securities. This takes them away from the kink in the value function $v(\cdot)$, making the market portfolio of Normally distributed securities seem slightly less risky, and lowering the market risk premium.

The analysis in this section also allows us to make predictions about the pricing of *negatively* skewed securities. Since a long (short) position in a positively skewed security is equivalent to a short (long) position in a negatively skewed security, the model of this section can be interpreted as saying that, in an economy with negatively skewed securities, these securities will earn a positive average excess return.

5.7 Can arbitrageurs correct the mispricing?

The economies we study in this paper do not contain arbitrage opportunities. Since cumulative prospect theory satisfies first-order stochastic dominance, investors with these preferences would immediately exploit an arbitrage opportunity: in equilibrium, then, there are no such opportunities.

At the same time, it is reasonable to ask how overpriced skewed securities would be in an economy with both cumulative prospect theory investors and more standard, risk averse expected utility agents. It is hard to give a definitive answer, because constructing such a model poses significant technical challenges. However, there is good reason to think that expected utility agents would not fully reverse the overpricing. While they could try to exploit the overpricing by taking a short position in a large number of skewed securities, such a strategy entails significant risks and costs, and these may blunt its impact.

One way to see this is to think about the model of Section 5.6. The cumulative prospect theory investors in that economy are already trying to exploit the overpricing: they are all shorting many of the skewed securities. Interestingly, however, this does not remove the overpricing: unless there are many skewed securities in the economy, the strategy of shorting skewed securities retains significant negative skewness; investors do not, therefore, short very aggressively, and the overpricing remains. Expected utility agents who attempt this strategy face exactly the same problem, and so it is likely that they, too, will fail to remove the overpricing.

Investors who short overpriced securities also face other risks and costs. They may have to pay significant short-selling fees. They run the risk that some of the borrowed securities are recalled before the strategy pays off, as well as the risk that the strategy performs poorly in the short run, triggering an early liquidation. Taken together, these factors suggest that investors may be unwilling to trade very aggressively against the overpricing of skewed securities, thereby allowing it to persist.

An interesting prediction of our model is that, since cumulative prospect theory investors value positively skewed securities so highly, we should see the creation of *new* skewed securities that cater to this demand. For example, there is an incentive to raise some capital and to issue N lottery tickets, each offering a $1/N$ chance of winning the capital. A firm with a subsidiary whose business model has positively skewed fundamentals has an incentive to spin that subsidiary off. And there is an incentive to issue out-of-the-money options.

In practice, of course, we *do* see the creation of new, positively skewed securities. New riskless lotteries do get initiated; and firms do spin off subsidiaries with positively skewed fundamentals – subsidiaries working with cutting-edge technologies, say. One interpretation of this activity is that intermediaries are catering to a preference for skewness, like that captured by cumulative prospect theory.

One concern is that the supply of new skewed securities may be so large as to dampen the premium that investors pay for them. We note, however, that there are important limits to the creation of skewed securities, and this suggests that existing skewed securities will remain overpriced. The creation of riskless lotteries, for example, is limited by federal regulation. The issuance of securities backed by positively skewed fundamentals is limited by the supply of businesses that have such fundamentals. And a market-maker who writes out-of-the-money options faces substantial risk: unless he writes options on a large number of uncorrelated assets, his returns will be strongly negatively skewed.

5.8 Alternative framing assumptions

In our analysis, we assume that agents apply cumulative prospect theory to gains and losses in overall wealth. A simpler way to derive the pricing of idiosyncratic skewness is to assume that agents apply cumulative prospect theory at the level of an individual stock: if agents overweight the tails of an individual stock's return distribution, it is intuitive that a positively skewed security will be overpriced and will earn a negative average excess return. When an agent gets utility directly from the outcome of a specific risk he is facing, even if it is just one of many that determine his overall wealth risk, he is said to exhibit "narrow framing" (Kahneman, 2003; Barberis, Huang, and Thaler, 2006).

In this paper, we retain the traditional assumption of portfolio-level framing in order to show that we do not need strong additional assumptions, such as narrow framing, to draw interesting implications out of cumulative prospect theory. At the same time, a framework which allows for both portfolio-level and stock-level framing might fit the data better. Our current model, for example, predicts that some investors hold a non-trivial fraction of their wealth in one positively skewed security. A model that allows for narrow framing would likely preserve the pricing of idiosyncratic skewness while also predicting a lower allocation to the skewed security. Such a model poses significant technical hurdles, however, and is beyond the scope of our current analysis.

6 Applications

The cumulative prospect theory models of Section 5 predict that a skewed security in small supply will earn a low average return. This result offers a unifying way of thinking about a number of pricing phenomena which, at first sight, may appear unrelated.

Our first application is the low average return on IPO securities (Ritter, 1991).[9] The distribution of IPO returns in the three years after issue is highly positively skewed, probably because IPOs are conducted by young firms, a large fraction of whose value is in the form of growth options. Our analysis therefore implies that, in an economy with cumulative prospect theory investors, IPOs can be overpriced and earn a low average return. By taking a substantial position in an IPO, the investor gives himself a chance, albeit a small chance, of a very large return on wealth. He values this highly, and is willing to hold the IPO even if it offers a low average return. Under cumulative prospect theory, then, the historical performance of IPOs may not be so puzzling.

Figure 5 raises a possible caveat. It shows that investors with cumulative prospect theory preferences overprice *highly* skewed securities, but not those with merely moderate skewness. We therefore need to check that there is *enough* skewness in IPO returns to support a heterogeneous holdings equilibrium and, thereby, to justify a low average return. A full analysis of this issue is beyond the scope of our paper, but, in preliminary calculations, available on request, we find that there *is* enough skewness in actual IPO returns to support a heterogeneous holdings equilibrium.

Our model may also be relevant to the "private equity premium puzzle" documented by Moskowitz and Vissing-Jorgensen (2002): the fact that the return on private business holdings is low, despite the high idiosyncratic risk these holdings entail. In their analysis, the authors find that the returns on private equity are highly positively skewed. Under cumulative prospect theory, then, a low average return is exactly what we would expect to see.

Campbell, Hilscher, and Szilagyi (2006) suggest that our framework may shed light on the average return of "distressed" stocks: stocks of firms with a high predicted probability of bankruptcy. Some theories of distress risk predict that such stocks will earn a high return, on average; but Campbell, Hilscher, and Szilagyi (2006) show that their average return is, in fact, very low. While investigating this puzzle, they also find that distressed stocks have high idiosyncratic skewness. In an economy with cumulative prospect theory investors, then, such stocks should indeed earn a low average return.

[9] Just how low this average return is, is a matter of debate. Most recently, researchers have investigated the extent to which the low average return may reflect "pseudo market timing" (Schultz, 2003; Dahlquist and De Jong, 2004).

Mitton and Vorkink (2006) point out that the pricing of idiosyncratic skewness predicted by our model may be relevant to the diversification discount: the fact that conglomerate firms trade at a discount to a matched portfolio of single segment firms (Lang and Stulz, 1994; Berger and Ofek, 1995). The idea is that investors may pay a premium for single segment firms if the returns of these firms are more positively skewed than those of conglomerates. Mitton and Vorkink (2006) confirm that the returns of single segment firms *are* more positively skewed, and that the diversification discount is most pronounced when the difference between the return skewness of a conglomerate and its matched single segment firms is particularly large.

