Hedge Fund Replication

A number of different approaches to replication have recently been postulated by practitioners and academics. This articles examines a number of them commenting on their rationale, feasibility and prospects for success. By Tammer Kamel.

Before embarking on a survey and subsequent critique of several replication approaches available to investors, there is a fundamental question that must first be addressed: What is it about hedge funds that is valuable and worth replicating?

For example, a 45 year old egotist with a penchant for fine wines pontificating about the housing market from his midtown office might constitute a hedge fund replica, but what use? The very objective of replication depends on a hedge fund's utility. Is that utility its alpha generating capabilities? Its correlation characteristics? Mitigated downside risk? Access to alternative risk premia? Each of these utilities leads to vastly different replication objectives and hence different approaches.

What and Why

Indeed, there is no consensus as to the true utility of a hedge fund. The reality is various investors value hedge funds for different reasons meaning the objective of replication also varies per investor. So we can already conclude that it is unlikely that one universal replication methodology will ever prevail unless it somehow can achieve all the objectives mentioned above and perhaps others. This seems unlikely, especially now, when current technology struggles to achieve even one. Thus investors should be methodical in their use of cloning technology. To begin, they must ask themselves exactly what they want from their hedge funds in the first place before embarking on a replication quest.

If for example the value of a hedge fund is its alpha generating capability then replication becomes extremely challenging because the mission is to match the ingenuity of the manager. At the other extreme, an investor who uses hedge funds as beta sources can have high expectation of success from simple, passive replication methodologies. Still an other investor might use hedge funds because of certain return properties such as muted downside risk, lower volatility and minimal correlation. This again frames the replication problem differently and creates its own challenges. But the important point is that replication efforts cannot and should not commence until an investor decides exactly what hedge fund attributes are to be cloned.

After clearly defining the replication objective, cloning can commence. But even then, there remains an important secondary question: Why replicate at all? Academics are perhaps excused from this question, but an investor has more than 10,000 hedge funds to choose from. What does she hope to achieve by cloning that cannot be found in the vibrant and competitive hedge fund market place?

On this question, there seems to be four areas where a clone can be superior to a conventional fund. The first is liquidity. An investor in a conventional hedge fund is bound by the often arbitrary redemption rules of the fund. An investor in a hedge fund often has to "manage" liquidity by attempting to anticipate the behavior of her fellow investors lest they redeem quickly, leaving her with the least liquid (and least attractive) remnants of a fund's portfolio. An investor who executes a replication strategy does not face these problems. She is guaranteed "fair" liquidity at all times. She will enjoy whatever liquidity is on offer

Real or Clone?

from the actual market and can never be hurt because other investors beat her to the exit.

The second advantage of a clone is transparency. An investor implementing a replication strategy knows at all times the precise composition of the portfolio. This has risk measurement and risk management advantages plus the corollary that style drift is impossible.

Perhaps the biggest advantage offered by in house replication is cost. Not only are management and performance fees eliminated, but the time cost of comprehensive due diligence is too. On a $100 million portfolio, the sum of these costs could easily be $4 million. That constitutes a 4% performance improvement which can be put directly in the "alpha" column.

The final advantage is one of capacity or scalability. An investor who clones a hedge fund via highly liquid capital markets will never be capacity constrained.

Any one of the these four benefits make the cloning exercise appealing. The four combined make a very strong case. And an investor gets all these benefits the moment a conventional hedge fund is replaced by a clone. However, they are only real benefits if the clone actually succeeds in its replication mission and there of course lies the rub. This article explores this challenge by examining several different replication approaches that have been postulated by a number of capable academics and practitioners.

Tell Dr Kat what you really want

Hedge funds might be of value to investors, but it's not because of alpha. This is the core assumption of the replication methodology Professor Harry Kat and his coauthor Helder Palaro espouse in a series of papers and related product offerings. In fact, they make the assumption that perpetual alpha generation is impossible because hedge funds operate in near efficient markets. This is bold assumption indeed, and one that many—probably all—hedge funds would dispute. But refuting it is difficult because, while there are plenty of hedge funds which, statistically, show

evidence of alpha there are just as many that don't. So on aggregate, total hedge fund alpha production might actually be close to zero and perhaps negative after all fees and biases are accounted for.

