How Risky Are Illiquid Investments?

A practical approach to estimating volatilities and correlations for non-traded assets.

Vineet Budhraja and Rui J.P. de Figueiredo, Jr.

In all asset allocation modeling, two primary principles must operate. First, the asset classes included have to be made comparable. Second, the inputs need to be as accurate or as reasonable as possible.

These two principles make it hard to include illiquid asset classes with long lock-ups and limited tradability when analysts are attempting to estimate long-run performance at the asset class level. One particular problem is that these asset classes—including leveraged buyouts, venture capital, and real estate—are not valued on the same basis as traditional and more liquid asset classes such as cash, bonds, and public equities.

On the one hand, this raises the issue of comparability. The fundamental problem is that while we measure the risk and correlation of the liquid asset classes on the basis of regular movements of publicly observable market prices, illiquid asset classes are not measured on a marked-to-market basis. The illiquid asset classes might have artificially smoothed return series, making them look both less variable and less correlated with other asset classes. On the other hand, because of inconsistencies in valuation and accounting methods, returns for illiquid asset classes, when they are reported, might be inaccurate.

The combination of these two problems means that inclusion of illiquid asset class forecasts of risk and correlation based on historically reported performance will lead to highly inappropriate allocations. Actually, it is not just that that allocation to these classes will be suspect, but because portfolio construction models are highly interactive, the weights suggested for *all* asset

Vineet Budhraja is a vice president at Citigroup Alternative Investments in New York City.
vineet.budhraja@citigroup.com

Rui J.P. de Figueiredo, Jr. is a research consultant to Citigroup Alternative Investments and an associate professor at the Haas School of Business at the University of California at Berkeley.

classes will may also be inappropriate.

Exhibit 1 shows as an example the annualized reported returns on the Nasdaq and venture capital, as reported by Venture Economics, from the beginning of 2000 through March 2002. Venture capital was reported as falling at a rate of roughly 7% annually while the Nasdaq was falling at a rate of over four times that. Most investors would look askance at these numbers; many expect the performance of these two markets to be more closely linked.

Consider the venture capital market. Two of the exit options for private investments, and thus determinants of valuation, have to do with the public equity markets: either acquisition by a public company, or an initial public offering. We might thus posit that the variation in one market should be reflected in the other. In the end, the long-run averages of reported returns on the Nasdaq and venture capital returns will likely equilibrate—at some point, the venture capital returns will have to be marked down if there is a sustained reduction in value.

Unfortunately, for those interested in constructing portfolios of these investments, or just in understanding their risk, the averages are not sufficient. Rather, since the emergence of modern finance, investors have understood that looking beyond averages to the *patterns* of returns is a crucial consideration in investment choice. And a pattern such as that shown in Exhibit 1 can be very misleading if we are comparing public and private market characteristics.

As Gompers and Lerner [2002] point out, such a pattern for venture capital, if taken on its face, will lead investors to believe that venture capital is both less risky and less correlated with traditional assets than it actually might be on a marked-to-market basis.

How to address this problem? The challenge is to develop an approach that lets us calculate volatility and correlation (to a range of other asset classes) using publicly available data.

One common approach is to use proxies for these asset classes. Recognizing that the data are flawed, this approach suggests using proxies such as real estate investment trusts for private real estate, or the Nasdaq or a small-capitalization index for venture capital. While this approach is easy to implement, it suffers from oversimplification.

It is true that these markets should be related, but using public markets as proxies perhaps goes too far; the public and the analogous private market are not substitutes. While the markets should be correlated, they should also have unique characteristics because of both sector and valuation biases and because of their reaction (for example) to liquidity shocks.

EXHIBIT 1

Reported Returns January 2000–March 2002

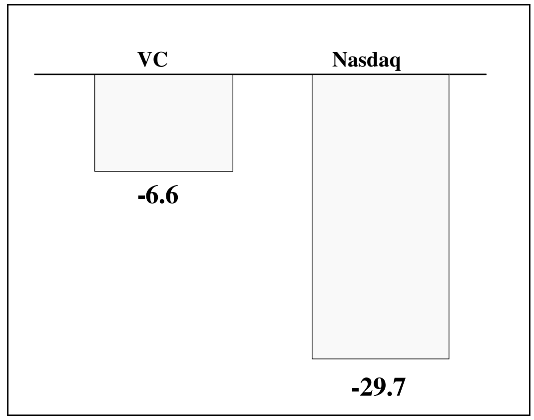

Pooled time-weighted IRR for VC and total return for Nasdaq.
Source: Venture Economics, Bloomberg.

A second way to deal with the problem of stale pricing has been to use data that have been aggregated over longer time periods. This approach focuses on time horizons over which valuations are reliably measured. As an example, while quarterly or even annual data are highly suspect, over longer time periods such as the vintage years in which most investments are realized one can accurately measure market-based returns. It is thus possible to examine the long-term volatility in the returns (and compare them, say, to those in traditional asset classes).

Indeed, one can even make inferences about annual volatility, with reasonable assumptions about independence in the returns over time. While these estimates are reasonable, such an approach suffers from the fact that it throws out much of the potential information in the data by aggregating over long time periods. Perhaps even more important, this approach makes calculation of correlations difficult, as it results in very few data points.

