

*Performance Persistence in the
Alternative Investment Industry*

James M. Park

Chairman & CEO

PARADIGM Capital Management, Inc.

Asst. Professor of Finance and Director of the Center for Financial Markets Research

Long Island University

Jeremy C. Staum¹

Director of Research

PARADIGM Capital Management, Inc.

Doctoral Student in Mathematical Finance

Columbia University

Please address communications to:

James Park

12 East 49th Street, 32nd Floor

New York, NY 10017

(212) 271-3388

jp@paradigm-ny.com

¹Jeremy Staum is supported in part by an NSF graduate research fellowship.

Performance Persistence in the Alternative Investment Industry

We construct an improved measure of skill among commodity trading advisors (CTAs) and hedge fund managers. The theoretical issues surrounding the possibility of internal leverage receive particular attention. The results strengthen previous findings that past performance, properly adjusted for risk, does give some indication of future performance. The statistical significance of the results is extreme, while the practical significance is moderate for CTAs and very great for hedge funds. Consequently, there should be noticeable rewards for implementing this criterion in portfolio construction by selecting traders with superior skill.

I. Introduction

It seems almost every piece of paper presenting historical performance data in the financial industry reiterates the ubiquitous (and mandatory) mantra “Past performance is not necessarily indicative of future returns.” While it is certainly laudable to protect the gullible from unsound statistical inference, this slogan leaves unanswered an important question: to what extent does past performance indicate future returns?

Some authors have found evidence for persistence of performance which would enable the investor to beat the market. (e.g. Hendricks, Patel, and Zeckhauser (1993), Goetzmann and Ibbotson (1994), Park (1995)) Lately, the academic focus has shifted to qualifying these findings, based on the heavy dependence of persistence on repeat losers, and the subtle ways in which survivorship bias can undermine historical studies of persistence. (e.g. Brown, Goetzmann, Ibbotson, and Ross (1992), (1997), Brown and Goetzmann (1995), Brown, Goetzmann, and Ross (1995)) Most recently, Goetzmann, Ibbotson, and Brown (1998) presented evidence that there is no performance persistence among hedge funds, which would be surprising for an industry with a reputation for manager skill. This paper will present evidence for performance persistence among CTAs and hedge funds.

First, consider the theoretical reasons for believing in manager skill in these industries. While it is plausible to imagine that the stock market is efficient, so mutual fund managers add no value and their relative returns should not be predictable, this is not the case for alternative investments such as CTAs or hedge funds. Mutual funds can make a profit in the stock market without skill because investment in production is, in the long term, profitable. On the other hand,

although CTAs do not exist in order to take net long positions in any commodity or security, as a whole they have generated positive returns over the last fifteen years. This is not the result of passive investment. It is a reward for CTAs' contribution to accurately pricing commodities, providing liquidity, and accepting risk for commodity producers and consumers who wish to hedge. This suggests that CTAs and hedge funds generate profits precisely by having skill at producing pricing information in an inefficient market. The questions then are how to measure this skill, and whether managers demonstrate persistent skill.

We address these questions by examining traders' performance relative to an industry index of their peers. This is the method which Schneeweis, Spurgin, and McCarthy (1997) advocated for extracting information from a CTA's track record, and with which Park (1995) analyzed skill among CTAs. The persistence of traders' tendencies to outperform or underperform the index is evidence that some traders have greater skill than others. Section II describes the data and addresses possible biases. Section III describes the proposed measure of skill and its theoretical justification. Statistical evidence for persistence of skill occupies section IV, while section V suggests directions for future research.

II. Data and Biases

The data, covering 1986-1997, comes from the published databases of TASS Management, a London-based information service. The two databases cover funds and CTAs separately. Unpublished TASS data on defunct funds has been added to the funds database. The TASS databases attempt to be as inclusive as possible, but because of the size of the industry and

the nature of voluntary reporting, the data is in truth only a very large nonrandom sample. Still, indices constructed as the returns of equal-weighted portfolios of all traders in the two databases may be taken as representative of the performance of the two industries respectively.

Data for funds of funds, i.e. funds which have multiple managers and invest in other funds, has been removed from the fund database, as suggested by Fung and Hsieh (1997a). The intent was to approximate a database of hedge funds, but many programs with a fund structure (as opposed to managed accounts) which trade primarily futures are included as funds. Conversely, many programs in the CTA database do not have a managed account structure. As Fung and Hsieh pointed out, in recent years there has been convergence between the two industries which makes it difficult to distinguish between them, especially in the case of programs with a fund structure trading futures on equity indices. Nonetheless, as the results below show, there is value to treating the two databases separately.

