Attack of the Clones

By Jasmina Hasanhodzic and Andrew W. Lo

Hedge funds are considered by many investors to be an attractive investment, thanks in large part to their diversification benefits and distinctive risk profiles. The major drawbacks are their high fees and lack of transparency. Research by Jasmina Hasanhodzic and Andrew W. Lo of the Massachusetts Institute of Technology raises the possibility of creating passive portfolios that provide similar risk exposures to those of hedge funds at lower costs and with greater transparency. Hasanhodzic and Lo find that for certain hedge fund strategies, these hedge fund "clones" perform well enough to warrant serious consideration.

As institutional investors take a more active interest in alternative investments, a significant gap has emerged between the culture and expectations of those investors and hedge fund managers. Pension plan sponsors typically require transparency from their managers and impose numerous restrictions on their investment mandates because of regulatory requirements such as ERISA rules; hedge fund managers rarely provide position-level transparency and bristle at any restrictions on their investment process, saying that restrictions can hurt performance. Plan sponsors require a certain degree of liquidity in their assets to meet their pension obligations and also desire significant capacity because of their limited resources in managing large pools of assets; hedge fund managers routinely impose lockups of one to three years, and the most successful managers have the least capacity to offer, in many cases returning investor capital once they make their personal fortunes. And as fiduciaries, plan sponsors are hypersensitive to the outsize fees that hedge funds charge and are concerned about misaligned incentives induced by performance fees; hedge fund managers argue that their fees are fair compensation for their

Illustrations by Edel Rodriguez for Alpha

unique investment acumen — and at least for now, the market seems to agree.

This cultural gap raises the natural question of whether it is possible to obtain hedge-fund-like returns without investing in hedge funds. In other words, can hedge fund returns be cloned?

In a series of recent papers, Harry Kat and Helder Palaro of the Cass Business School at City University in London show that sophisticated dynamic trading strategies involving liquid futures contracts can replicate many of the statistical properties of hedge fund returns. In fact, in a 2001 paper with Dimitris Bertsimas and Leonid Kogan of the Massachusetts Institute of Technology, we show that the risk/return characteristics of securities with very general payoff functions (like hedge funds or complex derivatives) can be synthetically replicated to an arbitrary degree of accuracy by dynamic trading strategies — called *epsilon-arbitrage* strategies — involving more liquid instruments. Although these results are encouraging for the hedge fund replication problem, the strategies are quite complex and not easily implemented by the typical institutional investor.

Table 1: Capital Decimation Partners*

The monthly returns of fictitious Capital Decimation Partners' simulated short-put-option strategy handily beat the Standard & Poor's 500 index.

Statistic	Standard & Poor's 500 index	Capital Decimation Partners
Monthly mean	1.4%	3.6%
Monthly standard deviation	3.6%	5.8%
Minimum month	–8.9%	–18.3%
Maximum month	14.0%	27.0%
Annual Sharpe ratio	0.98	1.90
Number of negative months	36	6
Correlation to S&P 500 index	100%	61%
Growth of $1 since inception	**$4**	**$26**

** January 1992 through December 1999.* *Source: Lo (2001).*

In this article we take a slightly different tack: We construct "linear clones" — buy-and-hold portfolios of common risk factors like the Standard & Poor's 500 and U.S. dollar indexes, with portfolio weights estimated by a linear regression of a fund's historical returns on market factors — of a large number of individual hedge funds in the TASS Hedge Fund Database. We then compare their characteristics to those of the corresponding funds from which the clones are derived.

If a hedge fund generates part of its expected return and risk profile from certain common risk factors, it may be possible to design a low-cost, buy-and-hold portfolio — not an active, dynamic trading strategy — that captures some of that fund's risk/reward characteristics by taking on just those risk exposures. For example, if a particular long-short equity hedge fund is 40 percent long growth stocks, it may be possible to create a passive portfolio that has similar characteristics through a long-only position in a passive growth portfolio coupled with a 60 percent short position in stock index futures.

The magnitude of hedge fund alpha that can be captured by a linear clone depends, of course, on how much of a fund's expected return is driven by common risk factors versus manager-specific alpha. This can be measured empirically. Although portable-alpha strategies have become fashionable lately among institutions, our research suggests that for certain classes of hedge fund strategies, portable beta may be an even more important source of untapped expected returns and diversification.

