PHANTOM PRICES & LIQUIDITY: THE NUISANCE OF TRANSLUCENCE

By Adil Abdulali, Leslie Rahl and Eric Weinstein, CMRA

"It is worth noting that the use of illiquid, and/or over-the- counter securities is particularly endemic to the hedge fund industry and is rarely satisfactorily addressed in academic studies of hedge fund performance. "Andrew Weisman, 2001

Introduction

"Liquidity" is a frequently used term in finance, with a wealth of different meanings. Market players readily agree to the importance of liquidity as a concept without agreeing on its meaning, a conceptual framework for its discussion, or an indicator for its measurement.

The many nuances of hedge fund liquidity support an evolving body of research that attempts to both explain and quantify liquidity phenomena. Even commonly held perspectives such as that which views illiquidity as a kind of market defect to be eschewed, may seem naive with the realization that the opportunity cost of being highly-liquid is likely to be substantially greater than zero. This chapter does not purport to be a taxonomy nor a comprehensive literature review, but instead attempts to articulate some aspects of liquidity, which the authors believe have particular importance in understanding risks in hedge funds, as many types of liquidity risk cannot be fully diversified away in the fund of funds structure.

This chapter has a dual purpose. The first is to suggest a general quantification scheme for liquidity risks. The second is to describe the genesis of the Phantom Price framework for quantifying valuation risk for translucent assets. This framework is an example demonstrating that the general scheme can be made concrete. Though some of the mathematical details of the Phantom Price framework have appeared elsewhere, the practical problems that inspired them have not.

Traditional definitions and examples of liquidity

While a specific and focused definition is lacking in Finance, the word "liquid" generally refers to collections of assets that have the property of being readily exchangeable for cash. Small positions in large cap equities, foreign hard currencies, or precious metals are but a few examples of highly liquid positions. Conversely, "illiquid" refers to collections of assets that, relative to the most convertible of assets, are not so easily exchanged. The most perfectly illiquid assets are those assets with intrinsic value that, for one reason or another, cannot be exchanged for cash at all (e.g. a transplantable kidney in a society which forbids the selling of organs). Many uses of the term "liquidity" make an assumption about how much the price of an asset moves away from some equilibrium level as different amounts of the asset are traded[1]. Perhaps the second most common use of the term refers to the extent of a discount from some assumed 'fair' valuation which must be offered if the asset is to be sold more quickly than is needed to realize the 'full value' of the position. Still another use of the term may refer to the depth of the market with respect to buyers and sellers where markets with many participants are thought to be more liquid than those with relatively few players.

At the high end of the liquidity continuum are US Treasury bills. They are perceived to be so liquid that the market for such bills is called the "money" market.[2]

At the opposite and low end of the liquidity continuum are restricted securities that can only be sold for cash in special and rare circumstances. Real estate, traditionally used as an example of an illiquid asset is in fact often more liquid than these highly illiquid securities. Hedge funds that have onerous redemption restrictions would be considered less liquid than a house of average price. To understand this, it may prove helpful to imagine a house that the owner is not allowed to sell for several years or for which the owner had to commit to selling the property months before knowing the price for which the house would sell.

In the gray middle range of the liquidity continuum, we have assets, such as corporate bonds and mortgage-backed securities. There are days when a billion dollar trade in mortgage securities would not move the market at all but other days when it would be impossible to trade such a volume without a major market disruption. Without frequent trades, the price of some of these securities becomes a guessing game. The guesses of major market participants may be significantly different from each other as illustrated in Table 1.

Liquidity and Pricing - two sides of the same coin

The following scale suggests the prevailing beliefs regarding the duality of liquidity in pricing and transparency in valuation.

[1] Market Liquidity: Research Findings and Selected Policy Implications, BIS, Basle 3 May 1999.

[2] "Money" instead of "cash" because the word "cash" is used to describe assets that have to be paid for, as opposed to derivatives which might have no exchange of actual cash.