Throughout the study, where r is the published percentage return, data is expressed on a logarithmic scale: $\log(1 + r/100)$. This is natural for return streams, because it shows that a loss of 33% and a gain of 50% (or vice versa) result in zero change. Furthermore, the log transformation reduces some of the data's unpleasant distributional properties such as outliers and right-skewness.

As usual, the data suffers from self-selection and survivorship biases. Park's (1995) left-truncation technique deals with the effect of self-selection bias. Each CTA's history has the first 28 months removed, and each hedge fund's history has the first 15 months removed. Notably, the switching regression procedure yields exactly the same cutoff for CTAs which Park found, using the MAR data set 1983-1992.

It is important to correct for self-selection bias before testing for persistence of performance. The self-selection arises when many traders enter the database with “instant histories”. Of course, predominantly those traders with good track records will choose to start reporting and soliciting money. For this reason, instant histories tend to be uniformly good and show a spurious internal persistence of good performance.

On the other hand, it is not as vital to address survivorship bias for this purpose. As shown by Brown, Goetzmann, and Park (1997), traders who exit the database usually do so because of long-term persistent underperformance. Disproportionate exclusion of such track records would weaken the evidence for the hypothesis that losers repeat. Also, the method of index construction insures that the cross-sectional sum of excess returns is always zero, ameliorating the troublesome effect of cross-sectional correlations described by Brown, Goetzmann, Ibbotson, and Ross (1992). Furthermore, results below (see Figure I) show that the survivorship-induced “J” shape of persistence, described by Hendricks, Patel, and Zeckhauser (1997), is not present in the data. In this case, survivorship bias should not result in overstated claims of persistent performance.

III. Measuring Skill

Both Goetzmann and Ibbotson (1994) and Park (1995) presented evidence that there is a statistically significant tendency for winners to repeat when “winners” are defined as those whose Jensen’s alpha is above the median. Jensen’s alpha, defined in the Capital Asset Pricing Model (CAPM), is an excess return above the rate of return which should accrue to an asset with a given

amount of nondiversifiable risk. The CAPM was designed for the world of stocks and mutual funds which Goetzmann and Ibbotson studied: they used the return on the S&P500 as an index. Park (1995) studied CTAs and managed futures funds and pools, using in each case an index of the industry's returns.

Plainly, the CAPM does not apply verbatim to a universe of hedge funds and CTAs, nor did it apply verbatim to the universe of equity and fixed income funds. For instance, it is no longer an equilibrium theory. A discussion of the nature of this theory as it applies to alternative investments is beyond the scope of this paper. Nonetheless, it is better to invest with the alpha winners than with the losers in constructing a portfolio of alternative investments.

Unfortunately, there is a problem with using alpha to measure skill: it is not invariant with respect to leverage. That is, imagine a hedge fund with return r , standard deviation σ , and correlation with the index ρ . It has beta $\beta = \sigma\rho/\sigma_m$ and alpha $\alpha = r - \beta r_m$ where r_m and σ_m are respectively the return and standard deviation of the index. Now imagine doubling the fund's internal leverage. This leaves ρ unchanged, but doubles r and σ , so the new beta is $\beta' = 2\sigma\rho/\sigma_m$, or twice the old beta. Likewise, the new alpha is $\alpha' = r' - \beta' r_m = 2r - 2\beta r_m = 2(r - \beta r_m) = 2\alpha$, or twice the old alpha. However, the hedge fund manager's skill is a function of his ability to make good trades, not the leverage of his fund. So alpha can not be a measure of skill.

The Treynor-Black appraisal ratio, which is alpha divided by the residual standard error of the CAPM regression, is leverage-invariant. To correct for bias, Park (1995) showed that Treynor-Black winners as well as alpha winners tend to repeat. However, the practical significance of this result is dubious: it is not clear that it is always best to have a high Treynor-Black appraisal ratio. As Lehmann and Modest (1987) point out, the appraisal ratio is

proportional to the t-statistic for the intercept in the CAPM regression. So the Treynor-Black ratio gives information about the statistical significance of a trader's excess returns, not about the practical significance. A trader who scores highly on this measure may have tracked the index closely and outperformed it slightly. On the other hand, those who outperformed the index greatly on a risk-adjusted basis, but whose returns were poorly explained by the index, may have a low score. To oversimplify slightly, if the Treynor-Black appraisal ratio persists, it measures the estimated probability of outperforming the index, not the expected amount of excess return.