A related line of reasoning suggests a link between our results and the recently-documented examples of "equity stubs" with remarkably low valuations (Mitchell, Pulvino, and Stafford, 2002; Lamont and Thaler, 2003). These are cases of firms with publicly traded subsidiaries in which the subsidiaries make up a surprisingly large fraction of the value of the parent company; in extreme cases, more than 100 percent of the value of the parent company, so that the equity stub – the claim to the parent company's businesses outside of the subsidiary – has negative value.

Our model cannot explain negative stub values, but it may explain stub values that are surprisingly low, albeit positive. If a subsidiary is valued mainly for its growth options, its returns may be positively skewed, leading investors to overprice it relative to its parent, and thereby generating a low stub value. Consistent with this, in most of the examples listed by Mitchell, Pulvino, and Stafford (2002), the subsidiary's business activities involve newer technologies – and therefore, in all likelihood, more growth options – than do the parent company's.[10]

Our results may also be relevant for understanding option prices. Deep out-of-the-money options have positively skewed returns and so, according to our analysis, may be overpriced. By put-call parity, deep *in*-the-money options will then also be overpriced. Our model therefore predicts a "smile" in the implied volatility curve.[11]

Options on individual stocks do indeed exhibit a smile (Bollen and Whaley, 2004). For index options, however, the implied volatility curve is downward-sloping, rather than U-shaped. It is not unreasonable that our model would fit individual stock option data better.

[10]Of course, since the subsidiary forms part of the parent company, its growth options will also give the returns of the parent company a positively skewed distribution. The parent company's returns will be less skewed than the subsidiary's returns, however, and, as we saw in Section 5.2, it is only high levels of skewness that are overpriced; more moderate skewness is fairly valued. The subsidiary can therefore be overpriced even relative to its parent.

[11]Implicit in this argument are two assumptions: first, that, as discussed by Bollen and Whaley (2004) and others, market makers face risks when they attempt to implement the dynamic strategy needed to enforce Black-Scholes pricing, and that prices can therefore deviate from those predicted by that formula; and second, that put-call parity, based as it is on a static strategy, does hold.

In our framework, investors buy options for speculative purposes: to give themselves a small chance of becoming very wealthy. A speculative trading motive is more plausible in the context of individual stock options than in the context of stock index options, where a large fraction of trading is thought to be driven by institutions.[12]

Our analysis also points to a possible cause for the lack of diversification in many household portfolios. Under cumulative prospect theory, investors may willingly take an undiversified position in positively skewed stocks in order to add skewness to their portfolios.[13] Goetzmann and Kumar (2004) and Mitton and Vorkink (2007) present some relevant evidence. Using data on the portfolios of individual investors, they show that the stocks held by undiversified investors have greater return skewness than those held by diversified investors.

Our framework suggests a link, then, between a number of seemingly unrelated phenomena. At the same time, it also offers a new empirical prediction: that positively skewed stocks should earn lower average returns. Unfortunately, this prediction is not easy to test, because it is hard to forecast a security's future skewness: past skewness, the most obvious potential predictor, has little actual predictive power. We therefore need to find more sophisticated ways of forecasting skewness.

Zhang (2006) suggests one way forward. He groups stocks into industries, and, for each industry in turn, records the return, over the past month, of each stock in the industry. For an industry with 100 firms, say, he therefore has 100 return observations. He then computes the cross-sectional skewness of those 100 returns, and uses it as a measure of the skewness of *all* stocks in that industry. The logic is that, if one stock in the industry did very well last month, leading to a high cross-sectional skewness measure for that industry, then, since stocks in the same industry are similar, we can conjecture that *all* stocks in that industry could potentially earn a high return in the coming months. Zhang (2006) shows that his measure of skewness does predict future skewness; but also, in line with the models in this paper, that it predicts returns, negatively, in the cross-section.

[12]Driessen and Maenhout (2004) offer a portfolio choice view of this argument. In a partial equilibrium setting, they show that, across several expected utility and non-expected utility specifications, the only preferences for which an agent would *not* short out-of-the-money puts on the S&P 500 index are cumulative prospect theory preferences. This suggests, as does our analysis, that such preferences may offer one way of understanding the high prices of out-of-the-money options.

[13]This point is also noted by Polkovnichenko (2005) who, in a portfolio choice setting, shows that an agent with cumulative prospect theory preferences may take an undiversified position in a positively skewed security.

7 Conclusion

We study the asset pricing implications of Tversky and Kahneman's (1992) cumulative prospect theory, with particular focus on its probability weighting component. Our main result, derived from a novel equilibrium with non-unique global optima, is that, in contrast to the prediction of a standard expected utility model, a security's own skewness can be priced: a positively skewed security can be "overpriced," and can earn a negative average excess return. Our results offer a unifying way of thinking about a number of seemingly unrelated financial phenomena, such as the low average return on IPOs, private equity, and distressed stocks; the diversification discount; the low valuation of certain equity stubs; the pricing of out-of-the-money options; and the lack of diversification in many household portfolios.

8 Appendix

Proof of Lemma 1: Since $w(1-P(\cdot))$ is right-continuous and $v(\cdot)$ is continuous, we can integrate (11) by parts to obtain

$$V(\hat{W}^+) = [-v(x)w(1-P(x))]_{x=0}^{x=\infty} + \int_0^\infty w(1-P(x))dv(x). \tag{51}$$

From Chebychev's inequality,

$$P\left[|\hat{W} - E(\hat{W})| \geq Z\right] \leq \frac{\mathrm{Var}(\hat{W})}{Z^2}, \tag{52}$$

which, in turn, implies

$$1 - P(x) \leq P\left[|\hat{W} - E(\hat{W})| \geq x - E(\hat{W})\right] \leq \frac{\mathrm{Var}(\hat{W})}{(x - E(\hat{W}))^2}. \tag{53}$$

Assumptions 1-5 then imply that there exists $A > 0$ such that

$$v(x)w(1-P(x)) \leq Ax^{\alpha - 2\delta} \to 0, \quad \text{as } x \to \infty. \tag{54}$$

The first term on the right-hand side of equation (51) is therefore zero, and so equation (15) is valid. A similar argument leads to equation (16).