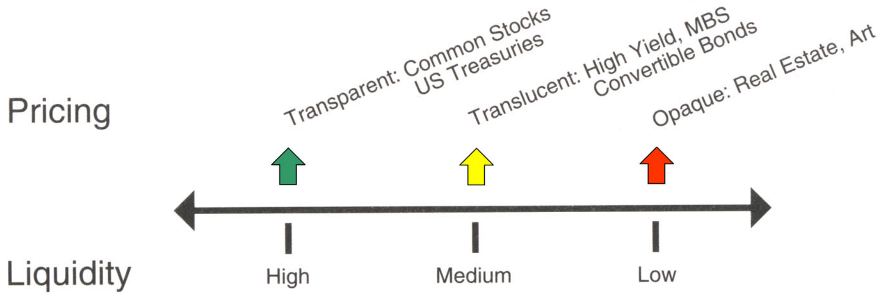

It is interesting to the authors that, in the summer of 2001, when the first draft of these ideas was circulated within the community, the inclusion of convertible bonds in the "medium" instead of "high" category was questioned by some. Since then, several high profile and unfortunately mis-marked convertible positions have convinced us that they do belong in the medium category.

Even stocks are not uniformly of high transparency. There are many small cap stocks that trade infrequently, in tiny quantities, and at questionable levels. It is no doubt still heresy to some to question the price of a stock quoted from an exchange, but this is no reason in itself to refrain from doing so. Of what significance is it that a single stock traded at some particular time during the day at some price? And why is the 4 pm New York closing time the right time to "mark-to-market" stock positions? We believe that this artificial construct is analogous to the use of a thin dime as the poor man's improvised screwdriver: until something better comes along, we as market participants simply tend to make do. However, as finance inexorably 'progresses' towards 24/7 trading around the world in a single sleepless market, new tools will likely need to be developed to replace the daily closing NAV concept.

Measuring Liquidity

Market participants generally view illiquidity in securities markets in one of three ways:

- As a defect which should be fixed or avoided.
- As representing opportunity that should be estimated and priced.
- As an annoyance which, in most periods, is more profitably ignored than engaged.

While each of these positions has its merits, it is the second of these positions that we find the most interesting. It is our view that liquidity, or more precisely the lack thereof, is a risk, and like all risks should afford rewards to the risk-taker. With this perspective in mind, the problem therefore becomes one of actively developing a framework for measuring the risk/reward tradeoff.

Volume

A common proxy for liquidity, market volume provides interesting problems and insights. The number of units traded is one measure of volume. However since different assets have different unit sizes in respect of investment amounts, the number of units times the size in dollars of each unit gives another important and different way to measure volume. This volume measure is sometimes referred to as "dollar volume". For equities, when comparing stocks with very different per share values, the two volume measures give different rankings. For example, on June 27, 2002, 980 shares of Berkshire Hathaway were traded as opposed to 1.13 million shares of Clorox. Since Berkshire Hathaway traded at an average price of $67,750 and Clorox at $41.5, we see that the dollar volume of Berkshire was approximately $20 million higher with a thousand times fewer shares traded. Despite such effects, it is common to hear the unadjusted trading volume reported on a given day in terms of the number of shares changing hands.

An interesting aspect of volume is the concept of the size of a portfolio holding vs. the average traded volume in some time interval. If a large portfolio consisted of half a million shares of stock A, than clearly the ease with which the position could be liquidated is significantly less in the case where the average daily volume is 100,000 shares as opposed to the case where the average daily traded volume was 15 million shares. There have been interesting recent theoretical steps taken to quantify this volume effect[3]

The classical micro-economic supply and demand curve model sheds some light on the liquidity puzzle. Oftentimes a highly liquid market is thought to be one where the price of an asset does not change perceptibly given moderate shifts in the quantities offered by a representative investor. If true, this would imply a set of nearly elastic demand curves that should be expected to be the exception and not the rule. Instead it is far more realistic (and is in fact assumed by most sophisticated practitioners) that large quantities of an asset should not be valued as the sum of the unit sizes. However a consistent theoretical framework as yet does not exist to implement these

---

[3] Cosandey, David " Adjusting value-at-risk for market liquidity", Risk, October 2001

ideas for purposes of portfolio valuation. As such, NAV's are often calculated in such a way as to ignore the corrections due to volume despite the obvious importance of such adjustments.