Instead we advocate using the ratio α/σ . The risk-free rate is not subtracted as in a Sharpe ratio; subtracting the risk-free rate makes the Sharpe ratio vary with leverage and renders it unsuitable for describing an industry in which internal leverage is possible. For the same reason, it is important that alpha and beta are calculated from a regression of returns not adjusted for the risk-free rate. This measure is leverage-invariant and is interpretable as the ratio of excess reward to risk.

Alternatively, one may consider dividing α by σ as parallel to the ideas of Schneeweis, Spurgin, and McCarthy (1997). They argued that the primary informational content of a CTA's historical performance is the beta relative to an industry index. Then an estimate of future performance is β times the index's return. However, the authors warned that historical data tends to overestimate the absolute value of β , and suggested using Bayesian techniques to deal with this problem.

The problem is essentially one of regression to the mean. When estimated β is extreme, this is probably because the true β is large, but not as extreme as the estimate, with the difference being the random estimation error. Of course, the same is true for α . If the distributions of α

were the same for all traders, there would be no need for Bayesian techniques. The solution would be simply to use empirical data to regress future α on estimated α to find a constant deflation factor. Unfortunately, as mentioned before, α is not leverage-invariant. In particular, such a regression is inappropriate because the conditional variance of α given σ is proportional to σ . This makes σ a good deflation factor, considering the attractive properties of α/σ .

This measure is also useful from the standpoint of portfolio construction. If a fund of funds manager has a broadly diversified portfolio of funds, it should have a correlation of close to one with the index of the fund universe from which he selects. Therefore α/σ roughly expresses how much a trader contributes to the portfolio's excess return over the market per unit of risk. A portfolio composed of traders with high skill statistics should deliver a risk-adjusted return superior to the market.

IV. Evidence

The study concentrates on the relationship between a trader's lifetime skill rank and his skill rank over one future year. For instance, in December 1990, calculate the skill statistic for all traders whose lifetime history to date is at least 12 months long, then rank them on this statistic. Next calculate α/σ from α and σ based on only data from 1991, and rank again. This is done for each of the years 1991 to 1997. Testing these pairs of lifetime and outcome skill statistics involves implicitly conditioning on survival over the future year after the lifetime skill is assessed. However, the final test of practical significance (see below) shows that the effect is real and not an artifact of the "look-ahead" bias described by Carhart (1997).

A. Hypothesis Tests

Nonparametric statistical tests are appropriate because the data suffers badly from right-skewness, heteroscedasticity, and outliers, although, as mentioned above, the use of the logarithmic return reduces these effects somewhat. For testing an alternative hypothesis of positive association between prior and outcome skill, a good nonparametric procedure is the Spearman rank correlation. The correlations are extremely strong for hedge funds, with p-values off the scale. This shows that a trader's skill relative to his peers tends to persist. The same is true for CTAs, but to a lesser extent. The strength of the persistence seems to vary substantially from year to year.

This annual variability of persistence resembles that reported by Brown and Goetzmann (1995), and is present in all tests and simulations in this study. In particular, the years 1994 and 1995 show the weakest effect, and in some cases a reversal of effect. Because 1994 was an extremely poor year for these alternative investments industries, this result is strikingly coincident with Brown and Goetzmann's finding of a persistence reversal among mutual funds in 1988, following the crash of 1987. This is not surprising: managers who prosper while their peers suffer from a rare event (e.g. stock market crash) are unlikely to be those who have the skills suited to beating the pack during periods of normal market conditions.

Next is a 2 by 2 chi-squared test, where one dimension is the trader's skill in the outcome year, and the other dimension is his skill over his lifetime prior to the outcome year. The null hypothesis is that past performance is of no use in predicting future results, so that traders' skill in prior and outcome periods are independent. Therefore it is legitimate to sum the tables to get an

aggregate result. The aggregate chi-squared tests powerfully reject this null hypothesis, even though the test has very low power due to discretization into just two groups.