Proof of Proposition 2: If we define $g(\cdot) : [0, \infty) \to \Re$ to be $g(x) = x^\alpha$, we can rewrite $v(\cdot)$ as

$$v(x) = \begin{cases} g(x) & \text{for } x \geq 0 \\ -\lambda g(|x|) & \text{for } x < 0. \end{cases} \tag{55}$$

For $\hat{W}_i$, $i = 1, 2$, with the same mean $\mu \geq 0$, we have

$$V(\hat{W}_i) = V(\hat{W}_i^-) + V_A(\hat{W}_i^+) + V_B(\hat{W}_i^+) + V_C(\hat{W}_i^+), \quad i = 1, 2, \tag{56}$$

where

$$\begin{aligned}
V(\hat{W}_i^-) &= -\lambda \int_{-\infty}^0 w(P_i(x))g'(|x|)\,dx, \quad i = 1, 2, && (57)\\
V_A(\hat{W}_i^+) &= \int_0^\mu w(1 - P_i(x))g'(x)\,dx, \quad i = 1, 2, && (58)\\
V_B(\hat{W}_i^+) &= \int_\mu^{2\mu} w(1 - P_i(x))g'(x)\,dx, \quad i = 1, 2, && (59)\\
V_C(\hat{W}_i^+) &= \int_{2\mu}^\infty w(1 - P_i(x))g'(x)\,dx, \quad i = 1, 2. && (60)
\end{aligned}$$

Applying the change of variable $x = 2\mu - x'$ to equations (59) and (60), and noting that, since the distributions are symmetric, $1 - P_i(x) = P_i(2\mu - x) = P_i(x')$, we have

$$\begin{aligned} V_B(\hat{W}_i^+) &= -\int_\mu^0 w(P_i(x'))g'(2\mu - x')\,dx', \\ &= \int_0^\mu w(P_i(x))g'(2\mu - x)\,dx, \qquad i = 1, 2, \end{aligned} \tag{61}$$

$$\begin{aligned} V_C(\hat{W}_i^+) &= -\int_0^{-\infty} w(P_i(x'))g'(2\mu - x')\,dx' \\ &= \int_{-\infty}^0 w(P_i(x))g'(2\mu + |x|)\,dx, \qquad i = 1, 2. \end{aligned} \tag{62}$$

Summing up equations (57), (58), (61), and (62), we have

$$\begin{aligned} V(\hat{W}_i) &= -\int_{-\infty}^0 w(P_i(x))[\lambda g'(|x|) - g'(2\mu + |x|)]\,dx \\ &\quad + \int_0^\mu [w(1 - P_i(x))g'(x) + w(P_i(x))g'(2\mu - x)]\,dx, \quad i = 1, 2, \end{aligned} \tag{63}$$

and

$$\begin{aligned} &V(\hat{W}_1) - V(\hat{W}_2) \\ &= \int_{-\infty}^0 [w(P_2(x)) - w(P_1(x))][\lambda g'(|x|) - g'(2\mu + |x|)]dx \\ &\quad + \int_0^\mu \{[w(1 - P_1(x)) - w(1 - P_2(x))]g'(x) - [w(P_2(x)) - w(P_1(x))]g'(2\mu - x)\}dx. \end{aligned} \tag{64}$$

Since $\widehat{W}_1$ and $\widehat{W}_2$ are symmetric, if condition (iii) holds at all, it must hold for $z = \mu$. This means that $P_1(x) \leq P_2(x)$ for $x < 0 \leq \mu$. Using this fact, as well as the fact that $w(\cdot)$ is increasing and that $\lambda g'(|x|) - g'(2\mu + |x|) > 0$, we see that the first term on the right-hand side of equation (64) is non-negative. Below, we show that

$$w(1 - P_1(x)) - w(1 - P_2(x)) \geq w(P_2(x)) - w(P_1(x)) \geq 0 \quad \text{for } x \in (0, \mu), \tag{65}$$

and that this holds strictly if $P_1(x) < P_2(x)$. Combining this with the fact that $g'(x) > g'(2\mu - x) > 0$ for $x \in (0, \mu)$, we see that the second term on the right-hand side of equation (64) is also non-negative. This implies $V(\hat{W}_1) \geq V(\hat{W}_2)$. Furthermore, given that all distribution functions are right continuous, it is straightforward to show that, if the inequalities in the single-crossing property hold strictly for some $x \in \Re$, then $V(\hat{W}_1) > V(\hat{W}_2)$.

To finish the proof, we need to show that (65) holds, and that it holds strictly if $P_1(x) < P_2(x)$. To do this, consider the function

$$h(p) \equiv w(p) + w(1 - p) = (p^\delta + (1 - p)^\delta)^{-(1-\delta)/\delta}. \tag{66}$$

Since

$$h'(p) = -(1 - \delta)(p^\delta + (1 - p)^\delta)^{-1/\delta}\left(\frac{1}{p^{1-\delta}} - \frac{1}{(1 - p)^{1-\delta}}\right) < 0 \quad \text{for } p \in (0, 1/2), \tag{67}$$

the function $h(p)$ is strictly decreasing for $p \in [0, 1/2]$. Since, for any $x \in (0, \mu)$, $P_1(x) \leq P_2(x) \leq 1/2$, we have $h(P_1(x)) \geq h(P_2(x))$, which then implies the first inequality in (65). The second inequality in (65) follows from the monotonicity of $w(\cdot)$. Finally, if $P_1(x) < P_2(x)$, then (65) holds strictly because $h(\cdot)$ is strictly decreasing in $[0, 1/2]$ and $w(\cdot)$ is strictly increasing.

We now discuss the intuition behind the proof. Proposition 2 implies that, within any class of symmetric distributions that have the same non-negative mean and that satisfy the single-crossing property, a mean-preserving spread is undesirable for an agent with the preferences in (10)-(12), in spite of the risk-seeking induced by the probability weighting function $w(\cdot)$ and the convexity of $v(\cdot)$ in the region of losses.

To see why, note first that for a distribution with positive mean $\mu > 0$ that satisfies conditions (i)-(iii) in the statement of the proposition, a mean-preserving spread means one of two things: either, (a), taking density around μ and spreading it symmetrically outwards towards losses and towards larger gains; or, (b), taking density around μ and spreading it symmetrically outwards towards smaller gains and towards larger gains.

For spreads of type (a), the convexity of $v(\cdot)$ in the region of losses is irrelevant, precisely because it applies to gambles involving only losses. Moreover, the probability weighting function $w(\cdot)$ is neutral to such spreads: while they do add mass to the right tail, which is attractive to an agent who overweights the tails of distributions, they also adds mass to the *left* tail, which is unattractive. The agent's attitude to type (a) spreads is therefore determined by the kink in $v(\cdot)$ at the origin, which, of course, generates aversion to these spreads.

For type (b) spreads, the convexity of $v(\cdot)$ in the region of losses is again irrelevant. The probability weighting function $w(\cdot)$ induces *aversion* to such spreads, because $w(\cdot)$ is more sensitive to differences in probabilities at higher probability levels; in particular, it is more sensitive over $p \,\epsilon\, (\frac{1}{2}, 1)$ – the relevant range for shifts in mass from μ to below μ – than it is over $(0, \frac{1}{2})$, the relevant range for shifts in mass from μ to above μ. The concavity of $v(\cdot)$ in the region of gains only compounds this aversion.

Proof of Proposition 3: Before proving Proposition 3, we first prove the following useful lemma.