Bid-Asked Spread

A departure from the "one-price" dogma has already occurred in some markets. Market participants have realized that where a bid-offer spread exists, long positions marked at the bid and short positions marked at the asked price is a more accurate description of the liquidation value or cash-equivalent value of a portfolio. An interesting variation occurs in Over the Counter ("OTC") markets where one side of the bid-offer is purely theoretical. For example, in markets for certain Collateralized Mortgage Obligations ("CMO"s), each asset is unique and cannot be readily shorted or borrowed. Therefore a market-maker who quotes a bid and asked price can in reality consummate only one side of the transaction depending on whether the market-maker is already long the quoted asset. It takes creativity to interpret the meaning of the other "theoretical" side and such creativity strains the credulity of all concerned.

Pricing Liquidity

There are many ways in which liquidity, or lack thereof, is priced in the market. There is the conventional understanding that, in a highly efficient market, the market price of an asset has "all" aspects of its risk efficiently incorporated within its price and no adjustments are necessary for valuation in the context of a portfolio. Unfortunately, this quickly falls apart even in the case of common stocks.

For example, if an entity holds more than a certain percentage of the outstanding stock of a corporation, then that entity enjoys greater rights and privileges in running the company than say the holder of a few shares of the outstanding stock. These greater rights (dubbed "control") would tend to push the value of a controlling block of stock to a figure greater than the simple sum of the values of each (noncontrolling) share considered individually.

But as previously suggested, there is an opposing argument that the large block might be priced at a discount relative to the sum of its parts whenever it would be difficult to liquidate the large block without severely depressing the market.

It seems under any serious theory that the valuation of the position must be expected to be dependent on the time allowed for and the circumstances surrounding the liquidation. It is evident that a richer metric is needed than is currently available to market participants. Such a new metric, one hopes, will provide a consistent alternative to the dogma of a single additive price for each asset in any context.

Four step program: valuation framework for quantifying liquidity

In an oft repeated aphorism, Albert Einstein famously quipped that ‘things should be made as simple as possible, but not simpler’. Unfortunately, in the case of the all too genuine complexities of hedge fund liquidity, the urge towards radical simplification has often proven an irresistible temptation.

In some sense, the liquidity problem can be viewed as representing a classic dichotomy between theory and experiment. An actual redemption of assets by a hedge fund represents a natural ‘valuation experiment’ while the reporting of its performance statistics (often in the absence of such informative transactions) represents an exercise in ‘theory’. When theory and experiment are extremely close together, as they are for small portfolios of actively traded common stock, we can perhaps be forgiven for ignoring this dichotomy when we speak of a single objective NAV. The small investor in blue chip stocks has every reason to be confident that the theory underlying his portfolio valuation will be in excellent agreement with experiment at any given moment during trading hours should he choose to conduct a transaction ‘experiment’.

The success of common stock valuation may be thought of as akin to the perfect parabola predicted by Newtonian mechanics when a feather is catapulted skyward in a sealed vacuum. The beautiful agreement with theory is quickly degraded when the artificial vacuum is replaced with more generic real world conditions as mild as a gentle breeze. Likewise in hedge fund valuation, it is the lack of antiseptic laboratory conditions provided by a rarified trading environment such as the NYSE which forces us to examine the shortcomings of our valuation theory. The accurate valuation of small positions in blue chip equities is distinguished by the relatively extreme experimental background conditions which may easily all be absent for even a small position in a hedge fund:

1. Trades occur openly in view of other market participants.
2. Trades occur in a centralized location rather than in many locations.
3. Trades are occurring nearly continuously with numerous active buyers and sellers.
4. Trades occur with bid-ask spreads being consciously and deliberately narrowed by specialist market-makers.
5. Trades can be counted upon to be executed promptly with minimal delay.
6. Trades involve small numbers of shares relative to the volumes being traded.

Part of the problem with current approaches to liquidity is that they tend to include effects which may conflate conceptual errors with actual liquidity effects. The most common such error is the failure to recognize when published scalars (e.g. NAV) become mere numerical stand-ins for non-scalar concepts.

Consider the army of different indicators needed to give us even a sketchy summary of a highly liquid mutual fund invested in largecap equities so that we might distinguish it from its competitors:

1. Present NAV
2. Average Growth Statistics
3. Beta
4. Sharpe Ratio
5. P/E Ratio
6. Price to Book Ratio

Etc... Why are so many numbers published when only the present and future value of the fund will determine its usefulness to investors? In essence, we have learned that numerical representations of mutual funds require arrays because such funds represent sufficiently complicated objects that a single scalar most always proves inadequate even at a summary level.

