Research Review

Investing in Private Equity

Andrew Ang, Ph.D
Columbia University
Morten Sorensen, Ph.D
Columbia University

1. Introduction

Private equity (PE) investments are investments in privately-held companies, that trade directly between investors instead of via organized exchanges. PE is often considered a distinct asset class, and it differs from investments in public equity in fundamental ways. There is no active market for PE positions, making these investments illiquid and difficult to value. PE funds typically have horizons of 10-13 years, during which the invested capital cannot be redeemed.

We survey the academic research concerning the risks and returns of institutional PE investments, as well as the optimal holdings of PE in an investment portfolio. Researchers have had limited access to information about the nature and performance of PE investments, and so research in this area is still preliminary and often inconclusive.

Institutional PE investments are typically made through a PE fund organized as a limited partnership, with the institutional investors (typically pension funds or university endowments) as limited partners (LPs) and the PE firm itself (such as Blackstone or KKR) acting as the general partner (GP). The GP manages the PE fund's acquisitions of individual companies (called portfolio companies). Depending on the type of portfolio companies, PE funds are typically classified as buyout, venture capital (VC), or some other type of fund specializing in other illiquid non-listed equity-like investments. Buyout funds invest in mature established companies, using substantial amounts of leverage to finance the transactions. VC funds invest in high-growth start-ups, using little or no leverage. Kaplan and Stromberg (2009) provide a detailed description of these investments and the development of the industry.

2. Estimating Private Equity Risk and Return

For traded assets, the risk and return are usually defined using the capital asset pricing model (CAPM) as the alpha and beta coefficients estimated in the one-factor linear regression (the expected return regression),

$$R_i(t) - R_f(t) = \alpha + \beta[R_m(t) - R_f(t)] + E_i$$

Here, $R_i(t)$ is the return earned by the investor from period t-1 to period t, $R_f(t)$ is the risk-free rate over this period, and $R_m(t)$ is the return on the market portfolio. The return earned on a financial asset from time t-1 to t is defined as:

$$R(t) = \frac{P(t) + CF(t)}{P(t-1)} - 1$$

where CF(t) is the cash flow paid out at time t and P(t) is the market price quoted at time t, immediately after payment of the cash flow.

The standard approach to evaluating the risks and returns of financial assets proceeds in two steps. First, alpha and beta are estimated by regressing the asset's observed returns on the corresponding market returns using the expected return regression. Second, invoking the CAPM, the estimated alpha is interpreted as an abnormal risk-adjusted return and the beta is interpreted as the systematic risk.

For PE investments, problems arise at both steps. At the first step, privately-held companies, by definition, do not have regularly observed market values, and the returns earned from investing in these companies are only observed at exit. Hence, period-by-period returns are unavailable, making it difficult to estimate the expected return regression directly. Better-performing privately-held companies may also be overrepresented in the data, creating sample selection problems that would cause the alpha coefficient to be overestimated and the beta coefficient to be underestimated. At the second step, after estimating alpha and beta, it is unclear whether or not these coefficients appropriately measure risks and returns. The assumptions of liquid and transparent markets underlying the CAPM are far from the realities of PE investing. To reflect the actual risks and returns facing LP investors, the estimated parameters may require adjustments to account for the cost of illiquidity, idiosyncratic risk, persistence, funding risk, etc. The lack of regularly quoted market prices and returns presents a fundamental challenge for empirical studies of the risk and return of PE investments. Alternative approaches have either used company-level performance data or fund-level data with the cash flow streams between the LPs and GPs.

2.1. Using Company-level Data

Company-level data contain information about investments by buyout or VC funds in individual companies. For these investments, the data typically contain the name of the company, the invested amount, the investment date, the exit date, and the exit amount. Such data are confidential and proprietary and researchers have mostly obtained them through direct contact with LPs and professional data providers. Studies using such data include Gompers and Lerner (1997), Cochrane (2005), and Korteweg and Sorensen (2010).

Compared to fund-level data, company-level data have two advantages. First, there are many more companies than funds, improving the statistical power of the analysis. Companies can be classified in terms of industries and types, allowing for a more nuanced differentiation of the risks and returns across industries and types and over time. Second, investments in individual companies have well-defined returns. The return can be calculated directly from the initial investment and the distribution of the proceeds at exit. As long as intermediate cash flows are few and small, as for buyout investments, this calculation provides a reasonable measure of the return.

A disadvantage of company-level data is that the return figures typically do not subtract management fees and carried interest paid by the LPs to the GPs. Hence, the estimated risks and returns reflect the total risks and returns of the investments (before fees), not those earned by an LP (net of fees). Translating between net-of-fee and before-fee returns typically requires additional assumptions and numerical simulations (see Metrick and Yasuda (2010) and Franzoni, Nowak, and Phalippou (2012) for two approaches).