Most hedge fund managers will concede that aggregate hedge fund alpha is nil because such a statement says nothing about their own abilities. However, this is a problem for investors because it means that the alpha expectation of any arbitrary hedge fund is zero. In this context, Professor Kat's assumption is entirely reasonable.

Under an efficient market assumption then, any replicator is liberated from pursuing alpha. But this then begs the question, why would one want to replicate a hedge fund if it contains no alpha let alone invest in it? Professor Kat would reply that a hedge fund may be of value to an investor not because of excess returns, but because of its return distribution. An investor may value mitigated downside risk, positive skewness, low or negative correlation, or decreased volatility. Even accompanied by no alpha, these performance properties of hedge funds can be attractive. Indeed, that investors continue to flock to hedge funds in spite of the low alpha expectation, argues for this motivation.

Hedge Funds in Efficient Markets

In this paradigm then, the utility of a hedge fund is its ability to transform the risk/return profile of an asset class. For example, a "good" equity long short fund takes the log normal distribution of the stock market and skews it to the right while reducing its volatility. Thus the value of this equity long short fund is its ability to generate returns with certain statistical characteristics. Thus the mission of a replicator is to produce returns with similar statistical characteristics.

The return characteristics of a hedge fund can be described parametrically using conventional statistical measures like volatility, skewness, kurtosis and correlation or more generally via a distribution function.

In either case, the mission of a clone is simply to produce a return distribution that matches that of the archetype hedge fund. It doesn't matter if the clone uses a different strategy or a completely different market, as long as the return distribution is per specification, replication has been achieved. Any vehicle that can achieve this would constitute a valid replica. Note that this means that the monthly returns of archetype and clone could be different, but in time, their aggregate statistical properties (volatility, skewness, kurtosis, correlation, etc) will match.

This ability to alter the risk/return profile of an asset or asset class is not something unique: options do precisely that same thing. Like a hedge fund, an option takes one distribution and transforms it into another. For example, a call option takes the lognormal distribution of a stock and transforms it into the positively skewed, asymmetric distribution of the option. The essence of Professor Kat's approach to find an option whose return distribution matches the archetype fund. For example, a correctly specified exotic option replicates an equity long short fund if it transforms the lognormal distribution of stocks in the same way the fund does. So the professor's methodology is ultimately an exercise in reverse engineering: find the exotic option with same return distribution as the hedge fund.

Finding the replicating option reduces to defining a payoff function. This function can be deduced given the distribution of the underlying and the required distribution of the option. But having done this, one cannot simply purchase such an option in the OTC market because even describing the payoff function contractually could be impossible let alone finding someone to sell it to you. Instead, the solution is to synthesize the option. This in can be done by combining a set of vanilla options to achieve the desired exotic option. But professor Kat takes an even simpler approach: He synthesizes the option by dynamically trading the underlying. Synthesizing options by

dynamic trading is actually a well developed science.

Risk and Return

In general, unless one can buy them cheap, options do not offer an investor higher risk adjusted returns. Indeed, the risk adjusted returns offered by a fairly priced call option on a stock are exactly the same as the risk adjusted returns offered by the stock itself. In other words, the expected risk adjusted return of a (correctly priced) derivative is only ever as good as its underlying. Given this axiom Professor Kat wisely suggests that when creating a "hedge fund exotic option" one should use an underlying asset with high expected risk adjusted returns. For example, if one replicates a equity long short fund using stocks, the replicating fund's risk adjusted returns can at best be as good as the underlying stocks'. But if instead one uses an asset class with a higher Sharpe Ratio than stocks, the clone will have better returns. Indeed, Professor Kat suggests using a maximally diversified portfolio of stocks and bonds as the building block of replication since such portfolio offers the highest rate of return per unit of risk. This is why Professor Kat's methodology may attempt to replicate some particular hedge fund strategy without ever touching any of the instruments the archetype fund would employ.

There is a further implication of this: The expected returns of the hedge fund clone will be the same as the building block asset one uses less the net cost of options. For example, if the target return distribution is to minimize downside risk then it will implicitly include put options which will have to be paid for. This in turn implies a lower return. Professor Kat's method gives the investor the ability to define the exact distribution she wants, but the expected return will be set by the market based on the building block portfolio's expected return and the cost of downside protection she is implicitly buying.