For moderately illiquid positions in hedge funds, the story is not much different conceptually. The key new problem is that traditional scalars, such as the NAV's that are sought, are now naturally modeled by probability distributions which, in turn, must be milked to provide the summary statistics. To see why this is the case, consider the value of 100 shares of Exxon stock vs. the sale of a world-famous painting. At any given moment, the owner of the exchange traded stock has a very precise idea of the liquidation value, while the owner of the painting has only a probabilistic guess as to the true value of the artwork, unless he or she has previously committed to selling it at the moment of its valuation.

The key here is to see that both valuations can be seen within the framework of distributions. In the case of the Exxon shares, the relative certainty of the price of the shares translates to a probability distribution which has very low variance as the true price almost certainly lies between the reported bid and ask quotes which will almost certainly be close together. Even with the most expert appraisal, the owner of the artwork will likely have to make due with an imputed probability distribution with a rather high variance due to the vagaries of the market for masterpieces. However, the same conceptual framework is now seen to cover both cases and the remaining difference is over parameters rather than principles.

A 4-step framework for the measurement of illiquidity:

Step 1: replace ill-posed appraisal questions with well posed valuation problems

We define a well-posed valuation problem to be a question about the value of a position that could be answered definitively by conducting a specific sale or purchase. Thus, while the abstract value of the Mona Lisa would not constitute a well-posed valuation problem, the question 'how much would the Mona Lisa fetch in cash at a sealed-bid auction open to qualified buyers to be announced today and closed in n days time' would constitute such a valuation problem.

In this example, there is but a single universally coveted object so that issues of liquidity are reduced to the discount that might result from being forced into a hasty sale with time being the only relevant variable. More generally, however, illiquidity may involve aspects of volume, distress, the difficulty of finding sufficiently numerous interested potential owners, or other such parameters. The answer to the valuation problem then becomes a function of several variables:

$$v(a1, a2, a3, \ldots) \qquad (1)$$

and the relative degrees of illiquidity can be represented by appropriate ratios.

Step 2: pick complete sets of parameters

In order to represent relative illiquidity as a ratio, it is necessary to specify 2 sets of parameters for a single well-posed liquidity problem. In such a case a ratio will express the ratio of values that is directly dependent on the liquidity parameters.

$$\frac{v(a1, a2, a3, \ldots)}{v(\hat{a}1, \hat{a}2, \hat{a}3, \ldots)} \qquad (2)$$

Thus if we imagine a single varying parameter $\boldsymbol{a}1$ representing the volume of shares to be sold with all other $\boldsymbol{a}i$ held constant, the liquidity gives a measure of the price elasticity. Such a ratio can then act as the required scaling factor to transform an NAV calculated for one parameter set into an NAV appropriate for a different well-defined valuation problem.

Step 3: pick appropriate distributional proxies

Since we do not expect to always have the luxury of holding or witnessing a sale for the purpose of valuing the positions of interest, it is foolish to think that we can always expect to know the answer to a well-posed valuation problem with certainty. Thus what we should attempt to do is to admit our uncertainty, and not search for a single price for our position. Instead, we should seek to put forward a probability distribution of all possible prices, which are assumed relevant. Thus, we call a valuation proxy, a method of producing families of distributions represented by probability density functions:

$$\psi(a1, a2, a3, \ldots) : \mathbb{R}_+ \longrightarrow \mathbb{R}_+ \qquad (3)$$

which are parameterized by our same sets of liquidity parameters. These distributions give us an idea of the probability that the answer to a well-posed valuation problem would lie in any given interval of possible prices.