A third approach is to use underlying investments in something like venture capital and estimate their true values to construct an index. Quigley and Woodward [2003], for example, use a method for correcting for both intermittent pricing and selection biases in their sample to estimate an index for venture capital based on reported company valuations. Cochrane [2001] provides an alternative way to estimate parameters of returns in venture capital based on maximum likelihood, but only for exiting firms. Peng [2001] attempts to build two indexes for venture capital separately (for firms that have exited and those that have not), and then uses an arbitrary weight to combine the two.

Gompers and Lerner [2002] also suggest the use of indexes built up from a deal-level analysis, having corrected for a range of biases. Gompers and Lerner [1997, 2000] provide an excellent discussion of the issues related to private equity valuations.

The common element of all these approaches is that they use a microanalytic (deal-level) foundation as a means to estimate market returns for the asset class. The weakness of this approach for a practitioner is twofold. On the one hand, these studies have primarily focused on one asset class—venture capital—and have not been generalized to other illiquid asset classes. Second, and more important, without the underlying data, most investors constructing portfolios will not be able to actually use these results to tailor analysis of risk and correlation to their particular asset allocation problem.

We outline a method that works with readily available and public data at the macroanalytic level. While we do not actually observe marked-to-market valuations to calculate returns, there is arguably some information in the reported returns. To take advantage of this information, we formulate the problem with the reported return series as a measurement error problem. A solution is to decompose this error into its systematic and random components.

The key step is to provide a framework for understanding the error, which we posit can be related to serial correlation, public market effects, and private market effects. We use this theoretical framework to estimate the model and provide methods to calculate adjusted marked-to-market volatilities and correlations.

Our approach can also be thought of as an attempt to include information beyond the reported return data. That is, we attempt to incorporate both temporal and other market information as a way of making the original data more robust.

The results provide a dramatic reevaluation of the risks and correlations in the illiquid asset classes. We estimate the actual volatility of venture capital at 43.3% annually, treated on a marked-to market basis, compared to the historically reported volatility of only 17.4%. For leveraged buyouts and private real estate, the figures are 20.4% versus 11.0%, and 10.6% versus 3.3%.

Correlations may also rise dramatically, particularly for private equity classes. We estimate that venture capital is correlated with the S&P 500, for example, at 0.55, compared to the historically reported 0.44. Similarly, for leveraged buyouts we estimate the correlation with the S&P to be 0.75 rather than the reported 0.51. For real estate, however, there is very little effect, as one might expect; the recalculated correlation is 0.01 versus a historical correlation of 0.00.[1]

These results are estimates. We have had to make a number of assumptions and judgments on the nature of the random components of the variables, on the use of models that are not fully specified, and on the accuracy of external measures of some of the statistics we calculate. That said, by incorporating all the information in a clearly defined and theoretically justified way, from a practical standpoint this approach provides a better estimate we believe than other approaches that ignore this information.

Further, since our purpose is both to aid forecasting and test hypotheses, we should emphasize that improved prediction (rather than statistical significance) is the standard by which to choose models. Although the results are still estimates, and should be viewed with the appropriate caution, we believe they provide new and more useful inputs for the purposes of asset allocation.

Finally, note that as our methods rely on an assumption of structural stability—for example, that the correlation structure between illiquid and liquid assets is relatively stable over time—we do not view the approach as suited to the analysis of all asset classes. Since hedge funds, for example, are often noted to have time-varying and changing correlations with other markets, we do not believe these assets necessarily fall within the domain of this study.

DATA

Data on venture capital and leveraged buyouts are drawn from the Venture Economics database of pooled time-weighted quarterly returns. For real estate, we use the NCREIF National Property Index (NPI) real estate quarterly return series. For traditional asset classes, we use total return data drawn from various sources for the S&P 500, 90-day T-bills, the Citigroup High Yield index, the Nasdaq index, the Russell 2000, and the NAREIT index.

In all cases, the series date from January 1988 through September 2003. We choose this period for two reasons. First, as one goes farther back in history, particularly with the private equity classes, samples for the aggregate return series become very small. Second, before 1988, the problems of sample size and selection biases were more acute.

Finally, for the private equity categories, we add two data sets. By its original method, Venture Economics included all reporting funds in the sample for a quarter. In May 2003, it changed the sampling method to exclude funds on a quarterly basis that do not report residual values. We find this to have created sampling instability that makes it difficult to base estimates on stable samples prior to 1993. After 1993, the two samples appear to converge, as the non-reporting problem appears to have diminished. Therefore, we use the data calculated by the earlier method prior to 1993 (which increases sample sizes substantially).[2]

CONCEPTUAL FRAMEWORK: ERRORS IN MEASUREMENT

We follow three principles in choosing appropriate models. First, we want the models to be as simple as possible while still capturing the theoretically posited effects we expect.

Second, we recognize these models will not necessarily be fully specified, and in some cases the fits of the model can be improved. In this case, we will prefer models based on our a priori theory of the structure of market relationships; in other words, we do not see this exercise as subject to a statistical horse race between factors.

Finally, we will accept some asymptotic biases in order to hone our predictions. In general, construction of these models will require us to trade off some biases in exchange for incorporation of particular information. Where there may be biases, we believe they are likely small compared to the gain of using the information against a null model of the unadjusted historical data.