Professor Kat is adamant about one important point: It is not sufficient to merely replicate the marginal distribution of the hedge fund. Its correlation characteristics vis-à-vis an investor's core portfolio are critical. Hence the reverse engineering problem is one that includes correlation as well. In general, one must find the payoff function of an exotic option whose joint distribution matches the joint distribution of the archetype fund and the investor's core portfolio.

One irony of the methodology is that a clone employing the technique might actually execute a trading strategy substantially more complex than its archetype. On the other hand, one of the appealing things is that the technique can be extended beyond just replication: Instead of using some real hedge fund to infer a desirable return distribution, one can dispense with the archetype fund entirely and simply define some ideal performance characteristics and proceed to produce that. Either way, a synthetic hedge fund in Professor Kat's methodology is nothing other than a fairly priced exotic option on a broadly diversified portfolio of stocks and bonds. The user chooses the exact statistical properties she wants—and gets them, but the expected return of the replicating fund is set by the market based on the expected return of the underlying and the price of risk.

Hedge Funds: What are they good for?

Professor Kat's approach is predicated on a the important assumption that the value of a hedge fund is its ability to transform return distributions. If this really is the only value of hedge funds, then they are most certainly expensive and overly risky. Much better to just use options to achieve the same thing and dispense with the hedge fund entirely, as Professor Kat suggests. His replication product invites investors to choose any return distribution and it will construct the option (and advise how to synthesize it via regular trading of the building block portfolio).

But regardless of what most hedge funds are, they should actually be more than just distribution transformers, even if they can't perpetually generate alpha. They should at least seek alternative returns, (hedge fund beta in the vernacular). Under an efficient market assumption, hedge funds would still add value if they were doing this. Only if these hedge fund beta factors joined Professor Kat's building block portfolio, would his methodology really subsume hedge funds and still only if his efficient market assumption is valid. But hedge fund beta is difficult to passively invest in. Hence the usefulness of the real thing, even sans alpha.

Good Replication of a Bad Investment?

While Professor Kat is fundamentally right that investors often get nothing out of their hedge fund portfolios other than a transformed risk/return profile on some set of beta factors, it is still the case that hedge funds promise to be much more. That they don't actually deliver is a failing of the individuals who run them. But it is not a flaw in the ideal they strive for.

While Professor Kat's product may do well to give investors what they are actually are getting out of their real world hedge fund investments, it still fails to deliver the promise of hedge funds—just like the real ones do. This begs the philosophical question, should one replicate what hedge funds *are* or what they *should be*?

Light on the kurtosis, extra skew please

Synthetic hedge funds are ultimately just hedge funds employing a new strategy. They will in time be judged on the exact same criteria "traditional" hedge funds are. And, despite the best advise of textbooks, professors and me, investors are return zealots. They are preoccupied by returns—and not risk adjusted returns—just good old fashioned bankable returns. I don't meet many hedge fund investors who put skew, kurtosis and correlation ahead of returns on their due diligence checklist. They should of course, but that they don't is a big marketing problem for Professor Kat because his methodology deemphasizes returns and that is a difficult sell in this industry. His replication methodology is sound and theoretically sensible, but alas it may be incompatible with the investment approach of many.

Plain Vanilla Replication: Lo and Hasanhodzic

Like Professor Kat, Professor Lo is not concerned with exactly what hedge fund managers are doing. For replication purposes, he too treats hedge funds as black boxes which is perhaps ironic since he is himself a competent traditional hedge fund manager when not at MIT. Writing with Jasmina Hasanhodzic in 2006, Professor Lo postulated a simple, linear factor model for replication which works well enough to call into question just how much value hedge fund managers add, yet not well enough to suggest firing them all.

Professors Lo and Hasanhodzic begin with the assumption that some fraction of hedge fund performance is not the product of manager skill but merely returns from simple, passive exposures to certain asset classes. Replication of this component of hedge fund performance is thus achievable by simply making certain passive investments. If manager skill is only responsible for a small component of hedge fund returns, then this approach can be expected replicate the bulk of hedge fund returns. On the other hand, if hedge fund performance is dominated by manager skill or even just attributable obscure beta factors, then his core assumption is wrong and clones will never come close to matching the performance of archetype hedge funds.

As a simple example, consider an equity long short fund which is net long the stock market. In such a fund, some component of performance will necessarily be replicable by simply being long a diversified set of stocks. The rest of the fund's performance

however, is governed by the manager's capability to buy eventual winners and short eventual losers. If the manager is not particularly good at this, then his efforts will translate to noise around his net long position. In this case, replicating the fund by simply being long the stock market will work quite well. However, the better the manager is, the less the clone's passive strategy will be able to keep up.