Lemma 2: *Consider the preferences in (10)–(12) and suppose that Assumptions 2-5 hold. If $\widehat{W}$ is Normally distributed with mean μ_W and variance σ_W^2, then $V(\hat{W})$ can be written as a function of μ_W and σ_W^2, $F(\mu_W, \sigma_W^2)$. Moreover, for any σ_W^2, $F(\mu_W, \sigma_W^2)$ is strictly increasing in μ_W; and for any $\mu_W \geq 0$, $F(\mu_W, \sigma_W^2)$ is strictly decreasing in σ_W^2.*

Proof of Lemma 2: Since every Normal distribution is fully specified by its mean and variance, we can write $V(\hat{W}) = F(\mu_W, \sigma_W^2)$. Proposition 1 implies that $F(\mu_W, \sigma_W^2)$ is strictly increasing in μ_W. Now consider any pair of Normal wealth distributions, $\hat{W}_1$ and $\hat{W}_2$, with the same mean but different variance. These two wealth distributions satisfy conditions (i)-(iii) in Proposition 2. That proposition therefore implies that, for any $\mu_W \geq 0$, $F(\mu_W, \sigma_W^2)$ is strictly decreasing in σ_W^2.

Our proof of Proposition 3 will now proceed in the following way. We first derive the conditions that characterize equilibrium, assuming that an equilibrium exists. We then show that an equilibrium does indeed exist.

We ignore the violation of limited liability and assume that all securities have positive prices in equilibrium. Consider the mean/standard deviation plane. For any set of positive prices for the J risky assets, Assumption 7 means that we can construct a hyperbola representing the mean-variance (MV) frontier for those assets. Since a risk-free asset is also available, the MV frontier is the tangency line from the risk-free asset to the hyperbola, plus the reflection of this tangency line off the vertical axis. The MV *efficient* frontier is the upper of these two lines. The tangency portfolio, composed only of the J risky securities, has return $\tilde{R}_T$.

By Lemma 2, each agent chooses a portfolio on the MV efficient frontier, in other words, a portfolio with return $\tilde{R} = R_f + \theta(\tilde{R}_T - R_f)$, where θ is the weight in the tangency portfolio. Since agents have identical preferences, they choose the same θ. Market clearing implies $\theta > 0$, and so the tangency portfolio has to be on the upper half of the hyperbola in equilibrium. This, in turn, implies $E(\hat{R}_T) > 0$: the risk-free rate R_f has to be lower than the expected return of the minimum-variance portfolio, the left-most point of the hyperbola.

At time 1, each agent's wealth is given by

$$\begin{aligned} \tilde{W} &= W_0(R_f + \theta(\tilde{R}_T - R_f)) && (68) \\ \hat{W} &= W_0 \theta \hat{R}_T. && (69) \end{aligned}$$

For $\theta \geq 0$, utility is therefore given by

$$U(\theta) = W_0^{\alpha} \theta^{\alpha} V(\hat{R}_T), \tag{70}$$

where

$$V(\hat{R}_T) = -\int_{-\infty}^{0} w(P(\hat{R}_T)) dv(\hat{R}_T) + \int_{0}^{\infty} w(1 - P(\hat{R}_T)) dv(\hat{R}_T). \tag{71}$$

The optimal solution for an agent is:

$$\theta = \begin{cases} 0 & \text{for } V(\hat{R}_T) < 0 \\ \text{any } \theta \geq 0 & \text{for } V(\hat{R}_T) = 0 \\ \infty & \text{for } V(\hat{R}_T) > 0. \end{cases} \tag{72}$$

To clear markets, we therefore need $V(\widehat{R}_T) = 0$.

In equilibrium, the aggregate demand for risky assets, given by (72), must equal the aggregate supply of risky assets, namely the market portfolio. Therefore, $\widehat{R}_T = \widehat{R}_M$. Earlier, we saw that $E(\widehat{R}_T) > 0$ and that $V(\widehat{R}_T) = 0$. Since $\widehat{R}_T = \widehat{R}_M$, we immediately obtain $E(\widehat{R}_M) > 0$ and $V(\widehat{R}_M) = 0$, as claimed in (19) and (20). Finally, $\tilde{R}_T = \tilde{R}_M$ also implies equations (17)-(18). For example, this can be shown by noting that a portfolio with return $\tilde{R}_M + x(\tilde{R}_i - R_f)$ attains its highest Sharpe ratio at $x = 0$.

So far, we have shown that, *if* an equilibrium exists, it is characterized by conditions (17), (19), and (20). We now show that an equilibrium does indeed exist; in other words, that we can find prices for the J risky assets such that conditions (17), (19), and (20) hold.

In conditions (17) and (19), we have J non-redundant equations in J non-redundant unknowns: the J non-redundant equations are condition (19) and any $J-1$ of the J equations in (17); the J non-redundant unknowns are the market price $p_M = \sum_j n_j p_j$ and any $J-1$ of the J prices $\{p_1, p_2, \ldots, p_J\}$. We can therefore solve the J non-redundant equations to obtain the risky asset prices.

It only remains to show that the risky asset prices also imply condition (20). To see this, note that

$$0 = F(0,0) > F(0, \sigma^2(\widehat{R}_M)), \tag{73}$$

where the inequality follows from Lemma 2, which also introduces the function $F(\cdot,\cdot)$. If $E(\widehat{R}_M) \leq 0$, Lemma 2 would then also imply $F(E(\widehat{R}_M), \sigma^2(\widehat{R}_M)) < 0$, contradicting condition (19), which says that $F(E(\widehat{R}_M), \sigma^2(\widehat{R}_M)) = 0$. We therefore have $E(\widehat{R}_M) > 0$, as in condition (20). The intuition is straightforward. Under conditions that apply here, cumulative prospect theory satisfies second-order stochastic dominance. An agent with these preferences therefore dislikes the variance of the market portfolio and only holds it if compensated by a positive risk premium.

The effect of introducing a small skewed security on the prices of existing securities

In Section 5, we noted a useful feature of the heterogeneous holdings equilibrium: the prices of the J original risky assets are *not* affected by the introduction of the skewed security. To see this, note that, after the introduction of the skewed security, the prices of the J original risky assets are determined by

$$E(\widetilde{R}_j) = R_f + \beta_j(E(\widetilde{R}_M) - R_f), \quad j = 1, \ldots, J, \tag{74}$$

or equivalently,

$$\frac{E(\widetilde{X}_j)}{p_j} = R_f + \frac{Cov(\widetilde{X}_j, \widetilde{X}_M)}{Var(\widetilde{X}_M)} \frac{p_M}{p_j} (\frac{E(\widetilde{X}_M)}{p_M} - R_f), \quad j = 1, \ldots, J, \tag{75}$$

where p_j is the price of asset j, $p_M = \sum_j n_j p_j$, and $x_M = \sum_j n_j x_j$.

In the economy of Section 4, the prices $\{p'_j\}$ of the J original risky assets are given by

$$\frac{E(\widetilde{X}_j)}{p'_j} = R_f + \frac{Cov(\widetilde{X}_j, \widetilde{X}_M)}{Var(\widetilde{X}_M)} \frac{p'_M}{p'_j} (\frac{E(\widetilde{X}_M)}{p'_M} - R_f), \quad j = 1, \ldots, J, \tag{76}$$

where $p'_M = \sum_j n_j p'_j$. From equations (19) and (23), we know that the return on the market portfolio formed from the J risky assets alone satisfies $V(\widehat{R}_M) = 0$, whether or not the skewed security is present. This implies $p_M = p'_M$, which, in turn, means that the equations for $\{p_j\}$ in (75) are identical to the equations for $\{p'_j\}$ in (76). The prices of the J original risky assets are therefore the same, whether or not the skewed security is present.