BEFORE TURNING TO OUR empirical analysis, we provide two concrete examples that span the extremes of the hedge fund replication problem. For one hedge fund strategy, we show that replication can be accomplished easily; for another strategy we find replication to be almost impossible using linear models.

The first example is a hypothetical strategy we proposed several years ago called "Capital Decimation Partners," or CDP, which yields an enviable track record that many investors would associate with a successful hedge fund: a 43.1 percent annualized mean return and 20.0 percent annualized volatility, implying a Sharpe ratio of 1.90, and with only six negative months over the 96-month simulation period from January 1992 through December 1999 (see Table 1).

So what is CDP's secret? The investment strategy involves shorting out-of-the-money S&P 500 put options on each monthly expiration date for maturities less than or equal to three months, and with strikes approximately 7 percent out of the money.

The essence of this strategy is the provision of insurance. CDP investors receive option premiums for each option contract sold short, and as long as the option contracts expire out of the money, no payments are necessary. From this perspective the handsome returns to CDP investors seem more justifiable: In exchange for providing downside protection, CDP investors are paid a risk premium in the same way that insurance companies receive regular payments for providing earthquake or hurricane insurance.

Given the relatively infrequent nature of 7 percent losses, CDP's risk/reward profile can seem very attractive in comparison to more traditional investments, but there is nothing unusual or unique about CDP. Investors willing to take on "tail risk" — the risk of rare but severe events — will be paid well for this service (consider how much individuals are willing to pay each month for their homeowner's, auto, health and life insurance policies). CDP involves few proprietary elements and can easily be implemented by most investors; it is one example of a hedge-fund-like strategy that can readily be cloned.

Now for the bad news. Consider the case of "Capital Multiplication Partners," or CMP, a hypothetical fund based on a dynamic asset-allocation strategy between the S&P 500 and one-month U.S. Treasury bills, where the fund manager can correctly forecast which of the two vehicles will do better in each month and invests the fund's assets in the higher-yielding instrument at the start of the month. (This example was first proposed by Robert Merton in his 15.415 Finance Theory class at the MIT Sloan School of Management in the 1970s.) The monthly return of this perfect-market-timing strategy is simply the larger of the monthly returns of the S&P 500 and T-bills.

The source of alpha is clear. Merton observes that this strategy is equivalent to a long-only investment in the S&P 500 plus a put option on the S&P 500 with a strike price equal to the T-bill return. The economic value of this perfect market-timing is equal to the sum of monthly put-option premiums over the life of the strategy.

There is little doubt that such a strategy contains significant alpha indeed: A $1 investment in CMP in January 1926 would have grown to more than $23 billion by the end of December 2004! Table 2 provides a more detailed performance summary of CMP, whose Sharpe ratio exceeds that of Warren Buffett's Berkshire Hathaway, arguably the most successful pooled investment vehicle of all time.

It should be obvious to even the most naive investor that CMP is a fantasy because no one can time the market perfectly. Therefore, attempting to replicate such a strategy with exchange-traded instruments seems hopeless. But suppose we try anyway. How close can we come? In particular, suppose we attempt to relate CMP's monthly returns to the monthly returns of the S&P 500 by fitting a straight line through a graph of their paired monthly returns, that is, a linear regression. The option-like nature of CMP's perfect-market-timing strategy, which is inherently nonlinear, cannot be captured by a straight line. However, the formal statistical measure of how well the linear regression fits the data — the R^2, a number between 0 and 100 percent that implies a perfect linear relationship at 100 percent and no relationship at all at 0 — is 70.3 percent in this case, which suggests a very strong linear relationship indeed. But when the estimated linear regression is used to construct a buy-and-hold portfolio of the S&P 500 and one-month T-bills, the results are not nearly as impressive as CMP's returns, as Table 2 shows.

This example underscores the difficulty of replicating certain strategies that have genuine alpha with linear clones, and it cautions against using the R^2 as the only metric of success. Despite the high R^2 achieved by the linear regression of CMP's returns on the market index, the actual performance of the linear clone falls far short of the strategy because a linear model will never be able to capture the option-like payoff structure of the perfect market-timer.