For a more accurate measure of the true risk and correlation in illiquid asset classes, we start by positing that the observed return series are measured with error. In particular, we assume

$$R_{it}^{T} = R_{it}^{O} + \varepsilon_{it} \tag{1}$$

where R_{it}^{T} is the true return, $i = (1, \ldots, N)$ indexes asset classes, $t = \{1, \ldots, T\}$ indexes periods, R_{it}^{O} is the observed or measured return for asset class i at time t, and ε_{it} is an error term. The problem then is to estimate the vector ε_{it}.

To gain a better understanding of the measurement error, we take a number of steps. First, because of staleness in pricing, we assume that part of the error is due to *autocorrelated disturbances*.[3] Second, we assume that private markets are *correlated with publicly traded markets*. Third, we also assume that there are *private market shocks* that affect all private markets in a similar way. An example would be shocks to overall market liquidity that affect private markets in different ways or degrees from public markets. Finally, we assume there are some residual idiosyncratic movements, which are noise.

Notice that the first three of these features might be termed *systematic* and the fourth *random*. To improve our estimate of the true return, we want to deal with each systematic component.

Our estimation strategy is to first deal with the first component, autocorrelated disturbances, and then to address the second two jointly in order to obtain estimates of the exposure of private markets to related public and private markets. We then use these estimates to infer, for example, the measurement error and idiosyncratic errors, which we then use to form estimates of correlations and volatilities.

In practice, one could model all three at once, but for purposes of clarity—particularly as to the contribution of the different error components—we do this in stages. The results are roughly invariant to this choice.

Autocorrelated Disturbances

Our first task is to remove autocorrelation in the data. While there are a number of potential reasons for serial correlation in financial return data—short sales through options, time-varying leverage, time-varying expected returns, and incentive-based fees—in the case of the asset classes we are dealing with, this autocorrelation is likely induced by accounting practices requiring particular treatments of the valuations and smoothing of returns (see Getmansky, Lo, and Makarov [2003]).

Whatever the reason, it is sensible to correct for this distortion in the quarterly data. Without estimating the covariance between periods, the volatility of a T-period return will not be appropriately estimated by the standard $\sqrt{T}$ transformation of the subperiod volatility. Only by explicitly estimating the subperiod correlation can we obtain an accurate estimate of the multiperiod volatility.

Our approach is one way to deal with this aggregation bias in estimating volatilities. As we will show, the estimates based on realized aggregated data over longer time horizons confirm this point.[4]

We estimate an autoregressive [AR(k)] model. For real estate and venture capital, we use AR(1) processes, and for leveraged buyouts, we use AR(2) processes. This choice is based on an evaluation of the autocorrelation functions of each of the data series. We thus estimate equations of the form (for example, for an AR(1) model):

$$R_{it}^{O} = \rho_o + \rho_1 R_{i(t-1)}^{O} + \upsilon_{it} \tag{2}$$

where υ_{it} is assumed to be a normally distributed error with mean zero and variance σ_{υ}^2.

Exhibit 2 reports estimated coefficients of the lagged returns. Because we combine two data sets for leveraged buyouts and venture capital, we present the autocorrelations separately for the two periods, although pooling yields very similar results.

The results of Exhibit 2 indicate substantial autocorrelation in the data. Given these estimates, we can now account for the autocorrelation in the return series. Specifically, following Georgiev [2002] and Desouza and Gokcan [2004], we can calculate an autocorrelation-corrected return series according to:

$$R_{it}^{A} = \frac{R_{it}^{O} - \rho_L R_{i(t-1)}^{O}}{1 - \rho_L}$$

where R_{it}^{A} is the adjusted series, and ρ_L is the coefficient of the relevant lag from Equation (2). These adjustments allow calculation of a new variance for each series. The adjusted series is shown in Exhibit 3.

It is apparent that removal of serial correlation by itself significantly raises the estimate of the volatility (and therefore risk) in these series. Correcting for autocorrelation increases the volatility of venture capital approximately 1.7 times, of leveraged buyouts approximately 1.4 times, and of real estate approximately 2.3 times.

Marking to Market Based on Public and Private Effects

For a more explicit marking to market of the data, we first attempt to estimate the factor exposures between markets. Again, our rationale is that since the illiquid data are not marked to market, but it is also hypothesized that they should be correlated somewhat with those markets, we want to develop a method for incorporating the information in correlated markets.

Then, making assumptions about measurement error, we use this information to estimate marked-to-market correlations and volatilities.

Methodological Framework

Consider first why a particular asset class might be related to both public and other private markets. In the case of the public markets, this link is at least twofold. First, these markets are subject to common shocks or variations. For example, as the demand for software increases, both public and private software companies benefit. Second, these markets are often explicitly linked economically.

EXHIBIT 2
Estimated Coefficients

Asset Class	Coefficient
Venture Capital	
1988-1992	0.43*
	(0.24)
1993-2003 (Q3)	0.54**
	(0.13)
Leveraged Buyouts	
1988-1992	0.35*
	(0.22)
1993-2003 (Q3)	0.29*
	(0.15)
Real Estate	0.69**
	(0.09)

*$p < 0.10$; ** $p < 0.05$ (two-tailed)* indicate the significance level.
Standard errors in parentheses.