The 2 by 2 chi-squared test is easy to interpret because it corresponds to ranking the traders into winners and losers. This allows some measure of the practical significance of the results. For instance, over the course of the period under consideration, if a money manager picked a hedge fund which was a winner over its lifetime to date, he had a 67.5% chance of having this choice be a winner in the next year. This is a powerful result: if a fund of funds manager can create a portfolio two-thirds of whose traders are in the top half of all traders, he will be ahead of the game. Even the 56.7% chance of choosing persistent winners among CTAs is enough to give a portfolio manager an edge. These results are not the sort where a large sample size turns a minuscule difference in behavior into a giant statistical effect.

It is interesting to investigate exactly how these effects manifest themselves. Grouping the traders into deciles shows how their relative ranks tend to evolve over time. Note that neither hedge funds nor CTAs exhibit the “J” shape of persistence which would be caused if the effect were entirely due to the form of survivorship bias studied by Hendricks, Patel, and Zeckhauser (1997). Hedge funds show an almost uniform increasing trend: those who have shown the least skill in the past tend to show the least in the future, and those who have been most skillful continue to be so. On the other hand, behavior is different among CTAs. Those with historical skill in deciles 0 through 6 (the bottom 70%) have on average a very similar future performance. That is, the prospects of a CTA whose historical skill was among the very worst are no poorer than those of a CTA whose historical skill was mediocre. Roughly, there are three tiers of CTAs: the bottom 70%, with poor average future performance, the top decile, with the best prospects,

and the remainder (deciles 7 and 8), with intermediate expected future skill.

This accounts for some of the weakness of the results for CTAs as compared to hedge funds. In a sense, a CTA's historical skill is not as informative as a hedge fund's. The expected future performance of a hedge fund varies more or less continuously with historical skill, so it is possible to identify groups with very low or very high expected future performance. On the other hand, it is impossible to identify a group of CTAs with very low average future skill. While the statistical results are thus far weaker for CTAs, not as much is lost in practical benefit. The goal is presumably to identify a group of CTAs with high expected future performance, and this is possible, albeit to a lesser extent than for hedge funds. The top decile of hedge funds (based on past skill) has expected future decile 7.38 or rank 78.8%, while the top decile for CTAs has expected future decile 6.18 or rank 66.8%.

B. Simulations

The most convincing demonstration of practical significance is showing that a portfolio of traders with historical α/σ in the top decile actually has a future risk/reward profile superior to a random or market portfolio. Here it is possible to avoid the complication of conditioning on survival mentioned above. The "elite" portfolio's monthly returns are the average of the returns of all surviving traders whose historical skill was in the top decile at the beginning of this year. In effect, the elite fund maintains an equal-weighted portfolio by redistributing capital from any trader who exits the database to all others, equally.

The contrast of the elite portfolio with the market portfolio is very favorable. Under the

null hypothesis of no skill effect, both portfolios would have the same expected returns, but the elite portfolio would have greater standard deviation than the market portfolio because it contains one tenth as many traders. However, just the opposite occurs. The elite portfolio always has much lower standard deviation than the market portfolio. At the same time, the elite portfolio's return is good enough that the ratio of monthly return to monthly standard deviation is almost always much, much greater than that of the market portfolio. As usual, the effect is stronger for hedge funds than for CTAs. Nonetheless, for both fund categories it is possible to construct a portfolio of skilled traders which will outperform the market index on a risk-adjusted basis.

V. Directions for Future Research

Clearly, there is room for much more research in what appears to be a very promising area for statistical modeling, with the potential to deliver great benefits to alternative investments industries. Perhaps the most important is to separate the effects of skill and style. If relatively young trading programs of different styles (e.g. FX trend-follower and emerging markets hedge fund) are analyzed together, one may have a much higher skill statistic than the other primarily because the market environment has recently been more favorable for his style. Style should not be confounded with skill, because the relative performance of styles is not very predictable.

A trader's skill should properly be measured only against his peers, others of the same style. As discussed in Brown and Goetzmann (1997), a major purpose of grouping traders into styles is making a trader's future returns highly correlated with an index constructed from all traders of his style. Fung and Hsieh (1997a) took an important step towards achieving this goal

for CTA funds and hedge funds, while acknowledging that their five style factors account for less than half of the variance among these funds. On the other hand, Fung and Hsieh (1997b) suggested that there is only one identifiable style among CTAs. Schneeweis and Spurgin (1996) showed that among both CTAs and hedge funds, loadings on explanatory factors vary. They recommended using daily data to determine styles among traders who use different strategies to profit from moves in various markets.