TO EXPLORE THE FULL RANGE of possibilities for replicating hedge fund returns illustrated by the two extremes of CDP and CMP, we investigate the performance of linear clones for a sample of individual hedge funds drawn from the TASS Hedge Fund Live Database over the sample period from February 1986 through September 2005. We start our analysis in February 1986 because this is the earliest date for which we have complete data for all of our factors (the TASS database goes back to 1977). Of these funds, we drop those that do not report net-of-fee returns, those that report returns in currencies other than the U.S. dollar, those that report returns less frequently than monthly, those that do not provide assets under management or provide only estimates and those that have fewer than 36 monthly returns. These filters yield a final sample of 1,610 funds.

For each fund we estimate a linear regression of its monthly historical returns on the following six risk factors: the U.S. dollar index return, the return on the Lehman Brothers corporate AA intermediate bond index, the spread between the Lehman BAA corporate bond index and the Lehman Treasury index, the S&P 500 total return, the Goldman Sachs commodity index return and the first-difference of the end-of-month value of the VIX Chicago Board Options Exchange volatility index, or DVIX. (Throughout this article all statistics, except for those related to the first-order autocorrelation, have been annualized to facilitate interpretation and comparison.)

We choose these six risk factors for two reasons: They provide a reasonably broad cross-section of risk exposures for the typical hedge fund (stocks, bonds, currencies, commodities, credit and volatility), and each of the factor returns can be realized through relatively liquid instruments

Table 2: Capital Multiplication Partners*

Based on a series of simulated monthly returns going back to 1926, the aptly named Capital Multiplication Partners' perfect-market-timing strategy easily beats its clone's.

Statistic	Standard & Poor's 500 index	Treasury bills	Capital Multiplication Partners	Clone
Monthly mean	1.0%	0.3%	**2.6%**	**0.7%**
Monthly standard deviation	5.5%	0.3%	**3.6%**	**3.0%**
Minimum month	–29.7%	–0.1%	**–0.1%**	**–16.3%**
Maximum month	42.6%	1.4%	**42.6%**	**23.4%**
Annual Sharpe ratio	0.63	4.12	**2.50**	**0.79**
Number of negative months	360	12	**10**	**340**
Correlation to S&P 500 index	100%	–2%	**84%**	**100%**
Growth of $1 since inception	$3,098	$18	**$2.3 x 10^{10}**	**$429**

** January 1926 through December 2004.* *Source: Authors' calculations.*

so that the returns of linear clones may be achievable in practice. In particular, there are forward contracts for each of the component currencies of the U.S. dollar index and futures contracts for the stock and bond indexes and for the components of the commodity index. Futures contracts on the VIX index were introduced by the CBOE in March 2004 and are not as liquid as the other index futures, but the over-the-counter market for variance and volatility swaps is quite well developed.

The linear-regression model provides a simple but useful decomposition of a hedge fund's expected return into two distinct components: beta coefficients multiplied by the risk premiums associated with various risk factors, and manager-specific alpha. The intuition for this decomposition is straightforward. Hedge funds generate their expected returns by taking on certain generic risks for which they are compensated, like market or credit risk, and also by taking advantage of insights and opportunities that are specific to the manager.

By "manager-specific alpha," we do not mean to imply that a hedge fund's unique source of alpha is without risk. We are simply distinguishing this source of expected return from those that have clearly identifiable risk factors associated with them. In particular, it may well be the case that manager-specific alpha arises from factors other than the six we have proposed, and a more refined list of factors — one that reflects the particular investment style of the manager — may yield a better-performing linear clone.

> *"For certain types of hedge fund strategies, a passive buy-and-hold approach may yield some of the same benefits as hedge funds."*
>
> — JASMINA HASANHODZIC AND ANDREW W. LO

A similar decomposition for a hedge fund's return variance can be derived that is the sum of three distinct components: the variances of the risk factors multiplied by the squared beta coefficients, the variance of the fund-specific sources of randomness or "residual" (which may be related to the specific economic sources of alpha) and the weighted covariances among the factors. This decomposition highlights the fact that a hedge fund can have several sources of risk, each of which should yield some risk premium — that is, risk-based alpha — otherwise, investors would not be willing to bear such risk. By taking on exposure to multiple risk factors, a hedge fund can generate attractive expected returns from the investor's perspective, as we saw with Capital Decimation Partners.