EXHIBIT 3
Standard Deviation of Illiquid Returns (annual)

Asset Class	Historical	Adjusted for Autocorrelation
Venture Capital	17.4%	29.4%
Leveraged Buyouts	11.0%	15.1%
Real Estate	3.2%	7.6%

While some of the variation in the public markets should be reflected in corresponding private markets, private markets do not march in lockstep with their public market counterparts. One reason is that asset liquidity risk differs across public and private markets. As illiquid private investments are often in markets that do not trade, changes in the underlying liquidity of the market might affect the private market in a way that does not affect the public market. Therefore, we also want to account for potential factors that are not captured by movements in the public markets.

Putting these two points together, we can then specify a model for returns. The general model is a system of equations in which, say, for two return series:

$$\begin{aligned} R_{1t}^{T} &= \alpha_1 + \boldsymbol{\beta}_1 \mathbf{x}_t + \gamma_1 R_{2t}^{T} + \eta_{1t} \\ R_{2t}^{T} &= \alpha_2 + \boldsymbol{\beta}_2 \mathbf{z}_t + \gamma_2 R_{1t}^{T} + \eta_{2t} \end{aligned} \tag{3}$$

where R_{it}^{T} is the market return, the vectors $\mathbf{x}$ and $\mathbf{z}$ are factors related to the specific return series, and η_{it} is a noise term.[5]

If we could estimate the parameters and the error in Equation (3), we would be able to calculate the "exact" returns. Unfortunately, recalling (1), we observe the return series with a random error. In other words, we assume a further additive error takes us from the autocorrelated-corrected series to the market return series:[6]

$$R_{it}^{T} = R_{it}^{A} + \varpi_{i} \tag{4}$$

where ϖ_{it} is a noise term. By construction, we are assuming that the observed variance is lower than the true variance. In other words, the model makes the assumption that the measurement error smooths the returns (i.e., the error is correlated with the true value).

If we substitute (4) into the equations in (3), we have:

$$\begin{aligned} R_{1t}^{A} &= \alpha_1 + \boldsymbol{\beta}_1 \mathbf{x}_t + \gamma_1 R_{2t}^{A} + \gamma_1 \varpi_{2t} + \eta_{1t} - \varpi_{1t} \\ R_{2t}^{A} &= \alpha_2 + \boldsymbol{\beta}_2 \mathbf{z}_t + \gamma_2 R_{1t}^{A} + \gamma_2 \varpi_{1t} + \eta_{2t} - \varpi_{2t} \end{aligned} \tag{5}$$

Now we have a model that uses all observable variables rather than variables that are unobservable. The problem is whether we can consistently estimate the parameters of this model. Two issues in particular are relevant.

Our first worry is that there is measurement error in the right-hand side variables. In other words, the return on the right-hand side is measured with error, so the coefficient is also acting on the error term for that return. This creates a problem for estimation of the coefficients, as the regressor is correlated with the disturbance term, and therefore violates one of the basic assumptions of ordinary least squares (OLS) regression.

Another issue with recovering the parameters from (5) is that the returns are themselves endogenous. In other words, some of the variables affect each other jointly. This will also lead OLS estimates of coefficients to be inconsistent, as it will provide another reason the error terms in each equation will be correlated with some of the explanatory variables.

These two points mean we need an alternative estimation procedure to traditional OLS. We choose a two-stage instrumental variables estimator (*two-stage least squares*). For a 2SLS model to be identified, we need enough excluded exogenous variables from each equation that we can use as instruments. In this case, we can use the variables **x** and **z** as instrumental variables in the equations from which they are omitted.

We first run the OLS regression:

$$\begin{aligned} R_{1t}^{A} &= \theta_1 + \boldsymbol{\varphi}_1 \mathbf{x}_t + \xi_{1t} \\ R_{2t}^{A} &= \theta_2 + \boldsymbol{\varphi}_2 \mathbf{z}_t + \xi_{2t} \end{aligned} \tag{6}$$

From (6), we then obtain predicted values of the endogenous return series, such that:

$$\begin{aligned} \hat{R}_{1t}^{A} &= \hat{\theta}_1 + \hat{\boldsymbol{\varphi}}_1 \mathbf{x}_t \\ \hat{R}_{2t}^{A} &= \hat{\theta}_2 + \hat{\boldsymbol{\varphi}}_2 \mathbf{z}_t \end{aligned} \tag{7}$$

As shown in Greene [1993], using the predicted values from (7) and substituting them into (5) for the endogenous values of the returns will yield asymptotically unbiased (or consistent) estimates of the parameters α, $\boldsymbol{\beta}$, and γ.

INITIAL MODEL RESULTS

We use the estimates from Equation (5) to better understand the connections between related public and private markets and later to help us formulate estimates of volatilities and correlations with a broad set of asset classes. For VC, LBOs, and real estate, we use the exogenous and endogenous markets shown in Exhibit 4. These are based on our ideas of the factors or markets that should be most closely related to the private market.

While we include markets that might not seem related (such as using real estate in the venture capital equation), this is for the specific purpose of modeling asset liquidity shocks. To the extent these shocks are uncorrelated in different markets, the model will allow for lower (zero) weights on these asset classes; inclusion simply reduces efficiency. In this case, we also examine full and reduced models to navigate this trade-off.

Finally, as Craft [2001] observes, the nature of the REIT markets changed dramatically during the late 1980s and early 1990s. Because of that, we specify a model for private real estate that allows for a slope change before and after 1993 through the use of an interaction between a dummy variable that indicates those two distinct periods and REITs.