Professor Lo identifies six passive investments that can be combined in some way in an attempt to replicate hedge fund performance. They are: stocks, bonds, currencies, commodities, credit and volatility. Having selected an archetype fund, his mission is then to determine what amount of exposure to each of these investments would best approximate the hedge fund and then measure just how well they do.

The canonical way to do this is via a mathematical technique called regression which determines the size of each of the six passive investments that minimizes tracking error between archetype and clone. Regression also measures what fraction of archetype performance is actually captured by the clone's passive investments. Note the key difference from Professor Kat: The Lo methodology is return centric. It aims to match returns on a per month basis as best as possible with little regard for anything else. The Kat methodology is not return centric at all, preferring to match aggregate return statistics. Professor Lo's method is actually the much simpler approach both conceptually and in implementation. His methodology might manage to capture some correlation characteristics of the archetype, but its linear, passive nature prevents it from ever achieving the other objectives of Kat's, like skewness, kurtosis, volatility, etc. The two different methods underscore how different replication objectives leads to completely different approaches and thus different clones.

Does Professor Lo's methodology achieve its objective? The answer is a qualified no. His analysis does reveal that a great deal of hedge fund returns are indeed nothing more than returns on simple, passive investments. Indeed, many a single hedge fund is easily replicated by passive beta exposure. However, across some 1600 hedge funds, replication failed to capture a substantial fraction of real hedge fund performance. In other words, on average, the clones underperformed.

But one should not take this result to mean passive replication can't work. Professor Lo was not setting out to develop an industrial strength clone. Rather, he was testing whether passive replication can work at all. And given how well his simple six factor approach worked, there is reason to be optimistic about the next generation of passive replication approaches. Indeed, if hedge funds were replicable by the relatively simple approach purposed by Professor Lo, the implication would not be a clone invasion but rather, one would hope, the demise of funds altogether; investors would surely go straight to the beta. Ironically, successful passive replication might precipitate the death of archetype and clone.

Lo's Alpha

In the classic alpha/beta paradigm, alpha is desirable because it represents performance from some unique risk source, presumably the manager's skill set. In practice, alpha is just the label for those returns that cannot be attributed to some set of beta factors. In Professor Lo's analysis, hedge fund returns which were not attributable to stock, bond, currency, commodity, credit and/or volatility markets were labelled as alpha. But what is the nature of Lo's alpha?

There are two possibilities. Firstly, Professor Lo might simply be missing some simple beta factors that might well mop up a lot of excess returns not explained by his original six factors. For example, the "small cap" beta factor is missing from the professor's analysis. Its presence might explain away some of what thus far looks like alpha. Indeed, other beta factors might do the same.

The other possibility is that the alpha is truly an artefact of manager skill. Interestingly, this does not necessarily mean the returns are can't be replicated. Technically, it depends on the precise nature of the manager's skill. Presumably there are two kinds: unique and common. Replicating the former should be impossible, but not so for the latter.

An example of a "common" skill would be executing an fx carry trade. There is no secret to how this is done: One shorts a low yielding currency and gets long a high yielding one. Furthermore, perpetually holding a diversified portfolio of fx carry trades tend to be profitable albeit risky. If a hedge fund is in fact doing this, the dividends would appear as alpha in any regression similar to Professor Lo's. And while this alpha is technically the product of "manager skill" it is hardly a unique or valuable skill. Rather, it is skill that any decent hedge fund manager will have. Practitioners are beginning to differentiate these common skills from truly unique skills. They are, quite rightly, no longer in awe of the former while still quite willing to pay high fees for the latter.

As such, returns which are the product of common skills have lately been stripped of their alpha status. Instead, these returns are being referred to as "alternative beta" or "hedge fund beta". Alternative beta is not beta in the classic sense because it is not achievable by a passive, zero maintenance investment. It is really a trading strategy. But it is a trading strategy that is simple, well defined and, on average, profitable. Most importantly, returns from common skills are easily replicated which is why "beta" is a more appropriate classification than "alpha".

Examples of alternative beta include carry trading, liquidity trading, certain types of event risk and certain types of spread trades. All these strategies are well understood. They require active management by a trader who is competent, but not necessarily brilliant.