Step 4: find statistics to reduce the distributions to scalars

We are now in a situation where we have proxied the numerator and denominator of our well-defined (but unknowable) liquidity ratio with objects which are no longer scalars. Because we cannot take ratios of distributions, we must seek ways to summarize the content of the distributions with statistics so that we may take ratios to proxy the true (but unknowable) ratios mentioned above. One possibility is to take the means of the various distributions produced. However, this is artificially restrictive in that it describes a risk neutral world where reward grows unacceptably slowly with increases in risk exposure. Another more general idea is to price the distributions using the calculus of expected value developed by Bernoulli, von Neumann and Morgenstern. If *E* is such a statistic (i.e. the risk sensitive expected value or risk neutral expected return), which takes the space of probability distributions to the positive real numbers, then it is possible to take ratios of the summary statistics to give a proxy for the degree of illiquidity according to

$$\frac{v(a1, a2, a3, \ldots)}{v(\hat{a}1, \hat{a}2, \hat{a}3, \ldots)} \approx \frac{E(\psi(a1, a2, a3, \ldots))}{E(\psi(\hat{a}1, \hat{a}2, \hat{a}3, \ldots))} \qquad (4)$$

Perhaps the first point to make is that when one is asked to calculate a liquidity ratio, it would become highly suspicious to argue that it is somehow always equal to unity. Yet that is precisely the default hypothesis being used implicitly whenever this scaling factor is neglected! Secondly, because generally accepted measures of illiquidity have been slow in coming, failures of valuation and appraisal are frequently blamed on illiquidity. Many difficulties of market agents which are colloquially blamed on 'illiquidity' are really grievances involving improper valuation. Referring back to our supply demand point analogy with (in)-elastic curves, it is clear that thoughtful market participants would not generally expect one share to trade at the same price as a million shares. However, with regard to an NAV, they have little choice given the current market dogma other than to value vast holdings for varying periods of time as a simple sum of the values of the constituent parts.

Example: Phantom Price framework for translucent assets

The general scheme of replacing point asset values with parameterized ones is especially effective for cases where point prices do not exist. In the Liquidity/Pricing continuum, the members of the translucent class of assets are characterized by multiple possible simultaneous prices for their fundamental observable units. Even for constant size and volume, translucent assets have a range of prices.

We are accustomed to being able to ask the price of many assets without having to make a commitment to buy or sell them. Small, uniform high-demand commodities like actively traded common stock or paper clips can, within bounds, be counted upon to have fairly definite market prices.

On the other hand, one-of-a-kind, expensive, indivisible, or infrequently traded items are often priced only when traded. A particular house or famous painting may be sold infrequently and it is difficult to know its precise market value at any given time. Particular categories of illiquid instruments form a kind of intermediate asset class between these extremes. With market prices possessing neither the transparency of common stock nor the opaqueness of a Michelangelo sculpture, these items may be priced regularly by a small collection of agents, e.g. broker-dealers, commercial pricing services who do not openly share information on pricing. As such, the true market value of these instruments is often viewed through a haze of approximate "indicative prices". Such valuations give the investor valuable approximate information without the sharp resolution of specific prices needed to compute performance statistics. Such instruments form a natural asset class, which can best be described as "translucent".

Consider, for example, the case of a mortgage derivative portfolio with assets that are priced using multiple "quotes" from active market makers. The table below illustrates the range of prices typically received for such assets:

Table 1

Translucent Pricing

MBO - 2000 Year End Dealer Quotes

	Security	#quotes	Min	Avg	Max	Range %
1	FHR 2064 G	5	54.31	78.47	85.62	39.90
2	FNR 99-15 SB	4	80.05	86.78	91.25	12.91
3	FHR 2131 SH	4	73.86	76.21	78.18	5.67
4	FHR 2136 SD	4	74.61	77.84	80.76	7.90
5	FNR 98-52 SA	5	78.20	81.57	86.05	9.62
6	FHR 2156 TS	5	78.94	88.33	94.30	17.39
7	FHR 2138 SB	5	82.36	88.24	92.30	11.26
8	GNR 00-7 ST	5	98.13	103.88	109.36	10.81
9	FNR 00-6 SK	5	99.30	108.17	115.03	14.55
10	FHR 2122 SD	5	5.34	6.05	7.52	35.89
11	FHR 2138 KS	5	4.73	5.12	6.19	28.36
12	FHR 2145 MS	5	3.88	4.61	5.28	30.48
13	FHR 2136 S	5	5.27	5.77	7.22	33.86
14	GNR 99-11 SC	5	3.77	4.71	5.33	33.21
15	GNR 99-30 SA	6	4.25	4.87	7.66	69.96
16	GNR 00-1 SD	5	0.58	0.68	0.84	38.78
17	GNR 99-40 SB	5	1.95	2.36	2.83	37.09
	Portfolio Millions		4.451	4.970	5.510	21.32