If traders can be grouped into styles which cluster more tightly, their excess returns relative to their style index will be a purer estimate of skill. Schneeweis, Spurgin, and McCarthy (1997) concurred that it is important to use a suitable index to forecast a trader's performance. Once this is done, a trader's skill when compared only to his peers will presumably exhibit even greater persistence than is evident in this study. Consequently, asset allocators can adopt a strategy of first deciding the style composition of their portfolios, and then making allocations within style on the basis of skill.

Another potential advance would be a forecasting model for a skill statistic or its rank. By examining aggregate behavior, this study has shown that the conditional distribution of future skill rank on lifetime skill rank does vary with the lifetime skill rank. A model providing useful point estimates or confidence intervals for a future skill statistic or rank is lacking. The ability to forecast future skill numerically would be particularly useful in the context of the framework for portfolio construction based on alpha, beta, and the correlations between traders. This resembles classical Markowitz efficient frontier theory, except that risk and return are both expressed not in absolute terms, but relative to an index.

VI. Conclusion

The evidence is strong that skill is a factor in CTA performance and a very important factor in hedge fund performance. It is possible for fund of fund managers to assess traders' skill and allocate assets accordingly in order to generate higher expected returns. Based on current research, we propose measuring skill with the ratio α/σ , which is interpretable, has some good theoretical properties, and shows significant persistence in practice.

This conclusion contrasts with that of Goetzmann, Ibbotson, and Brown (1998), who found no evidence for performance persistence among offshore hedge funds. This disagreement doubtless arises because of the different data sources used. Goetzmann, Ibbotson, and Brown used annual returns from a data set free of survivorship bias. While this makes the results of their study more reliable, it also allowed less scope for detailed study of the phenomena in question. In particular, their results about alpha and appraisal ratio are not based on regressions or the volatility of each fund, but only on differences in annual returns. These results confirm the findings of Park (1995), who showed that there is very little (if any) predictability to unscaled returns among CTAs and hedge funds.

On the other hand, this paper has argued that when skill is measured with appropriate risk adjustment, useful predictions are possible. This risk adjustment involves regressions and longitudinal standard deviations, for which it is best to have more than seven annual data points. It is impossible to evaluate the validity of this hypothesis about monthly return streams with annual data. While the subtle issues of survivorship bias have not been entirely resolved here, neither does the bias-free work of Goetzmann, Ibbotson, and Brown strictly contradict these

results.

If the above-mentioned research goals are attained, fund managers in alternative investment industries will have a powerful new tool for portfolio construction. There will be the potential for more science and less guesswork in asset allocation. Moreover, fund management will be able to demonstrate that it adds value by allocating assets to the most skillful traders. In turn, this means that managed alternative investments will be producing greater economic benefit, by channeling capital to those traders who use it to produce the best pricing information.

REFERENCES

Brown, Stephen, and William Goetzmann, 1995, "Performance Persistence," *Journal of Finance* 50:2, 679-698.

Brown, Stephen, and William Goetzmann, 1997, "Mutual Fund Styles," *Journal of Financial Economics* 43, 373-399.

Brown, Stephen, William Goetzmann, Roger Ibbotson, and Stephen Ross, 1992, "Survivorship Bias in Performance Studies," *Review of Financial Studies* 5:4, 553-580.

Brown, Stephen, William Goetzmann, Roger Ibbotson, and Stephen Ross, 1997, "Rejoinder: The J-Shape of Performance Persistence Given Survivorship Bias," *Review of Economics and Statistics* 74:2, 167-170.

Brown, Stephen, William Goetzmann, and James Park, 1997, "Conditions for Survival: Changing Risk and the Performance of Hedge Fund Managers and CTAs," working paper: www.fen.com.

Brown, Stephen, William Goetzmann, and Stephen Ross, 1995, "Survival," *Journal of Finance* 50:3, 853-873.

Carhart, Mark, 1997, "Mutual Fund Survivorship," working paper: www.fen.com.

Fung, William, and David Hsieh, 1997a, "Empirical Characteristics of Dynamic Trading Strategies: The

Case of Hedge Funds,” *Review of Financial Studies* 10:2, 275-302.

Fung, William, and David Hsieh, 1997b, “Survivorship Bias and Investment Style in the Returns of CTAs,” *Journal of Portfolio Management*, 30-41.