Next Generation Clones

Although Professor Lo falls well short of cloning Nirvana, he has demonstrated that a substantial fraction of hedge fund performance is indeed replicable by passive investing. The success he achieves with such a simple approach bodes well for the next generation which will reach to an expanded set of classical beta and perhaps alternative beta too. Armed with this extended tool set, these future clones will come closer to comprehensive replication.

Of Nickels and Steamrollers

In a 2005 paper, Professors Duarte, Longstaff and Yu explored the risk return characteristics of a number of fixed income strategies employed by hedge funds. Though the paper did not focus on replication, it offered a number of insights on the subject. This is because to achieve the measurements they wanted to take, the authors actually constructed hedge fund replicators. But their approach was distinct from anything thus far discussed. What the authors did was emulate the actual trading approach of fixed income arbitrage strategies by mechanized rules. This in essence is "white box" replication: Do what it is the hedge funds are doing.

Unlike the methods thus far discussed, one must understand the actual strategies employed by hedge funds to succeed in this approach. The deeper the understanding, the better the prospects for successful replication. The major advantage of this replication approach is that, if it is done correctly, it could capture everything a conventional fund has to offer—including alpha. The disadvantage is that the executor must himself be an expert on any strategy he aims to replicate.

In their paper, the authors constructed various "micro" hedge funds, each employing a particular strategy. For example, the authors constructed a swap spread arbitrage clone. Swap spread arbitrage is the art/science of betting on the yield spread between government bonds and

libor swaps. The authors' clone took swap spread positions based on simple rules consistent with how actual arbitrageurs behave. After "paper trading" the strategy, the authors measured its performance characteristics. Their emphasis was on what fraction of returns could be characterized as alpha. (They did not compare their clones' performance to real hedge funds, primarily because this was not their purpose. Also, it would be difficult to find the right benchmark fund to compare their highly focused, single strategy funds to.)

But even though replication was not the authors' objective, there are lessons for aspiring clones. Firstly, the authors did find that a substantial fraction of returns from "sophisticated" fixed income arbitrage is beta driven. This is encouragement for the likes of Professor Kat and Professor Lo who aim to replicate via beta only. However, there is bad news too: The authors found that the performance of some clones were not all beta driven. Some of their clones generated alpha too. This is anathema to beta based replicators because they, by construction, cannot produce alpha.

One interesting conclusion in the paper was that it seems that the more intellectually challenging hedge fund strategies are more difficult to replicate by simple beta approaches. For example, the authors finds that mortgage arbitrage, a notoriously complex strategy, can not be fully replicated by simple beta factors. Volatility arbitrage on the other hand, can be.

The key question is whether a combination of white box micro clones would subsume a real fixed income arbitrage fund—alpha and all.

Conclusion

The theoretical and practical value of hedge fund replication remains contentious. Nevertheless, several institutions have recently introduced replication products. In one sense, these offerings are ultimately just hedge funds implementing a strategy called "replication". As such, they will be judged as all hedge funds are, which probably means by track record: annual return and volatility.

The verdict of the market place will be easy to quantify because it can be measured in AUM terms. If, net of fees, hedge fund replicators out perform conventional hedge funds, they will attract capital. But many hedge fund investors typically demand multi year track records before they affirm a particular manager. And so one certainty is that, while replicas might fail any time, it will be some years before they are able to establish a reputation of capability. Investors will also need to be wary of a fresh dose of survivorship bias because, if enough replication products launch in the near future, there will inevitably be some left in three years from now—even if they are all poor.

Conventional hedge funds will not be a causality of replication products any time soon. But perhaps their fee structures will be. If replication research has demonstrated anything conclusively, it is that a substantial fraction of hedge fund performance is nothing but beta. And beta is not something that investors should pay high management fees for and certainly not something that they should pay performance fees on. The emergence of replication strategies might spurn investors to demand performance fees linked to alpha—otherwise known as benchmarking.

Replication is also forcing investors to face Professor Kat's most excellent question: What do you really really want from your hedge fund investment? For too long, with kudos to hedge fund marketers, hedge funds have been perceived as "better" than conventional investments. Replication should prompt investors ponder what they actually want from a hedge fund and in doing so figure out if hedge funds are "better" at achieving their objectives. This sort of sober approach to hedge fund investing is long overdue.

www.ilukacg.com