As Asness et. al. put it: "The absence of these prices may leave hedge funds with 'flexibility' in how they mark their positions for monthend reporting. In some cases, hedge funds use the last available traded price and in others (often with hard-to-price over-the-counter securities) they guess-timate the price, perhaps based on a model, along with broker-dealer input, in order to value their portfolio for month end performance reporting." One way of addressing the problem is by excluding the illiquid asset class from portfolios, which seek to report risk-reward statistics. Another possible approach is for portfolio managers to simply impute the prices that they claim would likely result from the sale of an illiquid asset. For obvious reasons, neither of these "solutions" is very appealing. There is, fortunately, another approach to this problem, which appears to be more fruitful when considering statistics, which aim to capture relationships between risk and reward (e.g. Sharpe ratios).

In the phantom pricing framework we discuss below, we have chosen to view an absence of consensus pricing as an additional source of risk for which investors must be able to expect higher returns than would accrue in the absence of such illiquidity.

Phantom Price framework for evaluating pricing risk for translucent assets.[4]

We start by stating some desirable properties of this framework:

1. Transparency: Reveals risk without revealing positions.
2. Provides analogs of all standard measures such as the Sharpe Ratio, which collapse to the usual measures in the absence of Translucent Assets.
3. Allows measurement of a Hedge Fund Manager's valuation bias.

We then propose the following thought experiment:

Figure 1

Price Quotes forming Market Beliefs

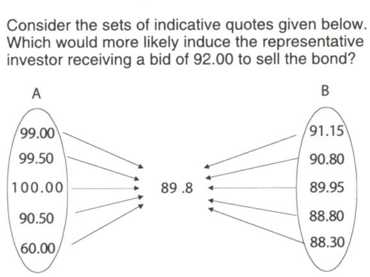

Figure 2

Phantom Prices

$$\mathbb{P}(v_t) = u_t^{-1}\left(\int_{\mathbb{R}_+} u_t(x)\ \mathcal{B}_{v_t}(x)\ dx\right)$$

$vt =$ Vector of indicative prices at time t.

$ut =$ Utility function representing an investors risk preferences.

$\mathcal{B}_{v_t} =$ Probability distribution that the investor believes generated the sample vt.

[4] In this article the authors attempt to give the story behind "the making of the movie" as this reveals the concerns and thoughts of practitioners in the field that led to phantom prices. The phantom pricing "movie" itself premiered in the June 2002 issue of Risk and is pitched to the slightly more mathematically inclined.

Presented with set A, many people would be unhappy with the 92.00 bid. Conversely, presented with set B, only the greedily optimistic would likely be displeased with a 92.00 bid. By design, the average price of either set is 89.8. The first lesson here is that both the dispersion and variance of the quotes shapes one's beliefs of the market more than the mean. The second lesson is that in the context of multiple quotes, the average price does not represent the views of most investors, which is just another way of saying that investors are rarely risk-neutral.

The essential idea at the heart of the phantom price framework is that one should formalize the ingredients of the above thought experiment. The mathematical underpinnings are the classic Risk-Calculus of Von-Neumann Morgenstern. The idea is to integrate a sub-utility function against a probability distribution. The result of the integration along with the inverse of the sub-utility function provides a "certainty equivalent" to the uncertain situation faced by an investor. It answers the question posed above: "what is the lowest single definite price I would accept for this asset in exchange for the ambiguous situation of the multiple indicative prices if my intention were to immediately liquidate the security in question?" In the most general terms the basic ingredients look like figure 2.

To calculate actual values of phantom prices from existing market data one has to pick a probability distribution and a utility function. Once these are defined, a simple integration and inverse mapping gives a risk-tolerance parameterized phantom price that is dependent on the price quotes. One can make any performance statistic calculated from price data into a *phantom* statistic that converges to the standard one if all the price quotes are the same.

A very interesting application is to recalculate the NAV series for a hedge fund with phantom prices instead of manager marks to see how much volatility smoothing is going on.