Using the linear-regression model to decompose a fund's expected returns, we can now reformulate the question of whether a hedge fund strategy can be cloned by asking how much of a hedge fund's alpha is due to risk premiums from identifiable factors. If it is a significant portion, then a passive portfolio with just those risk exposures — created by means of liquid instruments such as index futures, forwards and other contracts — may be a reasonable alternative to a direct investment in the fund.

Table 3 summarizes the empirical results of the expected-return decomposition for our sample of funds, grouped according to their style categories. Each row contains the average total mean return of funds in a given category, and averages of the percent contributions of each of the six factors and the manager-specific alpha to that average total mean return. For example, the most significant contributors to the investment return of convertible-arbitrage funds are the dollar index (67.1 percent), the bond index (34.9 percent), the commodity index (31.8 percent) and the credit spread (27.1 percent), and the average contribution of manager-specific alpha is –33.3 percent.

This implies that convertible-arbitrage funds, on average, earn more than all of their mean returns from the risk premiums associated with the six factor exposures, and that the average contribution of other sources of alpha is negative. Of course, this does not mean that convertible-arbitrage managers are not adding value. The results are averages across all funds in the sample; hence, the positive manager-specific alphas of successful managers will be dampened and, in some cases, outweighed by the negative manager-specific alphas of the unsuccessful ones. Moreover, all of the statistics reported in our study are estimates only and therefore subject to a certain amount of estimation error.

In contrast to the convertible-arbitrage funds, for the ten funds in the dedicated short-bias category, manager-specific alpha accounts for 225.6 percent of the total mean return, while the contribution of the S&P 500 factor is negative. This result is not as anomalous as it may seem. The bull market of the 1990s implies a performance drag for any fund with negative exposure to the S&P 500. Thus, dedicated short-bias managers that have generated positive performance during this period must have done so through other means.

A concrete illustration of this intuition is given by the decomposition of the annualized average return of the two most successful funds in the dedicated short-bias category. From 1997 through 2005 these two funds posted annualized net-of-fee returns of 15.56 percent and 10.02 percent, respectively, but the contribution of the S&P 500 factor to these annualized returns was negative in both cases. In fact, the six factors account for very little of the two funds' performance; hence, the manager-specific alphas are particularly significant for these two funds.

Between the two extremes of convertible-arbitrage and dedicated short-bias funds, there is considerable variation in the importance of manager-specific alpha for the other strategy categories. For the entire sample of 1,610 funds, 61.0 percent of the average total return is attributable to manager-specific alpha, implying that, on average, the remaining 39.0 percent is due to the risk premiums from our six factors. These results suggest that for certain types of hedge fund strategies, a passive buy-and-hold approach may yield some of the same benefits as hedge funds, but in a transparent, scalable and lower-cost vehicle.

HOW CLOSE CAN WE COME to replicating hedge fund returns? To answer this question, we construct linear clones of each fund in our sample by regressing the fund's returns on five of the six factors we considered above (we drop the DVIX volatility factor because its returns are not as easily realized with liquid instruments) and no intercept, and then rescaling the fitted regression equation so that the resulting buy-and-hold portfolio has the same sample volatility as the original fund's return series. We omit the intercept because our objective is to estimate a weighted average of the factors that best replicates the fund's returns. The motivation for rescaling the volatility of the clones is to create a fair comparison between the buy-and-hold portfolio and the fund, and is equivalent to changing the leverage of the clone portfolio.

Table 4 contains a comparison of the performance of these linear clones and that of the original funds from which the clones are derived. The results are striking — for several strategy categories the average mean return of the clones is only slightly lower than that of their fund counterparts, and in some categories the clones outperform. For example, the average mean return of the convertible-arbitrage clones is 8.15 percent, and the corresponding figure for the actual funds is 8.41 percent. For long-short equity hedge funds, the average mean return for clones and funds is 13.94 and 14.59 percent, respectively. And in the multistrategy category, the average mean return for clones and funds is 10.10 and 10.79 percent, respectively.