EXHIBIT 4
Model Specification for Equation (5)

Asset Class	Exogenous Markets	Endogenous Markets
Venture Capital	Nasdaq Russell 2000	Leveraged Buyouts Real Estate
Leveraged Buyouts	S&P 500 Russell 2000	Venture Capital Real Estate
Real Estate	NAREIT Citigroup High Yield 90-Day T-Bills	Venture Capital Leveraged Buyouts

EXHIBIT 5
2SLS Estimates

Independent Variable	VC (1)	VC (2)	LBO (1)	LBO (2)	RE (1)	RE (2)
Nasdaq	-0.26 (0.54)	-0.38 (0.36)				
Russell 2000	0.52** (0.37)	0.55** (0.36)	-0.22 (0.27)	-0.17* (0.16)		
S&P 500			0.42** (0.30)	0.44*** (0.26)		
NAREIT					-0.29 (0.32)	-0.21* (0.20)
Citigroup High Yield					0.28* (0.22)	0.24** (0.17)
90-Day T-Bills					2.65*** (1.41)	2.67*** (1.27)
1993+					0.03* (0.02)	0.03*** (0.01)
NAREIT 1993+*					0.46* (0.40)	0.36** (0.22)
VC			0.19 (0.29)	0.13* (0.12)	0.04 (0.11)	
LBO	1.54** (1.01)	1.70*** (0.85)			-0.29 (0.42)	-0.17* (0.13)
Real Estate	-1.23 (4.21)		0.74 (2.86)			
Constant	-0.02	-0.04	0.01	0.02	-0.03	-0.04
F-statistic (p)	4.86 (0.00)	6.34 (0.00)	3.36 (0.02)	5.48 (0.00)	1.71 (0.12)	2.44 (0.04)

**p < 0.15, ** p < 0.10, *** p < 0.05 (one-tailed).*
Dependent variables corrected for autocorrelation; standard errors are in parentheses.

Exhibit 5 presents the results of the second-stage estimation of a series of models for each asset class. A number of elements stand out.

First, the private markets do not necessarily covary very strongly together, with the exception of the effect of leveraged buyouts on venture capital. Indeed, when two endogenous variables are included, the effect of the efficiency loss from the 2SLS approach seems particularly severe, as few of the coefficients are estimated very precisely. When models eliminate one of the endogenous variables, estimates are much more precise, although in the leveraged buyout and real estate models these effects are marginally significant in one-tailed tests.

This is not to say there is no effect, but in general the effects are weak at best. This could in part be due to the inefficiency of the estimator, but when we execute the model proposed in (5), most of the effect will be felt by the public markets, as the coefficients are much higher for these markets.

Second, the major public market indexes are all significant predictors in the model. There is substantial positive covariation—controlling for other factors—between venture capital and small-cap public equities, between leveraged buyouts and large-cap public equities, and between private real estate and short-term interest rates and REITs from 1993 onward (with a less robust impact of credit as represented by high-yield). This means that the public market indexes certainly provide additional information.

To further evaluate these models, we use two sets of additional criteria. First, how good is the fit of these models (including the less significant coefficients)? And, second, how good are the instruments in the first-stage regressions?

Model fit is not the only criterion one should apply to instrumental variables models, as the models combine two projections rather than just a single one, as in a typical regression framework. Indeed, they are often believed to be of limited value. That said, in all cases, the F-statistics of the reduced models show a very good fit; the p-values are all significant at less than a 5% level of significance.

The second set of criteria relies on the first-stage results. Staiger and Stock [1997] suggest that the bias in the estimates induced by the 2SLS procedure is proportional to the inverse of the F-statistic from the regression of the left-hand side variable on the instruments. In this case, these biases appear to be fairly small; the F-statistics are 5.5, 2.7, and 2.2, respectively for venture capital, leveraged buyouts, and real estate.

We can also evaluate the quality of the instruments following the approach suggested by Bollen [1996], where regressions of the endogenous variables on the instruments should yield R-squares higher than 0.1. In this case, the R-squares for all three are higher: 0.37, 0.22, and 0.25, respectively, for venture capital, leveraged buyouts, and real estate.

While the models appear to fit reasonably well by these standards, then, the results do indicate that both the model specification and the quality of instruments might be improved.

ESTIMATING THE CORRELATIONS

How might we use these models to estimate the correlations between these asset classes and other markets? From (7) above, we have predicted values of each of the return series. This gives us a basis for tracking the general patterns of the return series. Yet including only the fitted values will overfit the data, as this excludes the second part of (3), the error term, which is also volatile.

If we treat the predicted value as the base component of the variation in the returns, we still need an estimate of the error in the regression in order to estimate correlations. In other words, because the error in the regression reflects both the idiosyncratic error (as represented by the η_{it}) and measurement error (as represented by the ω_{it}), we need some additional identifying information in order to separate the two.

In this formulation, we want to estimate the quantity $R_{it}^T = \hat{R}_{it}^T + \eta_{it}$ where $\hat{R}_{it}^T$ is the fitted value from (3) and η_t is a mean zero white noise error. If for some other return series k, we use the correlation between $\hat{R}_{it}^T$ and R_{kt}^T, we will systematically underestimate the absolute value of the correlation since

$$\left|Corr(\hat{R}_{it}^T, R_{kt}^T)\right| > \left|Corr(\hat{R}_{it}^T + \eta_{it}, R_{kt}^T)\right| \tag{8}$$

if the distribution of η_{it} is non-degenerate.