2.1.1. Continuous-Time Specifications

A technical disadvantage of company-level data is that the returns are measured over different periods. Returns are measured from the time of the initial investment to the time of exit, and the duration varies substantially across investments. The standard (discrete-time) CAPM is a one-period model, where the period may be a day, a month, or a quarter. This model does not compound, however, and the returns must all be calculated over periods of the same length.

A standard solution is to use a continuous-time version of the CAPM, which compounds and allows calculating risks and returns of investments of different durations. In the continuous-time CAPM, the expected return regression is restated in log-returns (continuously-compounded returns) as:

$$\ln[1 + R_i(t)] - \ln\left[1 + R_f(t)\right] = \delta + \beta\left(\ln[1 + R_m(t)] - \ln\left[1 + R_f(t)\right]\right) + \epsilon_i$$

One limitation is that the estimated intercept cannot be interpreted as an abnormal return, as in the discrete-time CAPM. Under additional distributional assumptions, the abnormal returns can be calculated using the following adjustment.

$$\alpha = \delta + \frac{1}{2}\sigma^2$$

With this adjustment, a high volatility leads to a high alpha. Cochrane (2005) reports an annual volatility around 90%, resulting in an alpha of 32% per year. This alpha seems unreasonably high compared to studies using fund-level data, casting some doubts on the appropriateness of the distributional assumptions underlying this model.

2.1.2. Selection Bias

Another problem with company-level data is sample selection. Data sets with valuations of individual VC rounds are dominated by better-performing companies, and failing start-ups are usually not formally liquidated but are instead left as shell companies without economic value (zombies). Hence, when observing old companies without new financing rounds or exits, these companies may be alive and well or they may be zombies, in which case it is unclear when the write-off of the company's value should be recorded.

This selection problem is illustrated in Exhibit 1 (Korteweg and Sorensen (2010)). The universe of returns is illustrated by all the dots. The observed data, however, only contain the good returns above the x-axis (in black). Worse returns (shaded gray) are unobserved. A simple estimation of the expected return regression gives an estimate of alpha that is upward biased, an estimate of beta that is downward biased, and a total volatility that is too low. An analysis that does not correct for these biases will be too optimistic about the performance of these investments.

Heckman (1979) introduced corrections for such selection biases. Cochrane (2005) estimates a dynamic selection model using VC data and finds a large effect of selection. The selection correction reduces the intercept of the log-market model, denoted δ above, from 92% to -7.1%. Korteweg and Sorensen (2010) estimate an extended version of Cochrane's model. Without selection bias, the estimate of the intercept, δ, is -19% per year, but correction for selection reduces this estimate to -68% (note, these intercepts cannot be interpreted as returns).

Korteweg and Sorensen estimate separate alphas for the 1987-1993, 1994-2000, and 2001-2005 periods. Alphas in

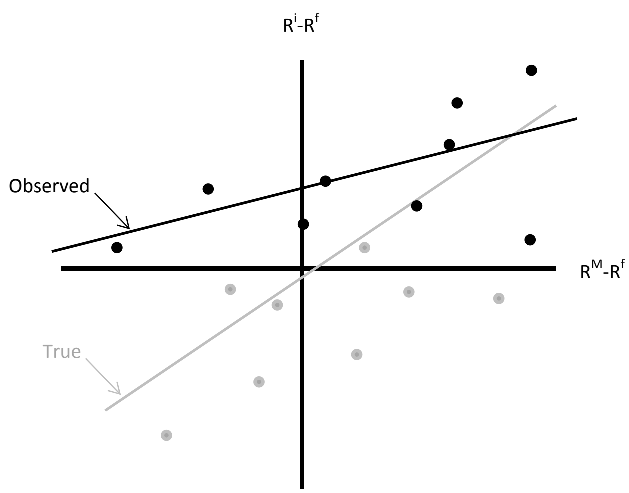

Exhibit 1 Illustration of Selection Bias
Source: Authors

the early period were positive but modest; alphas in the late 1990s were very large; and alphas in the 2000s were negative. Finally, they report beta estimates for VC investments of 2.6-2.8.

2.2. Using Fund-level Data

Fund-level data are typically obtained from LPs with investments across many PE funds. Each observation represents the performance of the entire portfolio of companies held by a fund. In addition to information about the fund, such as its investment focus and vintage year, the data may contain the cash flow stream between the LP and the fund or a performance measure calculated from this cash flow stream. When individual cash flows are available, however, they are typically not tied to individual portfolio companies.

There are several advantages to fund-level data. First, fund-level data reflect actual LP returns, net of fees, resulting in estimates of the risks and returns actually realized by the LPs. The sample selection problem is smaller, since the performance of companies that ultimately never produce any returns for the investing funds (zombies) is eventually reflected in the fund-level cash flows. However, other sample selection problems may arise. Fund-level performance is typically self-reported, and better-performing funds may be more likely to report their performance. Finally, since funds have similar lifetimes (typically 10 years), the expected return equation can be estimated directly, avoiding the problems with the continuous-time log-return specification used for company-level data.