In three cases the average mean return of the clones is higher than that of the funds: global macro (14.43 percent versus 11.38 percent), managed futures (23.47 percent versus 13.64 percent) and fund of funds (8.63 percent versus 8.25 percent). However, these differences are not statistically significant because of the variability in mean returns across funds within each category. Even in the case of managed futures, the difference in average mean return between clones and funds — almost 10 percentage points — is not statistically significant because of the large fluctuations in average mean returns of the managed-futures clones. Nevertheless, these results suggest that for certain categories, the performance of clones may be within shouting distance of their corresponding funds.

One category of hedge funds that seems particularly difficult to replicate is event-driven strategies. The average performance of the event-driven clones, at 9.60 percent, is considerably lower than the 13.03 percent average for the event-driven funds. This large gap is understandable, given the idiosyncratic and opportunistic nature of most event-driven strategies. Moreover, a significant source of the profitability of event-driven strategies is the illiquidity premium that managers earn through their willingness to provide capital in times of distress. This illiquidity premium will clearly be missing from a clone portfolio of liquid securities; therefore, we should expect a significant performance gap in this case.

For dedicated short-bias funds, the average mean return of the clones and the funds is 3.58 and 5.98 percent, respectively. This may seem somewhat counterintuitive in light of the expected-return decomposition in Table 3, where we observed that dedicated short-bias funds were responsible for more than 100 percent of the average total returns of funds in this category. The fact that dedicated short-bias clones have positive average performance is due entirely to the clone of a single fund, No. 33735 in the TASS database, and when this outlier is dropped from the sample, the average mean return of the remaining nine clones drops to –0.35 percent. The underperformance of the clones in this category is also intuitive — given the positive trend in the U.S. stock market during the 1980s and '90s, a passive strategy of shorting the S&P 500 is unlikely to have produced attractive returns when compared to the performance of more nimble discretionary short-sellers.

Another metric of comparison is the average Sharpe ratio, which adjusts for the volatilities of the respective strategies. Given our rescaling process, the standard deviations for the clones are identical to their fund counterparts, so a comparison of Sharpe ratios reduces to a comparison of mean returns. However, the average Sharpe ratio of a category is not the same as the ratio of that category's average mean return to its average volatility, so the Sharpe ratio statistics in Table 4

Table 3: Breaking Down Hedge Fund Returns by Strategy

Not all alpha is created equal. An analysis of the total mean returns for more than 1,600 hedge funds in the TASS database from February 1986 to September 2005 shows which asset classes and factors make the biggest contribution to their investment performance.

Category description	No. of funds	Average expected return (%)	Average of percentage contribution of factors to total expected return (%)						
			Credit spread	Dollar index	S&P 500 index	Bond index	VIX index	Commodity index	Alpha
Convertible arbitrage	82	8.4%	27.1%	67.1%	–19.3%	34.9%	–8.4%	31.8%	–33.3%
Dedicated short-bias	10	6.0	12.2	19.4	–108.2	7.0	8.9	–64.9	225.6
Emerging markets	102	4.9	–0.3	–3.2	19.3	0.1	–0.4	6.2	78.3
Equity market-neutral	83	20.4	0.2	3.6	4.0	3.9	1.3	6.3	80.8
Event-driven	169	8.1	2.1	3.0	4.3	9.4	–0.7	3.1	79.0
Fixed-income arbitrage	62	13.0	–1.4	3.3	2.7	18.5	–0.5	4.4	73.1
Global macro	54	9.5	2.0	8.1	9.7	25.0	–3.3	10.0	48.6
Long-short equity hedge	520	11.4	1.1	1.9	17.8	2.1	–1.8	8.4	70.5
Managed futures	114	14.6	1.9	23.4	–3.4	53.8	–1.5	53.2	–27.5
Multistrategy	59	13.6	0.5	3.5	5.7	10.1	–1.9	3.2	78.9
Fund of funds	355	10.8	0.5	5.4	9.7	8.8	–2.8	7.3	71.1
All funds	1,610	8.3	2.3	7.8	8.5	11.3	–1.9	10.9	61.0

Source: Authors' calculations.

do provide some incremental information. The average Sharpe ratio of the funds in the convertible-arbitrage category is 2.70, which is almost twice the average Sharpe ratio of 1.54 for the clones, a significant risk-adjusted performance gap between the funds and their clones. However, there is virtually no difference in average Sharpe ratios between clones and funds for equity market-neutral, long-short equity hedge and fund-of-funds categories. As we discussed above, the apparent similarity of dedicated short-bias clones to their funds is the result of a single outlier. And for global macro and managed futures, the average Sharpe ratios of the clones are, in fact, higher than those of the funds.

