Economics and Portfolio Strategy

Peter L. Bernstein, Inc. ·· 575 Madison Avenue, Suite 1006 ·· New York, N.Y. 10022
Phone: 212·421·8385 FAX: 212·421·8537

October 15, 2004

SKEW YOU, SAY THE BEHAVIORALISTS[1]

By Mark Kritzman

Hedge funds and behavioral finance are hot topics in portfolio management. Hedge funds have grown to become a trillion dollar industry and the focus of seemingly endless journal articles. And behavioral finance devotees, once looked upon as members of the radical fringe, are now firmly ensconced in the mainstream, thanks in part to Daniel Kahneman's Nobel Prize. It turns out that these two topics intersect in a particularly interesting way, which gives rise to my R-rated title. Let's begin with hedge funds.

Statistical Summary of Hedge Fund Returns

By now almost everyone knows that hedge fund return distributions differ from the distributions produced by traditional assets. Traditional assets, such as stocks and bonds, produce approximately normal distributions, resembling a bell shaped curve. In contrast, hedge funds often generate negatively skewed distributions, which means they produce a greater number of above-average returns than a normal distribution and fewer though more extreme negative outcomes. To put it more prosaically, hedge funds with negative skewness perform reasonably well most of the time, but every now and then, they do really badly.

Hedge fund return distributions also display a feature known as leptokurtosis. Although leptokurtosis may sound like a condition you would treat with penicillin, it merely refers to the hedge funds' fat tailed distributions. Compared to a normal distribution, a larger proportion of the returns are located near the extremes than the mean of the distribution.

In our study, we examined four different types of hedge funds: 25 equity funds, 10 convertible arbitrage funds, 19 event-driven funds, and 7 merger arbitrage funds.[2]

Equity hedge funds maintain a long position in equities perceived to be undervalued and hedge these holdings by selling short individual stocks or stock index futures perceived to be overvalued or neutrally valued. Often the managers of these funds employ leverage in order to raise the expected return, which causes these funds to experience a relatively high incidence of extreme outcomes; hence they display a higher level of kurtosis than a normal distribution.

[1] The results in this essay come from a paper I coauthored with Jan-Hein Cremers and Sebastien Page, which has the decidedly more prosaic title, "Optimal Hedge Fund Allocations: Do Higher Moments Matter," Revere Street Working Paper Series, Financial Economics 272-13.
[2] We use monthly hedge fund returns for the 10-year period from January 1994 through December 2003, provided by the Center for International Securities and Derivatives Markets (CISDM), which is associated with the Isenberg School of Management at the University of Massachusetts.

Convertible arbitrage funds typically purchase convertible bonds and hedge the equity risk by selling short the company's underlying common stock. These funds display both higher levels of kurtosis than a normal distribution and negative skewness.

Event driven funds search for opportunities created by spin-offs, mergers and acquisitions, bankruptcy reorganizations, recapitalizations, and share repurchases. They employ long and short positions in common and preferred stock, debt securities, and options. Because they attempt to capitalize on significant events, their return distributions display higher than normal kurtosis. They are also negatively skewed on average. Some investors refer to this strategy as corporate life cycle investing.

Merger arbitrage funds attempt to profit by acquiring the stock of companies that are takeover targets and by selling the stock of the acquiring companies. These funds also display higher than normal kurtosis.

Table 1 shows the average skewness and kurtosis for each hedge fund style. A normal distribution has skewness equal to 0 and kurtosis equal to 3. Table 1 also shows the fraction of funds in each category that failed the Jarque Bera test for normality. This test determines if the degree of skewness and kurtosis is sufficiently severe to render the distribution "statistically non-normal."[3] Suffice it to say, the hedge fund returns in our study are overwhelmingly non-normal.

Table 1: Summary of Hedge Fund Non-normality

	Average Skewness	Average Kurtosis	Percent Failing JB
Equity Hedge	0.19	6.10	80%
Convertible Arbitrage	-0.61	5.42	90%
Event Driven	-0.37	7.38	84%
Merger Arbitrage	0.16	6.51	100%

As you will soon discover, the non-normality of hedge fund returns is critically important to my story. But let's shift our focus for a moment to behavioral finance.

Behavioral Finance

Neo-classical finance assumes that investors act rationally and, moreover, are averse to risk. Daniel Bernoulli introduced the notion of risk aversion as far back as 1738 in his classic paper on risk measurement.[4] He argued that utility for the typical investor equals the natural logarithm of wealth. This relationship is shown in Figure 1 and reveals that utility rises with wealth but at a diminishing rate.

---

[3] The Jarque-Bera test is a non-parametric test of normality that is based on skewness and kurtosis. It is non-parametric in the sense that it tests for normality without specification of a particular mean or variance. The actual statistic is a sum of two independent parts, one built from the third moment, the other from the fourth moment. Each of these parts is distributed as a standard normal. The Jarque-Bera statistic is their sum of squares and thus should be compared to a chi-square distribution with two degrees of freedom.