Proof of Proposition 4: The Gateaux derivative in (38) follows from

$$\frac{\partial V(\hat{R} + x)}{\partial x}|_{x=0} = \frac{\partial}{\partial x}|_{x=0} \left[- \int_{-\infty}^{0} w(P(R - x)) dv(R) + \int_{0}^{\infty} w(1 - P(R - x)) dv(R) \right]. \tag{77}$$

To show the main result, note that, to the first order of ε,

$$\begin{aligned} \delta V &\equiv V(\hat{R} + \varepsilon \hat{R}_n) - V(\hat{R}) \\ &\approx - \int_{-\infty}^{0} w'(P(R)) \delta P(R) dv(R) - \int_{0}^{\infty} w'(1 - P(R)) \delta P(R) dv(R), \end{aligned} \tag{78}$$

where, again to the first order of ε,

$$\begin{aligned} \delta P(R) &\equiv P(\hat{R} + \varepsilon \hat{R}_n \leq R) - P(\hat{R} \leq R) \\ &= E\left[1(\hat{R} + \varepsilon \hat{R}_n \leq R) - 1(\hat{R} \leq R) \right] \\ &= E\left[-1(\varepsilon \hat{R}_n > 0) 1(R - \varepsilon \hat{R}_n < \hat{R} \leq R) + 1(\varepsilon \hat{R}_n < 0) 1(R < \hat{R} \leq R - \varepsilon \hat{R}_n) \right] \\ &\approx E\left[-1(\varepsilon \hat{R}_n > 0) f(R|\hat{R}_n)(\varepsilon \hat{R}_n) + 1(\varepsilon \hat{R}_n < 0) f(R|\hat{R}_n)(-\varepsilon \hat{R}_n) \right] \\ &= E\left[f(R|\hat{R}_n)(-\varepsilon \hat{R}_n) \right] \\ &= -\varepsilon E(\hat{R}_n) f(R), \end{aligned} \tag{79}$$

where $f(R|\hat{R}_n)$ and $f(R)$ are the conditional and unconditional probability densities of $\hat{R}$ at R, respectively, and where the last equality in equation (79) follows from the independence assumption. Substituting equation (79) into equation (78), we obtain equation (37).

9 References

Barberis, N., Huang, M., and T. Santos (2001), "Prospect Theory and Asset Prices," *Quarterly Journal of Economics* 116, 1-53.

Barberis, N., Huang, M., and R. Thaler (2006), "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing," *American Economic Review* 96, 1069-1090.

Benartzi, S. and R. Thaler (1995), "Myopic Loss Aversion and the Equity Premium Puzzle," *Quarterly Journal of Economics* 110, 75-92.

Berger, P. and E. Ofek (1995), "Diversification's Effect on Firm Value," *Journal of Financial Economics* 37, 39-66.

Bollen, N. and R. Whaley (2004), "Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?" *Journal of Finance* 59, 711-753.

Brunnermeier, M. and J. Parker (2005), "Optimal Expectations," *American Economic Review* 95, 1092-1118.

Brunnermeier, M., Gollier, C., and J. Parker (2007), "Optimal Beliefs, Asset Prices, and the Preference for Skewed Returns," *American Economic Review (Papers and Proceedings)*, forthcoming.

Campbell, J.Y., Hilscher, J. and J. Szilagyi (2006), "In Search of Distress Risk," Working paper, Harvard University.

Dahlquist, M. and F. De Jong (2004), "Pseudo Market Timing: Fact or Fiction?" Working paper, Stockholm Institute for Financial Research.

De Giorgi, E., Hens, T., and H. Levy (2003), "Prospect Theory and the CAPM: A Contradiction or Co-existence?" Working paper, University of Zurich.

Driessen, J. and P. Maenhout (2004), "A Portfolio Perspective on Option Pricing Anomalies," Working paper, INSEAD.

Epstein, L. and S. Zin (1990), "First-order Risk Aversion and the Equity Premium Puzzle," *Journal of Monetary Economics* 26, 387-407.

Goetzmann, W. and A. Kumar (2004), "Why do Individual Investors Hold Undiversified Portfolios?" Working paper, Yale University.

Harvey, C. and A. Siddique (2000), “Conditional Skewness in Asset Pricing Tests,” *Journal of Finance* 55, 1263-1290.

Kahneman, D. (2003), “Maps of Bounded Rationality: Psychology for Behavioral Economics,” *American Economic Review* 93, 1449-75.

Kahneman, D. and A. Tversky (1979), “Prospect Theory: An Analysis of Decision under Risk,” *Econometrica* 47, 263-291.

Kane, A. (1982), “Skewness Preference and Portfolio Choice,” *Journal of Financial and Quantitative Analysis* 17, 15-25.

Kraus, A. and R. Litzenberger (1976), “Skewness Preference and the Valuation of Risk Assets,” *Journal of Finance* 31, 1085-1100.

Lamont, O. and R. Thaler (2003), “Can the Market Add and Subtract? Mispricing in Tech Stock Carve-outs,” *Journal of Political Economy* 111, 227-268.

Lang, L. and R. Stulz (1994), “Tobin’s Q, Corporate Diversification, and Firm Performance,” *Journal of Political Economy* 102, 1248-1280.

Mitchell, M., Pulvino P., and E. Stafford (2002), “Limited Arbitrage in Equity Markets,” *Journal of Finance* 57, 551-584.

Mitton, T. and K. Vorkink (2006), “Why Do Firms with Diversification Discounts Have Higher Expected Returns?” Working paper, Brigham Young University.

Mitton, T. and K. Vorkink (2007), “Equilibrium Undiversification and the Preference for Skewness,” *Review of Financial Studies*, forthcoming.

Polkovnichenko, V. (2005), “Household Portfolio Diversification: A Case for Rank-dependent Preferences,” *Review of Financial Studies* 18, 1467-1501.

Rabin, M. (2000), “Risk Aversion and Expected Utility Theory: A Calibration Theorem,” *Econometrica* 68, 1281-1292.

Ritter, J. (1991), “The Long-run Performance of Initial Public Offerings,” *Journal of Finance* 46, 3-27.

Schultz, P. (2003), “Pseudo Market Timing and the Long Run Performance of IPOs,” *Journal of Finance* 58, 483-518.

Tversky, A., and D. Kahneman (1992), “Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty* 5, 297-323.

Zhang, Y. (2006), "Individual Skewness and the Cross-section of Average Returns," Working paper, Yale University.

Table 1: The table describes the properties of equilibrium when agents with cumulative prospect theory preferences are allowed to invest in N identical, independent, positively skewed securities in addition to a Normally distributed market portfolio: μ_M is the average excess return on the market portfolio, $E(\hat{R}_n)$ is the average excess return on each positively skewed security; $E(\hat{R}_n^{NS})$ is the average excess return on each positively skewed security in an economy that is identical in structure except that short sales are not allowed.