Table 4 provides one more comparison worth noting: the average first-order autocorrelation coefficients of clones and funds. The first-order autocorrelation, $\hat{\rho}_1$, is the correlation between a fund's current return and the previous month's return, and in our previous studies we show that a positive value for $\hat{\rho}_1$ in hedge fund returns is a proxy for illiquidity risk. The clones have much lower average autocorrelations than their fund counterparts, with the exception of the managed-futures category, for which both clones and funds have very low average autocorrelations. For example, the average autocorrelation of convertible-arbitrage funds is 42.2 percent, and the corresponding average value for convertible-arbitrage clones is only 10.4 percent. A more formal statistical analysis shows that for every single category the average level of autocorrelation in the funds is higher than that in the clones, confirming our intuition that, by construction, clones are more liquid than their fund counterparts.

A PORTION OF EVERY hedge fund's expected return is risk premiums — compensation to investors for bearing certain risks. One of the most important benefits of hedge fund investments is the nontraditional types of risks they encompass, such as tail risk, liquidity risk and credit risk. Most investors would do well to take

Table 4: A Comparison of Hedge Funds and Their Clones*

The promise of replicating hedge fund returns varies greatly by strategy. As shown by the following performance comparison of linear clones with the corresponding hedge funds in the TASS database, the technique is very effective for convertible-arbitrage, global macro, long-short equity hedge and managed-futures strategies.

		AVERAGE MEAN RETURN		STANDARD DEVIATION RETURN		SHARPE RATIO		FIRST-ORDER AUTOCORRELATION ($\hat{\rho}_1$)	
Category description	No. of funds	Mean (%)	Standard deviation (%)	Mean (%)	Standard deviation (%)	Mean	Standard deviation	Mean (%)	Standard deviation (%)
LINEAR CLONES									
Convertible arbitrage	82	8.15%	5.15%	6.20%	5.28%	1.54	0.62	10.4%	10.7%
Dedicated short-bias	10	3.58	13.09	28.27	10.05	0.16	0.54	1.2	4.4
Dedicated short-bias**	9	–0.35	4.40	28.75	10.53	0.00	0.17	1.9	4.2
Emerging markets	102	17.91	16.51	22.92	15.16	0.97	0.61	0.7	8.8
Equity market-neutral	83	7.45	6.81	7.78	5.84	1.14	0.76	1.8	9.6
Event-driven	169	9.60	6.79	8.40	8.09	1.39	0.52	3.5	11.3
Fixed-income arbitrage	62	8.55	6.04	6.56	4.41	1.43	0.64	2.5	8.2
Global macro	54	14.43	9.65	11.93	6.10	1.25	0.55	3.9	8.9
Long-short equity hedge	520	13.94	10.34	15.96	9.06	0.96	0.59	0.1	9.5
Managed futures	114	23.47	15.94	21.46	12.07	1.11	0.46	5.7	8.5
Multistrategy	59	10.10	7.66	8.72	9.70	1.50	0.68	1.8	10.0
Fund of funds	355	8.63	5.88	6.36	4.47	1.46	0.48	–0.3	11.2
ACTUAL FUNDS									
Convertible arbitrage	82	8.41%	5.11%	6.20%	5.28%	2.70	5.84	42.2%	17.3%
Dedicated short-bias	10	5.98	4.77	28.27	10.05	0.25	0.24	5.5	12.6
Dedicated short-bias**	9	4.92	3.58	28.75	10.53	0.20	0.20	3.4	11.3
Emerging markets	102	20.41	13.01	22.92	15.16	1.42	2.11	18.0	12.4
Equity market-neutral	83	8.09	4.77	7.78	5.84	1.44	1.20	9.1	23.0
Event-driven	169	13.03	8.65	8.40	8.09	1.99	1.37	22.2	17.6
Fixed-income arbitrage	62	9.50	4.54	6.56	4.41	2.05	1.48	22.1	17.6
Global macro	54	11.38	6.16	11.93	6.10	1.07	0.58	5.8	12.2
Long-short equity hedge	520	14.59	8.14	15.96	9.06	1.06	0.58	12.8	14.9
Managed futures	114	13.64	9.35	21.46	12.07	0.67	0.39	2.5	10.2
Multistrategy	59	10.79	5.22	8.72	9.70	1.86	1.03	21.0	20.1
Fund of funds	355	8.25	3.73	6.36	4.47	1.66	0.86	23.2	15.0