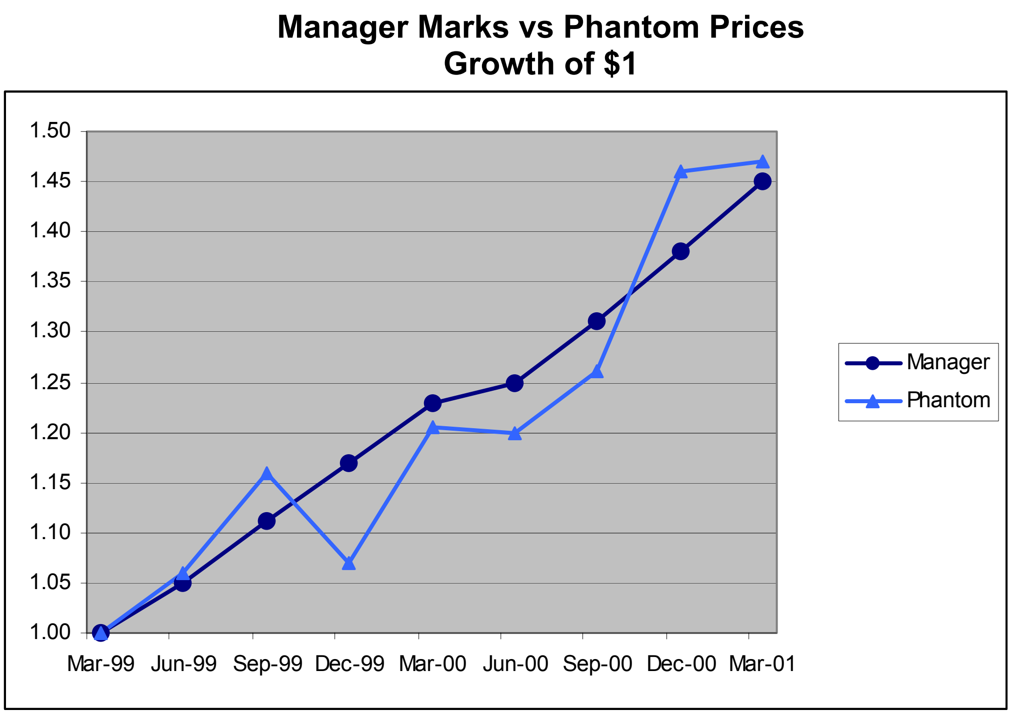

Figure 1: Manager's Marks vs. Phantom Price Performance Evaluation. The manager claims a Sharpe ratio of 2.34 while the phantom price evaluation reveals a significantly lower value of 0.45. The value for `a' has been chosen here to match the manager's average risk aversion in NAV reporting over the period in question.

Redemption Restrictions in Hedge Funds

The subject of performance and risk measurement for hedge funds is too vast to address here. Instead, what follows is a brief discussion of the liquidity factors that should influence hedge fund valuation and performance measurement. We include this section, as it is an obvious area where the one-size-fits-all scalar NAV concept seems most obviously artificial.

First, it is customary for hedge funds to have a lock-up period, usually lasting at least 1 year, during which investors cannot redeem their investments. Second, after the lock-up period the hedge fund usually restricts withdrawals to once a quarter with a long notice period.

Before the current popularity of hedge funds, hedge fund managers reported quarterly results at best, and some in fact only communicated with investors on request. The current imperative of monthly reporting has to be understood in the context of the lack of ability to actually monetise the NAV's reported by managers. The measurement of risk measured by such monthly statistics needs to be approached with appropriate caution. Perhaps the hedge fund world would do well to consider importing some insights from the banking industry, which uses vastly different techniques for valuing non-transferable loan assets as opposed to tradable bonds.

What does a 15% return with a 6% annualized standard deviation as calculated from monthly return series mean? Even if one were to ignore the statistical inference problems of impoverished data, one has to come to grips with the hypothetical nature of a hedge fund NAV. It is not a number that represents the cash-equivalent value of the hedge fund, because the investor cannot, in reality, turn his/her investment into cash. From this point of view, the NAV does not represent liquidity, as there is none to be had by the investor.