N	μ_M	$E(\hat{R}_n)$	$E(\hat{R}_n^{NS})$
1	0.075	-0.046	-0.047
2	0.075	-0.046	-0.047
5	0.075	-0.046	-0.047
10	0.074	-0.046	-0.047

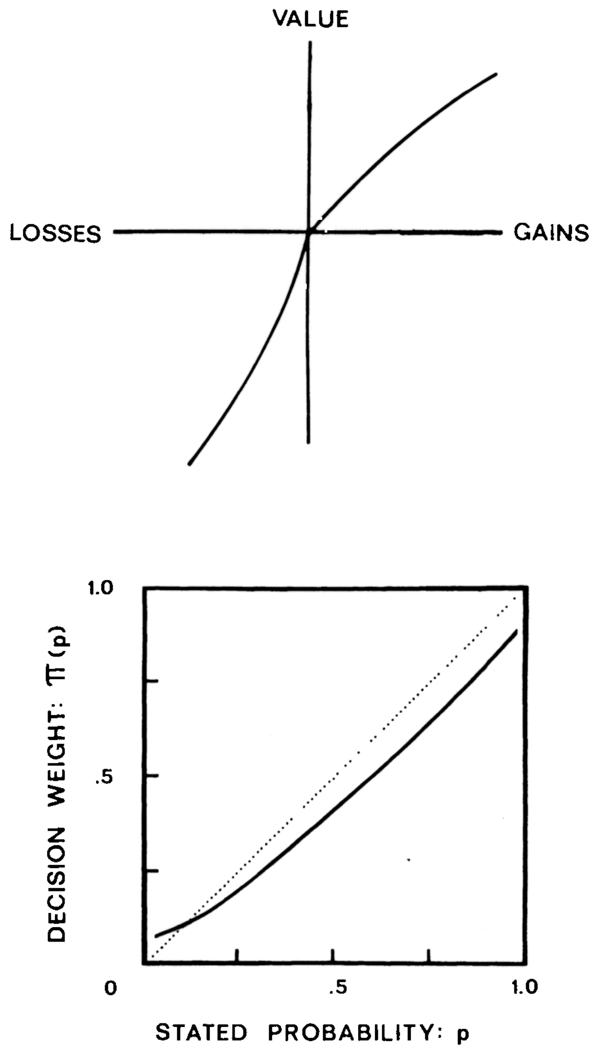

Figure 1. The two panels show Kahneman and Tversky's (1979) proposed value function $v(\cdot)$ and probability weighting function $\pi(\cdot)$.

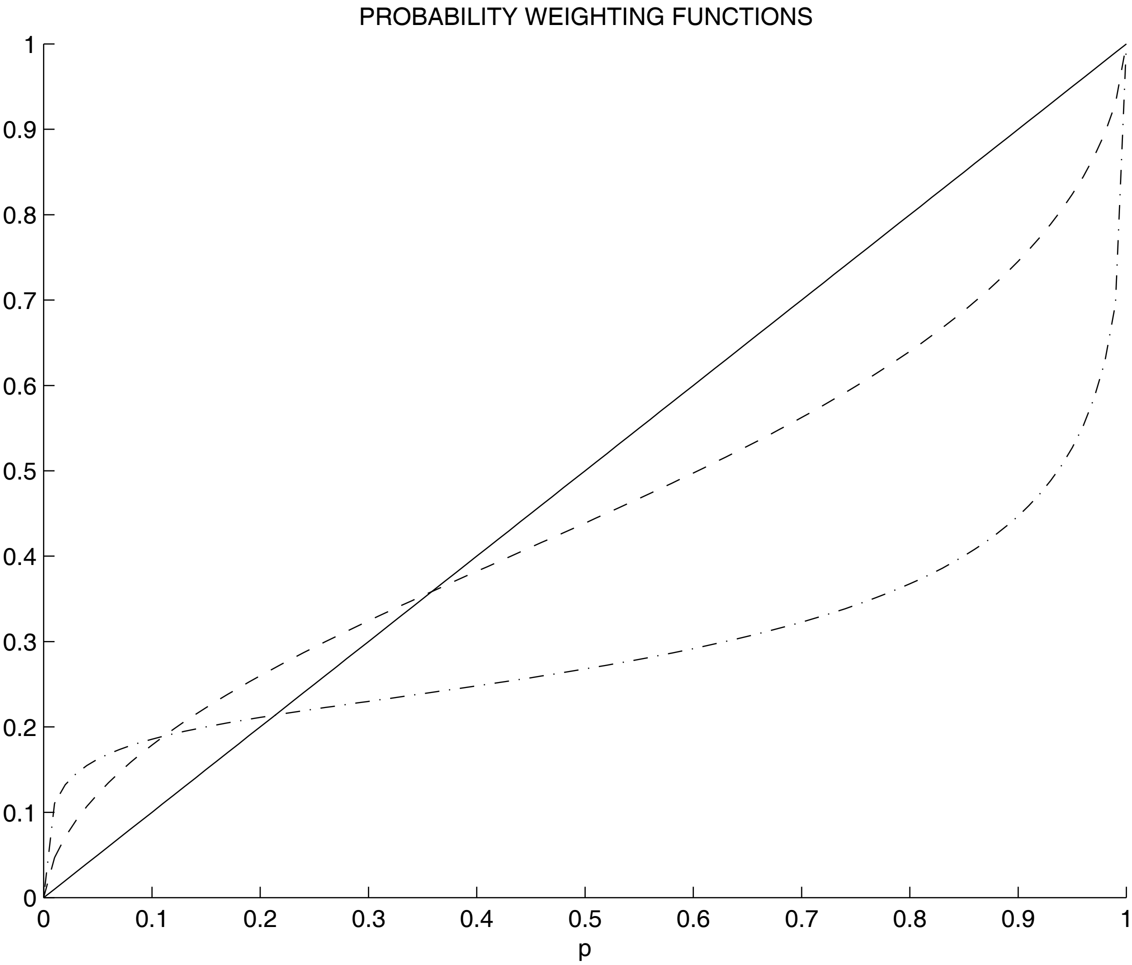

Figure 2. The figure shows the form of the probability weighting function proposed by Tversky and Kahneman (1992), for parameter values $\delta = 0.65$ (dashed line), $\delta = 0.4$ (dash-dot line), and $\delta = 1$, which corresponds to no probability weighting at all (solid line). p is the objective probability.