** February 1986 to September 2005.*
*** Fund No. 33735 has been dropped from this sample of dedicated short-bias funds.*

Source: Authors' calculations.

on small amounts of such risks if they are not already doing so because these factors usually yield attractive risk premiums, and many of these risks are not highly correlated with those of traditional long-only investments. Although talented hedge fund managers are always likely to outperform passive buy-and-hold portfolios, the challenges of manager selection and monitoring, the lack of transparency, the limited capacity of such managers and the high fees may tip the scales for the institutional investor in favor of clone portfolios. In such circumstances, portable beta may be a reasonable alternative to portable alpha.

Our empirical findings suggest that the possibility of cloning hedge fund returns is real. For certain hedge fund categories, the average performance of clones is comparable — on both a raw-return and a risk-adjusted basis — to that of their hedge fund counterparts. For other categories, like dedicated short-bias and event-driven, the clones are less successful.

As encouraging as these results may be, several qualifications must be kept in mind. First, we have used the entire sample of return histories to construct our clones, which is a particularly naive approach to replicating a dynamic strategy and also imparts a "look-ahead bias" to the results. Any practical cloning process must employ rolling or expanding windows to estimate the portfolio weights. This allows the clone-portfolio weights to change over time and in response to changing market conditions, a particularly important feature in the hedge fund context. Although the look-ahead bias may not be that severe in this case because we did not select the best-performing clone among many trials, nevertheless, a more realistic simulation is an important extension of our analysis.

Second, despite the promising properties of linear clones in several style categories, it is well known that certain hedge fund strategies contain inherent nonlinearities that cannot be captured by linear models (for example, Capital Multiplication Partners). Therefore, more sophisticated nonlinear methods — including nonlinear regression, regime-switching processes, stochastic volatility models and Kat and Palaro's copula-based algorithm — may yield significant benefits in terms of performance and goodness-of-fit. However, there is an important trade-off between goodness-of-fit and the complexity of the replication process, and this trade-off varies from one investor to the next. As more sophisticated replication methods are used, the resulting clone becomes less passive, requiring more trading and risk-management expertise, and eventually becoming as complex as the hedge fund strategy itself.

Third, the replicating factors we proposed are only a small subset of the many liquid instruments that are available to the institutional investor. By expanding the universe of factors to include options and other derivative securities and customizing the set of factors to each hedge fund category (and perhaps to each fund), it should be possible to achieve additional improvements in performance, including the ability to capture tail risk and other nonlinearities in a buy-and-hold portfolio. In fact, an earlier study by Lo and Martin Haugh shows that a judiciously constructed buy-and-hold portfolio of simple put and call options can yield an excellent approximation to certain dynamic trading strategies, and this approach can also be used to create better clones.

Finally, a number of implementation issues remain to be resolved before hedge fund clones become a reality: the estimation methods for computing clone portfolio weights, the implications of the implied leverage required by our volatility rescaling process, the optimal rebalancing interval, the types of strategies to be cloned and the best method for combining clones into a single portfolio. We are cautiously optimistic that the promise of our initial findings will provide sufficient motivation to take on these practical challenges.

Andrew W. Lo

Jasmina Hasanhodzic is a Ph.D. candidate in the Department of Electrical Engineering and Computer Science at the Massachusetts Institute of Technology. Andrew W. Lo is the Harris & Harris Group Professor of Finance at the MIT Sloan School of Management and founder and chief scientific officer of AlphaSimplex Group, a quantitative investment management company based in Cambridge, Massachusetts. *The views and opinions expressed in this article are those of the authors only and do not necessarily represent the views and opinions of AlphaSimplex Group, MIT or any of their affiliates and employees. The authors make no representations or warranty, either expressed or implied, as to the accuracy or completeness of the information contained in this article, nor are they recommending that this article serve as the basis for any investment decision. For a complete list of references used in preparing this article, as well as a more detailed explanation of the analyses and additional results and tables, please refer to Lo's home page, web.mit.edu/alo/www.*