If we knew the distribution of η_{it}, we would be able to explicitly calculate the covariance (and thus correlation) of the illiquid time series with other return series and each other. The true correlation of two series i and k will be:

$$\rho_{ik}^T = \frac{Cov(\hat{R}_{it}^T + \eta_{it}, R_{kt}^T)}{SD(\hat{R}_{it}^T + \eta_{it})SD(R_{kt}^T)} \tag{9}$$

What we observe from the fitted values is the correlation:

$$\rho_{ik}^O = \frac{Cov(\hat{R}_{it}^T, R_{kt}^T)}{SD(\hat{R}_{it}^T)SD(R_{kt}^T)} \tag{10}$$

Since the error term is noise, it is uncorrelated with either the fitted value or other series. Therefore, we can rewrite (10), the true correlation, as:

$$\rho_{ik}^T = \frac{Cov(\hat{R}_{it}^T, R_{kt}^T)}{SD(\hat{R}_{it}^T + \eta_{it})SD(R_{kt}^T)} \tag{11}$$

which implies

$$\frac{\rho_{ik}^T}{\rho_{ik}^O} = \frac{SD(\hat{R}_{it}^T)}{SD(\hat{R}_{it}^T + \eta_{it})} \tag{12}$$

Using (12), we can calculate the correlation of the actual return with a third return series, indicated by l, as:

$$\rho_{il}^T = \frac{Cov(\hat{R}_{it}^T, R_{lt}^T)}{SD(\hat{R}_{it}^T)SD(R_{lt}^T)} \frac{\rho_{ik}^T}{\rho_{ik}^O} \tag{13}$$

Finally, we can rewrite (13) as:

$$\rho_{il}^T = \rho_{il}^O \frac{\rho_{ik}^T}{\rho_{ik}^O} \tag{14}$$

The expression in (14) means that if we have a measure of the ratio of the overfitted correlation to the actual correlation for any series, we can apply a constant factor to correct the overfitting of all correlations. In fact, we do have such information.

Using an alternative methodology that estimates marked-to-market valuations for a sample of underlying investments, Gompers and Lerner [2002] estimate that the correlation between private equity and the S&P 500 is 0.74. If we assume that the approach provides an accurate assessment of this single correlation, then the data indicate we should apply a factor of 0.93 to the correlations to include the noise in the assessment of correlations. We can then, by (14), apply this factor to all the correlations between illiquid and other asset classes, because we assume the errors are the same across the illiquid asset classes.

When we make these adjustments to the fitted data (which provide the underlying basis for the correlation), we end up with dramatically adjusted correlations. Exhibit 6 reports three types of correlation with the S&P 500. In the first column is the historically reported correlation, based on completely unadjusted data. The second column reports the correlation based only on the initial fitted values of the regressions. The last column reports the correlation after correcting for overfitting as we have described.

Two points stand out. On the one hand, we have dramatically raised our assessment of the true correlation between private equity and the public markets corresponding to it. Real estate, on the other hand, starting from a low base, continues to be relatively less correlated with the S&P 500.

Second, if we take only the results from the regressions, the result might be too high—these markets should be correlated, but they are not identical. Therefore, after we make our adjustments, the correlations are slightly lower, if still much higher than previously believed from the historical data.

EXHIBIT 6
Correlation with the S&P 500

Asset Class	Historical	Fitted Values	Fitted Values Adjusted for Overfitting
Venture Capital	0.45	0.59	0.55
Leveraged Buyouts	0.51	0.81	0.75
Real Estate	0.00	0.02	0.01

EXHIBIT 7
Standard Deviation of Illiquid Returns (annual)

Asset Class	Historical	Adjusted for Autocorrelation	Adjusted for Model Estimates
Venture Capital	17.4%	29.4%	43.3%
Leveraged Buyouts	11.0%	15.1%	20.4%
Real Estate	3.2%	7.6%	10.6%

ESTIMATING THE VOLATILITY

To estimate the volatility of private assets, we can use the estimates of (3) above. We want to compare the observed variance of the autocorrelation-corrected series to the true one. From (4), we know the variance of the return series for asset *i*:

$$V(R_{it}^{T}) = V(R_{it}^{A}) + V(\varpi_{it}) \qquad (15)$$

since we assume the error term is uncorrelated. This in turn implies that the difference between the true variance and the variance of the autocorrelation-corrected series is simply:

$$V(R_{it}^{T}) - V(R_{it}^{A}) = \sigma_{\varpi}^{2} \qquad (16)$$

where σ_{ϖ}^{2} is the variance of the measurement error.

One question that emerges with (16) is what we should use for the term σ_{ϖ}^{2}. Once again, the problem is one of identification. As is apparent in (5), the error we estimate consistently, even after removing the measurement error from the right-hand side variables, is still a combination of idiosyncratic and measurement error. We propose using the *observed* standard error of the regression (whose corresponding variance we will call σ_{R}^{2}), since it is a conservative estimate of the measurement error.[7]

Notice that the adjustment to the volatility is a function of the overall model fit (as represented by σ_{R}^{2}). In other words, as the model fits better, there is less adjustment (beyond the autocorrelation adjustments), which comports with our intuition that if the adjusted series tracks well with the theoretical correlates, there is less error to be accounted for.