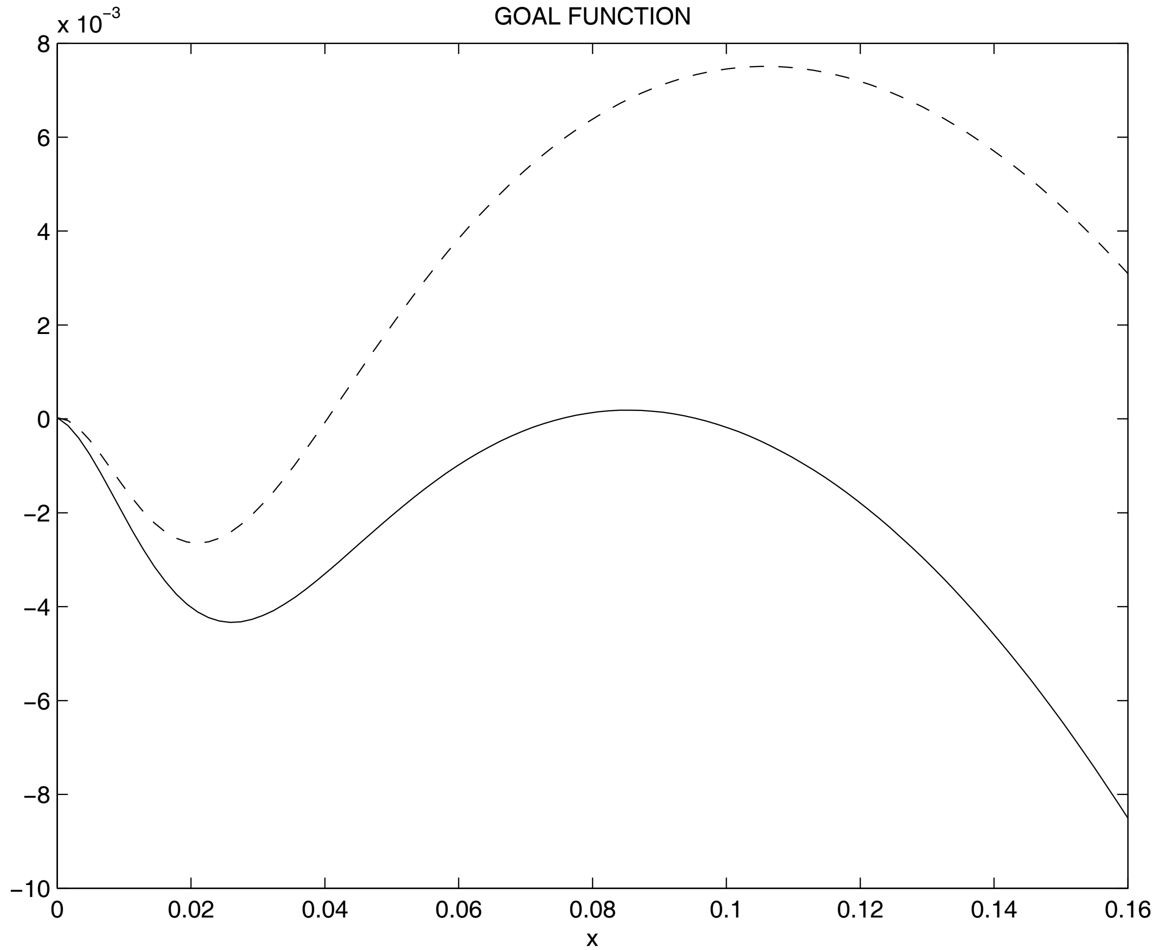

Figure 3. The figure shows the utility that an investor with cumulative prospect theory preferences derives from adding a position x in a positively skewed security to his current holdings of a Normally distributed market portfolio. The two lines correspond to different mean returns on the skewed security.

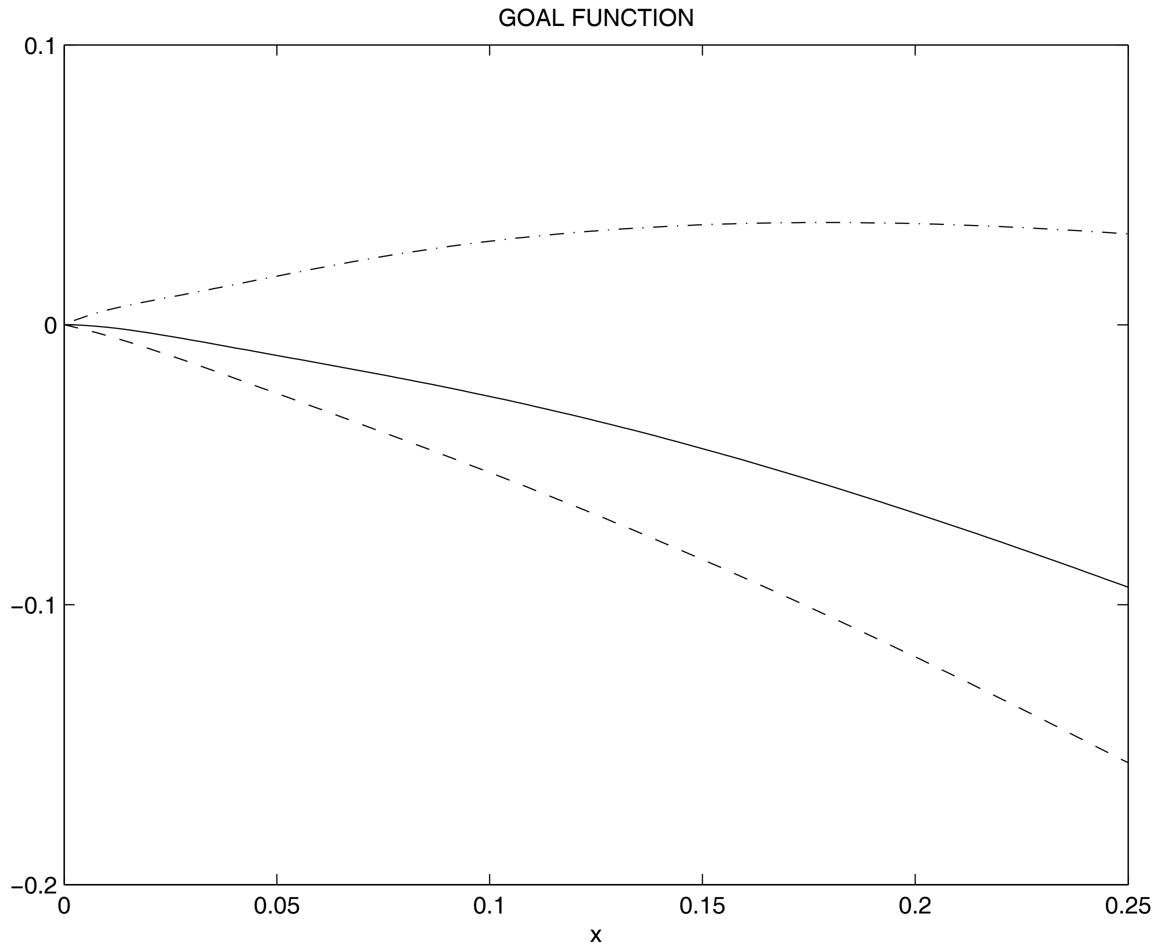

Figure 4. The figure shows the utility that an investor with cumulative prospect theory preferences derives from adding a position x in a positively skewed security to his current holdings of a Normally distributed market portfolio. The three lines correspond to different mean returns on the skewed security.