[4] Bernoulli, Daniel. 1954. "Exposition of a New Theory on the Measurement of Risk." *Econometrica*, (Translation from 1738 version).

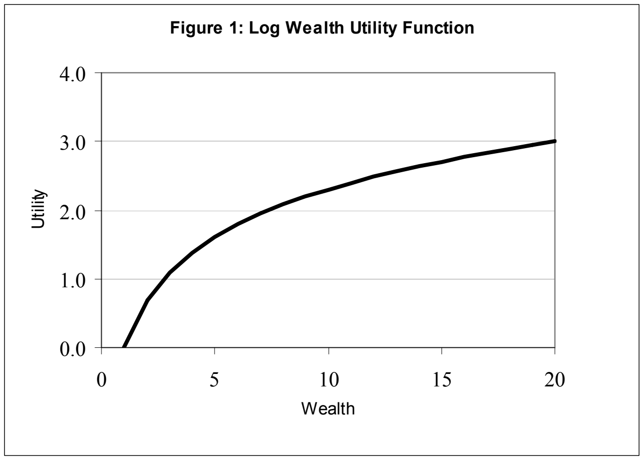

It is easy to see how diminishing marginal utility leads to risk aversion. Figure 1 shows that 20 units of wealth (let's call them dollars) convey 3.0 units of utility (ln(20) = 3.00). Now consider a risky gamble which has an equal chance of returning \$50 or nothing at all. Although the expected outcome is greater than \$20 (.5 x \$50 + .5 x \$0 = \$25), it conveys only 1.96 units of utility (.5 x ln(50) + .5 x ln(0) = 1.96). Hence, a log wealth investor will prefer \$20 for sure to a risky gamble that has an uncertain expected value of \$25. The preference for a lower certain value to a higher uncertain value is the definition of risk aversion.

Risk aversion served as one of the central tenets of financial theory until Daniel Kahneman and Amos Tversky came along. In 1979 they published an influential article[5] in which they demonstrated that people in certain circumstances seek risk rather than avoid it.

Consider the following choice:

A: 80% chance of gaining \$4,000
20% chance of gaining nothing

B: \$3,000 gain for sure

Which do you prefer, A or B?

Now consider a second choice:

A: 80% chance of losing \$4,000
20% chance of losing nothing

B: \$3,000 loss for sure

Which do you prefer now?

---

[5] Kahneman D. and Tversky, A. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 1979, 47, 263-290.

When faced with the first choice, most respondents select B, a sure gain of $3,000, rather than A, an uncertain outcome with a greater expected value of $3,200 (80% x $4,000 + 20% x $0 = $3,200). This choice is consistent with the notion of risk aversion. When faced the second choice, however, most respondents select A over B. They prefer a risky gamble with an expected loss of $3,200 to a certain loss of $3,000. These contradictory preferences imply that investors are risk averse when faced with gains but risk seeking when faced with losses.

Risk aversion yields a concave utility function as shown in Figure 1, so it must follow that we represent risk seeking behavior with a convex utility function. When we combine risk aversion in the domain of gains with risk seeking in the domain of losses, we end up with what behavioralists call an S-shaped value function, shown in Figure 2.

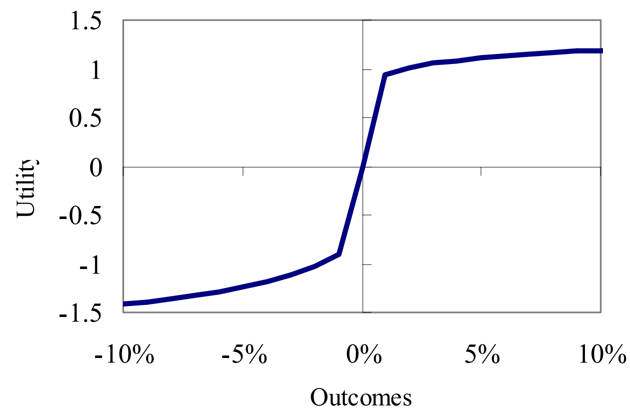

Figure 2. S-Shaped Value Function

Let's catch our breath. So far we have determined that hedge funds are unlike traditional investments; they generate skewed distributions with fat tails rather than normal distributions. We have also determined that most investors are not consistently risk averse. When it comes to losses they prefer to roll the dice. Now let's see how these observations intersect to justify my catchy title.

Full-Scale Optimization

Many sophisticated investors employ mean-variance optimization to form portfolios. This wonderful innovation by Harry Markowitz allows us to identify the optimal blend of assets based only on their expected returns, standard deviations, and correlations. This procedure works extraordinarily well, even for assets with non-normal distributions, as long as investors have log wealth or similar utility functions.[6] Remember, log wealth utility reflects Bernoulli's view of risk aversion and assumes that we always want more wealth, but each addition has less utility than the previous. But mean-variance optimization ignores skewness and kurtosis and therefore performs poorly if investors have utility functions that are more sensitive to these features, such as the behavioralists' S-shaped value functions.