Using ideas from derivatives markets, one could state the option equivalent position for the manager vs. the investor. In many traditional money management frameworks, other than hedge funds, the manager is short, to the investor, an American put, exercisable daily, on the investments undertaken, with the strike price determined at the end of each day. In contrast, the hedge fund manager (typically) is short a put to the investor which has a one year exercise lock-out and thereafter has a quarterly exercise with the strike price determined as of the 45th day passed the exercise date. In addition, when the investor decides to exercise his option and put the fund back to the manager, he has to decide to exercise at a future price that the manager has some influence over. An interesting thought experiment is to try and figure out how much less this hedge fund ‘option’ is worth than its traditional counterpart built on the same underlying. This would put a price on the redemption restrictions, which are a form of liquidity risk in hedge funds.

Conclusion

From what we have argued above, it may be useful to think of illiquidity as coming in two basic flavors:

- <u>TYPE A</u>: Removable forms of illiquidity resulting from poorly structured markets which can and should be excised.
- <u>TYPE B</u>: Irremovable forms of intrinsic illiquidity which must be disclosed so that they can reward those investors (and only those investors) both able and willing to price and assume the risk.

Our concern is that some of the current anxiety over hedge-fund illiquidity may potentially do more harm than good. When hedge funds are over-regulated to make them appear artificially like mutual funds, we engage in a fools errand to remove intrinsic Type B illiquidity and while losing marvelous opportunities for profit in the service of greater market efficiency. Likewise, when hedge funds operate in illiquid markets which are poorly regulated, we may actually be pointlessly creating liquidity risks out of removable TYPE A inefficiencies and market failures.

Both of these responses go against the basic position taken in this paper which may be stated simply:

"The proper response to risk that cannot be removed is to price and trade it."

Lamentably, this is made unnecessarily difficult in the case of liquidity concerns when the true risks are not properly identified. As we see it, the principal danger to hedge funds is two-fold. First, when significant liquidity risks are hidden from investors (e.g. within an innocuous looking scalar NAV) in a fund lacking position level transparency, investors may be unwittingly assuming liquidity risks which they would wish to forgo in a truly transparent market. Second, because our abilities to price and model such liquidity risks have yet to mature, many of these risks appear more dire than they would if it were possible to view them through mature valuation models enjoying smaller ‘error bars’. In such situations, it must be expected that funds holding reasonably ‘well behaved’ illiquid instruments (e.g. translucent assets) may be attracting inefficiently low levels of investor capital compared to their true risk/reward profile.

Liquidity and valuation often appear as two sides of the same coin. There is no way to specify one without reference to the other. There are multiple variables that need to be considered together in the context of a market to properly quantify and price liquidity. The reward of developing the quantification schemes will no doubt lead to a change in the liquidity of the assets they attempt to quantify. Success in modeling any of the various types of liquidity risks will lead to a better understanding of the markets. In fact, our initial exploration into the land of uncertain NAVs seemed (at least to the authors) quite encouraging as we found tools that track valuation bias, and provide a scheme for fixing the "smoothing" of performance statistics by hedge funds.

References

Arfken, G., Mathematical Methods for Physicists, Second Edition, 1970, Academic Press, New York.

Arrow, KJ., Aspects of the Theory of Risk Bearing, 1965, Yrjo Hahnsson Foundation, Helsinki.

Asness, C., Krail, R., Liew, J. "Do Hedge Funds Hedge", forthcoming in journal of Portfolio Management, 2001.

Hanna, S. & Gutter, M., A Theory Based Measure of Risk Tolerance" 1999, Working Paper, Ohio State University.

Pratt, J.W, "RiskAversion in the Small and in the Large", Econometrica, Vol. 32, 1964, p. 122-136.

Weinstein, E. & Abdulah, A., "Phantom Prices: Valuation Transparency for Illiquid Assets" 2001, Working Paper available at www.transparentrisk.com.

Weinstein, E. & Abdulah, A," Towards Hedge Fund Transparency: Quantifying Valuation Bias for Illiquid Assets", Risk, Spring 2002
Weisman, A, "Dangerous Attractions: Informationless Investing and Hedge Fund Performance Measurement Bias", to appear in The Journal of Portfolio Management, Spring 2002.
BIS Study, "Market Liquidity: Research Findings and Selected Policy Implications" Basle, 3 May 1999
Cosandey, David, Adjusting value-at-risk for market liquidity", Risk, October 2001
Persaud, Avinash, " The Liquidity Puzzle", Risk, June 2000