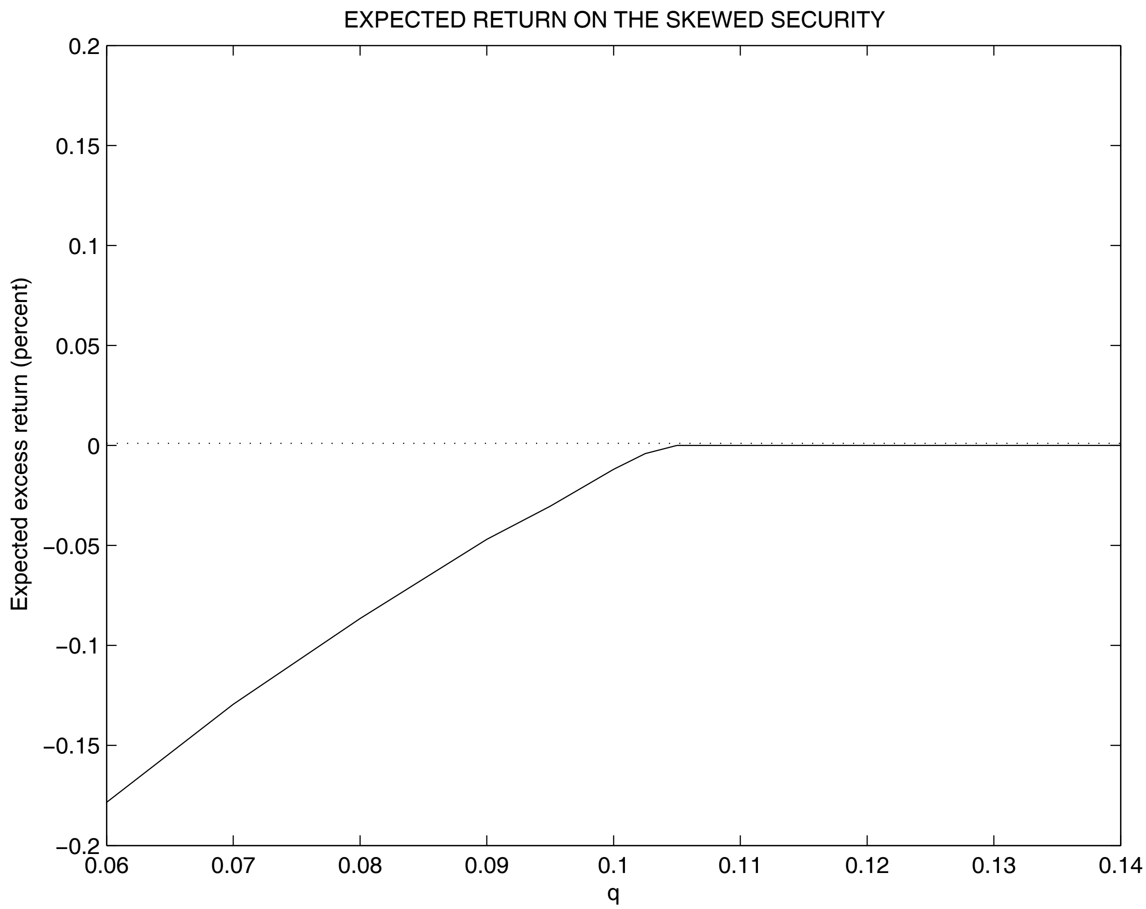

Figure 5. The figure shows the expected return in excess of the risk-free rate earned by a small, independent, positively skewed security in an economy populated by cumulative prospect theory investors, plotted against a parameter of the the security's return distribution, q, that determines the security's skewness (a low q corresponds to high skewness).

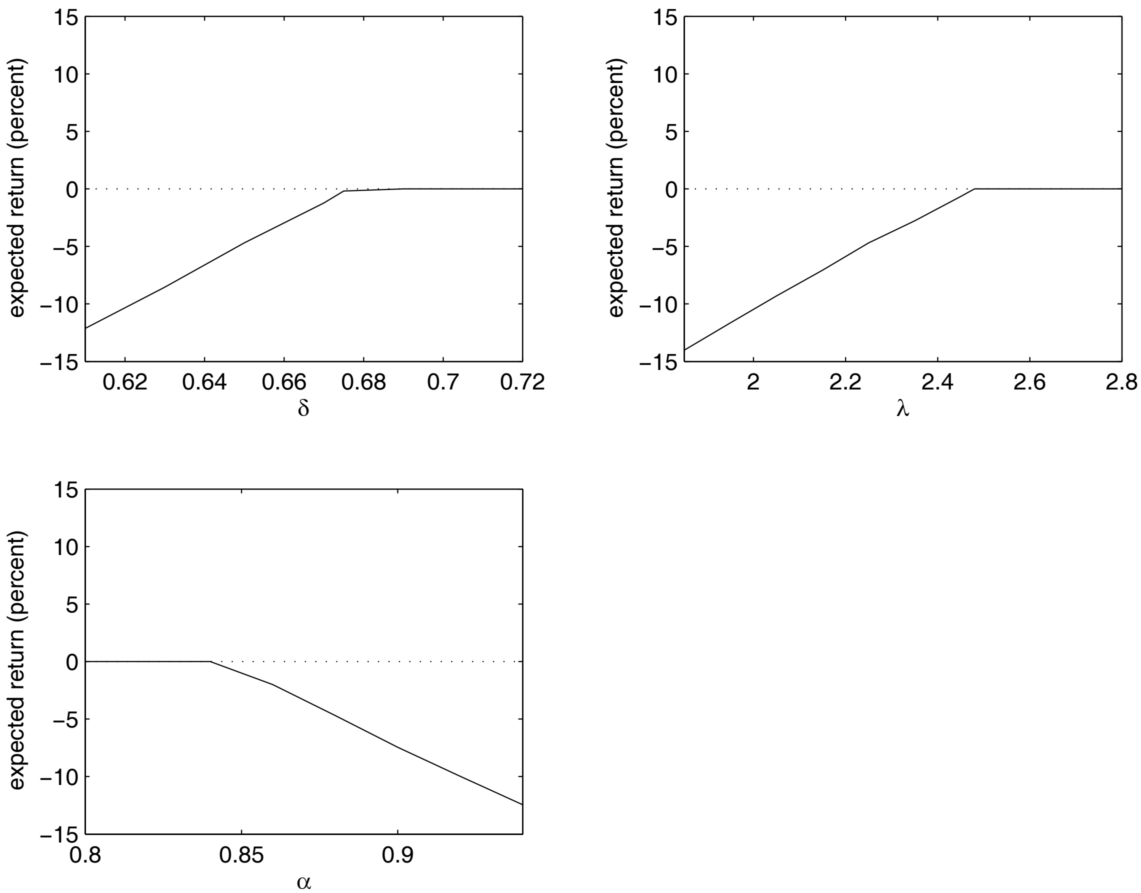

Figure 6. The figure shows the expected return in excess of the risk-free rate earned by a small, independent, positively skewed security in an economy populated by cumulative prospect theory investors, plotted against parameters of investors' utility functions. As δ falls, investors overweight small probabilities more heavily; as λ increases, they become more sensitive to losses; and as α falls, their marginal utility from gains falls.