According to this estimator, and the models estimated in Exhibit 5, we have a new estimate of the variance of the return series. We show these in Exhibit 7.[8]

To assess the validity of our approach, we also look at the one set of return series that we feel confident accurately reflects realized returns. Namely, if there is intertemporal smoothing of returns in the quarterly series, another way of examining this problem is to look at the longer-period returns, such as the vintage year returns reported by Venture Economics for leveraged buyouts and venture capital. We compare the volatility of these returns to the volatility of returns for the S&P 500 calculated on a seven-year vintage year basis.[9]

Exhibit 8 reports this analysis compared to estimated volatilities over the same period using our method for private equity and venture capital. According to the two-stage least squares method, the ratio of venture capital volatility to S&P 500 volatility is 3.0, while according to the aggregated vintage year data, the ratio is 3.7, indicating a similar order of magnitude difference. For leveraged buyouts, the figures are even closer—1.4 and 1.1, respectively. The similarity of these two methods indicates that the two-stage least squares estimates should be reasonable.

Of course we have assumed that the bias in the estimates of the coefficients of the models is not that great; that the standard error of the regression is a reasonable estimate of σ_{ϖ}^{2}; and that the models and instruments are appropriate specifications. While we recognize that these issues almost certainly impact our results, from a practical point of view, we think that on balance these estimates will be much closer to the true values than the historically reported ones.

CONCLUSION

Investors face innumerable choices when they are determining how to deploy their resources. They must make comparisons: On the margin, which investment will make the greatest improvement in a portfolio? To answer this type of question, a crucial parameter is the risk entailed in a particular investment. Investors have to be able to measure risk in a comparable way across the full set of investment options.

While market prices (and therefore market returns) are regularly observable in publicly traded assets such as equities and fixed-income, the same cannot be said for

EXHIBIT 8
Comparison of Alternative Volatility Calculations

Asset Class	Annual Basis	Vintage Year Basis
Venture Capital	43.3%	21.9%
Leveraged Buyouts	20.4%	6.3%
S&P 500	14.5%	5.9%

illiquid assets such as private equity and real estate. Unobservable values raise significant problems for investors who want to measure risk in a comparable way using marked-to-market returns.

We have outlined a method to estimate such comparable measures of volatility and correlation in illiquid asset classes. The results lead to a dramatic reevaluation of what we understand to be both the risk-adjusted returns and the diversification benefit in these asset classes.

Despite this advance, some questions remain unanswered. First, we do not know the most appropriate instruments (or factors). Despite the better predictions of our simple model, using broader or more appropriate instruments could yield even stronger results.

Second, our assumptions require further testing and evaluation. Of particular concern is the specification of the model—both in terms of time-varying coefficients and the regressors to use—as well as the choice of the estimates of the error.

Finally, instrumental variable regressions are generally fairly inefficient and not very robust. The coefficients in these models are somewhat unstable in both subsamples and when additional regressors are included or deleted.

Despite these limitations, our results represent a significant departure from conventional wisdom as to these asset classes. And, more important, our method provides one important element in a new way for investors to incorporate these asset classes in an overall portfolio construction process.

ENDNOTES

The authors thank Neil Brown, Jacques Chappuis, Stephen Coyle, Clifford Desouza, Ahmed Fahour, Suleyman Gokcan, Paul Gompers, Martin Lee, Joshua Lerner, Ryan Meredith, Mikhail Smirnov, and Pablo Spiller for helpful comments and conversations.

[1]The period for all these data is 1988 through third quarter 2003. Sources are Venture Economics pooled average time-weighted returns for leveraged buyouts and venture capital, the National Council of Real Estate Investment Fiduciaries NPI for real estate, and total returns for the public indexes. One venture capital observation is excluded.

[2]Following Hadi's [1992, 1994] method for identifying multivariate outliers, we also exclude Q3 1999 from our analysis of venture capital. This observation is roughly eight standard deviations away from the mean in the data, and there are a number of other reasons we make this move. First, for estimating second moments, inclusion will significantly increase volatility, but all on the positive side. Second, this observation seems to be the result of measurement error that violates our basic model parameters. Finally, there is significant doubt that this result is repeatable.

[3]This effect is similar to that observed in hedge fund strategies where pricing discretion also can lead to artificial smoothing of returns (Desouza and Gokcan [2004], and Getmansky, Lo, and Makarov [2003]). The same approach has also been applied to real estate data (Geltner [1993], Georgiev [2002], and Giliberto [2003]).

[4]An alternative approach to generating non-correlated returns (to correct for serial correlation) would be to reestimate the returns using the average or trend return, and then randomly draw from an error distribution with prespecified volatility. This will generate: 1) a given volatility by construction; and 2) uncorrelated estimated returns. It will also lead to a gaming of the correction for serial correlation. The new series may have the desirable characteristics of both the correct mean or trend and no serial correlation; the volatility will be completely determined by construction. Our approach is different for two reasons. First, we use moments beyond the average from the original data. Second, the serial correlation correction is derived from the data and not based on a simulated sample.

[5]Equation (3) might strike some at first glance as a standard factor model. There are two important differences from the standard model. First, some of the variables are *unobservable* (or latent) or in this case, *measured with error*. Second, the inclusion of dependent variables from one equation as an explanatory variable in other equations means these models suffer from *endogeneity* (or feedback), which is a significant departure from the traditional factor approach.