Goetzmann, William, and Roger Ibbotson, 1994, “Do Winners Repeat? Patterns in Mutual Fund Performance,” *Journal of Portfolio Management* 20, 9-17.

Goetzmann, William, Roger Ibbotson, and Stephen Brown, 1998, “Offshore Hedge Funds: Survival & Performance 1989 - 1995,” working paper: www.fen.com.

Hendricks, Darryll, Jayendu Patel, and Richard Zeckhauser, 1993, “Hot Hands in Mutual Funds: The Persistence of Performance, 1974-1988,” *Journal of Finance* 48, 93-130.

Hendricks, Darryll, Jayendu Patel, and Richard Zeckhauser, 1997, “The J-Shape of Performance Persistence Given Survivorship Bias,” *Review of Economics and Statistics* 74:2, 161-166.

Lehmann, Bruce, and David Modest, 1987, “Mutual Fund Performance Evaluation: A Comparison of Benchmarks and Benchmark Comparisons,” *Journal of Finance* 42:2, 223-265.

Park, James, 1995, *Managed Futures as an Investment Class Asset* (UMI Dissertation Services, Ann Arbor, Michigan).

Schneeweis, Thomas, and Richard Spurgin, 1996, “Multi-Factor Models in CTA/Hedge Fund Return

Estimation,” working paper: www.umass.edu/som/cisdms.

Schneeweis, Thomas, Richard Spurgin, and David McCarthy, 1997, “Informational Content in Historical CTA Performance,” *Journal of Futures Markets*.

Table 1

Lifetime Skill: Spearman Correlation

Outcome Period	Correlation	P-value		Correlation	P-value	
	Hedge Funds			CTAs		
Aggregate	0.505	0.0001	**	0.232	0.0001	**
1991	0.602	0.0001	**	0.253	0.0067	**
1992	0.745	0.0001	**	0.352	0.0001	**
1993	0.691	0.0001	**	0.313	0.0001	**
1994	0.232	0.0002	**	0.260	0.0007	**
1995	0.380	0.0001	**	0.041	0.5922	
1996	0.532	0.0001	**	0.231	0.0021	**
1997	0.547	0.0001	**	0.102	0.1995	

Table 2

Lifetime Skill: 2x2 Chi-Squared

Outcome Period	Chi-Squared	P-value		P(win win)	Chi-Squared	P-value		P(win win)
	Hedge Funds				CTAs			
Aggregate	218.5	1.9E-50	**	0.675	18.8	9.3E-06	**	0.567
1991	38.2	2.9E-10	**	0.826	4.2	0.030	*	0.596
1992	46.6	3.3E-12	**	0.806	4.4	0.026	*	0.592
1993	59.9	2.9E-15	**	0.781	4.3	0.028	*	0.582
1994	2.1	0.092		0.548	4.1	0.031	*	0.578
1995	32.6	8.4E-09	**	0.657	0.8	0.223		0.535
1996	37.8	5.6E-10	**	0.650	2.1	0.099		0.557
1997	44.8	1.4E-11	**	0.672	1.0	0.193		0.543

Table 3

Future Performance of Top Hedge Fund Decile

Outcome Period	Top Decile			Market		
	Return	Std. Dev.	Ratio	Return	Std. Dev.	Ratio
1991	1.03	0.32	3.3	1.58	3.51	0.5
1992	1.25	0.48	2.6	0.91	1.66	0.5
1993	1.42	0.53	2.7	1.61	1.44	1.1
1994	0.21	0.92	0.2	-0.09	1.10	-0.1
1995	0.85	0.67	1.3	1.14	1.26	0.9
1996	1.33	0.57	2.3	1.07	1.78	0.6
1997	0.94	0.83	1.1	1.09	2.12	0.5

Table 4

Future Performance of Top CTA Decile

Outcome Period	Top Decile			Market		
	Return	Std. Dev.	Ratio	Return	Std. Dev.	Ratio
1991	0.98	1.13	0.9	0.81	4.95	0.2
1992	0.59	0.81	0.7	0.29	3.68	0.1
1993	1.22	1.44	0.8	1.11	2.82	0.4
1994	0.82	0.94	0.9	-0.11	2.61	-0.0
1995	-0.06	1.25	-0.0	1.23	3.06	0.4
1996	0.58	1.19	0.5	1.00	3.77	0.3
1997	0.33	0.71	0.5	0.90	2.75	0.3

Figure I

Skill: Lifetime vs. Outcome