[6] See, for example, my January 15, 2004 missive to *Economics and Portfolio Stratergy*, entitled, "Response to Samuelson's Critique of Mean-Variance Analysis."

Computational efficiency now allows us to perform full-scale optimization as an alternative to mean-variance optimization. With this approach we calculate a portfolio's utility for every period in our sample considering as many asset mixes as necessary in order to identify the weights that yield the highest expected utility, given any utility function including S-shaped value functions. Suppose, for example, we wish to find the optimal blend between two funds given their historical returns and assuming the investor has log wealth utility. We compute utility each period as $\ln[(1+R_A) \times W_A + (1+R_B) \times W_B]$, where R_A and R_B equal the returns of funds A and B, and W_A and W_B equal their respective weights. We then shift the fund's weights using a numerical search procedure until we find the combination that maximizes expected utility. This approach implicitly takes into account all of the features of the empirical sample, including skewness, kurtosis, and any other peculiarities of the distribution. In a nutshell, it identifies the truly optimal portfolio based on the entire empirical distribution. Mean-variance optimization, in contrast, yields an *approximately* optimal portfolio based only on mean and variance.

We applied both mean-variance optimization and full-scale optimization to determine optimal hedge fund portfolios, using the hedge fund returns described earlier. In all cases, the full-scale results refer to the truly optimal portfolio, whereas the mean-variance results apply to the mean-variance efficient portfolio with the same expected return as the truly optimal full-scale portfolio.

Table 2 depicts the fraction of the mean-variance efficient portfolios one would need to trade in order to invest the portfolio according to the full-scale optimal weights. It reveals that mean-variance optimization closely approximates the truly optimal weights, assuming the investor has log wealth utility. For investors with S-shaped value functions, however, mean-variance optimization doesn't come close to the truly optimal portfolio. The reason for the huge miss is that skewness and kurtosis really matter to investors with S-shaped preferences.

Table 2: Turnover Required to Shift from MV to FS Efficiency

	Log Wealth Utility	S-Shaped Value Function
Equity Hedge	4%	40%
Convertible Arbitrage	0%	56%
Event Driven	0%	31%
Merger Arbitrage	0%	12%
All Hedge Funds	15%	57%

Table 3 shows the skewness of the mean-variance and full-scale portfolios. Remember, the portfolios determined by full-scale optimization are the truly optimal portfolios. The mean-variance portfolios are approximations.

Table 3: Skewness

	Log Wealth Utility		S-Shaped Value Function	
	MV	FS	MV	FS
Equity Hedge	-0.23	-0.08	-0.08	-0.34
Convertible Arbitrage	4.78	4.78	-0.40	-0.32
Event Driven	1.02	1.02	-1.58	-1.82
Merger Arbitrage	0.17	0.17	-1.02	-1.10
All Hedge Funds	-0.56	-0.55	-1.10	-1.98

Table 3 confirms that mean-variance optimization comes close to the correct full-scale results, as long as investors have log-wealth utility, but it produces portfolios with significantly incorrect degrees of skewness for S-shaped investors. The more important -- perhaps even startling -- insight is that *S-shaped investors like negative skewness*. They have this preference not because they are fond of extremely negative returns, but because they do not dislike them that much more than they dislike mildly negative returns. And they really like the high concentration of returns just above the threshold associated with negative skewness.

Finally, Table 4 shows the kurtosis of the mean-variance and full-scale portfolios for log wealth and S-shaped investors. It also reveals that mean-variance optimization performs well for log wealth investors, essentially because these investors are relatively insensitive to kurtosis. It does not produce the right amount of kurtosis for S-shaped investors, however. Kurtosis appears not to be problematic for these investors, again because they are not especially averse to extremely negative results.

Table 4: Kurtosis

	Log Wealth Utility		S-Shaped Value Function	
	MV	FS	MV	FS
Equity Hedge	4.70	4.65	2.86	3.80
Convertible Arbitrage	4.78	4.78	5.38	5.88
Event Driven	8.53	8.53	9.24	12.56
Merger Arbitrage	4.92	4.92	7.59	8.70
All Hedge Funds	3.86	3.85	6.62	11.08

Summary

- Hedge funds typically do not generate normally distributed returns. Often their returns are negatively skewed and fat tailed.
- Neo-classical finance assumes that most investors have log wealth utility or similar preferences, which implies they are consistently risk averse.
- Mean-variance optimization closely approximates truly optimal portfolios revealed by full-scale optimization, as long as investors have log wealth or similar preferences.
- Behavioral finance, however, holds that many investors are risk averse when faced with gains but risk seeking when faced with losses. We can represent these preferences graphically with an S-shaped value function. The risk seeking behavior in the domain of losses implies that investors do not dislike extremely negative returns that much more than they dislike mildly negative returns.
- S-shaped investors prefer a negatively skewed distribution because it offers a high concentration of outcomes above the loss threshold and because these investors are relatively insensitive to significantly negative returns.
- It is this preference for negative skewness by S-shaped investors that inspired my title.

oOo