[6]Note this means we are assuming there is variance added besides that in the autocorrelated series.

[7]This estimate is essentially an upper bound, or a conservative estimate, as it assumes that all of the regression error is measurement error, which it is not. An alternative would be to use information concerning the idiosyncratic error from the analysis of correlations to identify the measurement error. Namely, from (3) above we have an estimate of

$$\sigma_\eta^2 = (\frac{1}{k^2} - 1)V(\hat{R}_{it}^T)$$

where k is the ratio of the correlations in (14). Further, the following also holds:

$$R_{it}^A + \varpi_{it} = \alpha_i + \beta_i x + \gamma_i R_{jt}^A + \eta_{it}$$

This implies that the regression we actually run (and obtain errors for) is simply (noting that the two-stage process gives consistent estimates without measurement error of the returns for series j):

$$R_{it}^{A} = \alpha_i + \beta_i x + \gamma_i R_{jt}^{A} + \eta_{it} - \varpi_{it}$$

Thus, the observed standard error of the regression is the square of the sum of the variances of the error terms, or

$$\sigma_R^2 = \sigma_\eta^2 + \sigma_\varpi^2$$

where σ_R^2 is the square of the standard error of the regression.

This then implies we can recover the variance of the measurement error:

$$\sigma_\varpi^2 = \sigma_R^2 - \sigma_\eta^2$$

In other words, we can identify σ_ϖ^2 since we have an estimate of the overall model fit and the true model fit based on our calculation of the correlations. Using this alternative estimate of the variance of the measurement error yields very similar results to those reported in Exhibit 7. This alternative estimate reduces the volatilities by 160 basis points, 30 bp, and 20 bp for venture capital, leveraged buyouts, and real estate, respectively. Since the difference is so minor, in the interest of parsimony, we use the standard error of the regression on an unadjusted basis in Exhibit 7.

[8]We also examine the temporal patterns implied by this analysis. The pattern in the original data is generally preserved here. Venture capital had two distinct periods in our sample: a lower-volatility period over 1988-1994, and a higher-volatility period over 1995-2002. Leveraged buyouts exhibited the opposite characteristics: a higher-volatility period in the first years of the sample, and a lower-volatility period in the later years. Real estate exhibited more stability across the periods, although volatility did decline somewhat in the later period. To capture full cycles or complete histories requires examination of the entire period.

[9]We use vintage year returns to 1997 (in order to have at least seven years of realizations from each vintage year). The seven-year choice is based on an examination of where the vintage year returns stabilize with respect to inception year.

REFERENCES

Bollen, Kenneth. "An Alternative 2SLS Estimator for Latent Variable Models." *Psychometrika,* 61 (1996), pp. 109-121.

Cochrane, John H. "The Risk and Return of Venture Capital." Working paper, NBER, 2001.

Craft, Timothy P. "The Role of Private and Public Real Estate in Pension Plan Portfolios." *Journal of Real Estate Portfolio Management*, 7 (1) (2001), pp. 17-23.

Desouza, Clifford, and Suleyman Gokcan. "Allocation Methodologies and Customizing Hedge Fund Multi-Manager Multi-Strategy Products." *The Journal of Alternative Investments*, 2004.

Geltner, D. "Temporal Aggregation in Real Estate Return Indices." *Journal of the American Real Estate and Urban Economics Association*, 21 (2) (Summer 1993).

Georgiev, Georgi. "The Benefits of Real Estate Investing." Working paper, CISDM, 2002.

Getmansky, Mila, Andrew W. Lo, and Igor Makarov. "An Econometric Model of Serial Correlation and Illiquidity in Hedge Funds." Working paper, Massachusetts Institute of Technology, Sloan School of Business, 2003.

Giliberto, S.M. "Assessing Real Estate Volatility." *The Journal of Portfolio Management*, 29 (5) (2003).

Gompers, Paul A., and Joshua Lerner. "Money Chasing Deals? The Impact of Fund Inflows on Private Equity Valuations." *Journal of Financial Economics,* 55 (2000), pp. 281-325.

——. "Private Equity and Asset Allocation: Clues to a Puzzle." Working paper, Citigroup Alternative Investments, 2002.

——. "Risk and Reward in Private Equity Investments: The Challenge of Performance Assessment." *The Journal of Private Equity,* Winter 1997, pp. 5-12.

Greene, William H. *Econometric Analysis*. East Brunswick, NJ: Prentice-Hall, 1993.

Hadi, A.S. "Identifying Multiple Outliers in Multivariate Data." *Journal of the Royal Statistical Society*, B 54 (1992), pp. 761-771.

——. "A Modification of a Method for the Detection of Outliers in Multivariate Samples." *Journal of the Royal Statistical Society*, B 56 (1994), pp. 393-396.

Peng, Liang. "Building a Venture Capital Index." Working paper, Yale University Department of Economics, 2001.

Quigley, John M., and Susan E. Woodward. "An Index for Venture Capital." Working paper, University of California at Berkeley Department of Economics, 2003.

Staiger, Douglas, and James H. Stock. "Instrumental Variables Regression with Weak Instruments." *Econometrica*, 65(3) (1997), pp. 557-586.

To order reprints of this article, please contact Ajani Malik at amalik@iijournals.com or 212-224-3205.