2.2.1. Fund-level Performance Measures

A main problem for fund-level data is how to define a return. Period-by-period returns, as previously defined, require valuations of the PE investment at regular intervals, and the reported net asset values (NAVs) are noisy substitutes for these market values. Absent regularly quoted returns, alternative measures have been proposed. None of these measures is a return, as previously defined, however, and their relationships to asset pricing models are tenuous. For example, Phalippou (2009) offers a critical assessment of these measures.

2.2.1.1. Internal Rate of Return (IRR)

A natural starting point is to interpret the internal rate of return (IRR) of the cash flows between the LP and GP as a return earned over the life of the fund. Denoting the cash flow at time *t* as CF*(t)*, and separating those into capital calls paid by the LP to the GP, denoted C*all(t)*, and distributions of capital from the GP back to the LP, denoted D*ist(t)*, the IRR is defined as the solution to the equation,

$$PV = \sum \frac{CF(t)}{(1+IRR)^t} = \sum \frac{Dist(t) - Call(t)}{(1+IRR)^t} = 0$$

Thus,

$$\frac{\left(\Sigma \frac{Dist(t)}{(1+IRR)^t}\right)}{\left(\Sigma \frac{Call(t)}{(1+IRR)^t}\right)} = 1$$

A recent survey by Harris, Jenkinson, and Kaplan (2011) summarizes the academic studies using fund-level data from various data providers. For buyout funds, they report weighted average IRRs of 12.3%-16.9%. For VC funds, the weighted average IRRs are 11.7%-19.3%. The performance of buyout funds has been reasonably stable over time, with weighted average IRRs of 15.1%-22.0% in the 1980s, 11.8%-19.3% in the 1990s, and 5.8%-12.8% in the 2000s. VC fund performance has been more volatile, with weighted average IRRs ranging from 8.6% to 18.7% in the 1980s, 22.9% to 38.6% in the 1990s, and -4.9% to 1.6% in the 2000s.

Overall these figures reveal substantial variation in IRRs across studies and data sources. Moreover, the IRR is a problematic measure of economic performance. It is an absolute performance measure that does not calculate performance relative to a benchmark or market return. Moreover, the IRR calculation implicitly assumes that invested and returned capital can be reinvested at the IRR rate. If a fund makes an early small investment with a large quick return, the investment can largely define the IRR for the entire fund, regardless of the performance of subsequent investments.

2.2.1.2. Total Value to Paid-in Capital (TVPI) multiple

Another performance measure, which is less susceptible to manipulation than the IRR, is the total value to paid-in capital (TVPI) multiple. This multiple is calculated as the total amount of capital returned to the LP investors (net of fees) divided by the total amount invested (including fees). Formally, the TVPI multiple is defined as:

$$TVPI = \frac{\sum Dist(t)}{\sum Call(t)}$$

This calculation is performed without adjusting for the time value of money. While the IRR is calculated under the implicit assumption that capital can be reinvested at the IRR rate, the TVPI multiple is calculated under the implicit assumption that capital can be reinvested at a zero rate. Harris, Jenkinson, and Kaplan (2011) report weighted average TVPIs of 1.76-2.30 for buyout funds and 2.19-2.46 for VC funds. This multiple varies substantially over time, however. For buyout funds, they report a multiple of 2.72-4.05 for the 1980s, 1.61-2.07 for the 1990s, and 1.29-1.51 for the 2000s. For VC funds, they report a multiple of 2.31-2.58 for the 1980s, 3.13-3.38 for the 1990s, and 1.06-1.09 for the 2000s.

2.2.1.3. Public Market Equivalent (PME)

Both the IRR and TVPI measures are absolute performance measures. The public market equivalent measure (PME) is used to evaluate performance relative to the market. It is calculated as the ratio of the discounted value of the LP's inflows divided by the discounted value of outflows, with the discounting performed using realized market returns:

$$PME = \frac{\left(\sum \frac{Dist(t)}{\prod(1 + R_m(t))}\right)}{\left(\sum \frac{Call(t)}{\prod(1 + R_m(t))}\right)}$$

Kaplan and Schoar (2005) argue that a PME greater than one is equivalent to a positive economic return for the LPs when PE investments have the same risk as the general market (a beta equal to one). This report average equal-weighted PMEs of 0.96. Value-weighted, the PME for VC funds is 1.21 and the PME for buyout funds is 0.93.

More recently, Harris, Jenkinson, and Kaplan (2011) report weighted PMEs of 1.16-1.27 for buyout funds and 1.02-1.45 for VC funds. For buyout funds, the PMEs varied from 1.03-1.11 in the 1980s to 1.17-1.34 in the 1990s and 1.25-1.29 in the 2000s. For VC funds, the PMEs were 0.90-1.08 in the 1980s, 1.99-2.12 in the 1990s, and 0.84-0.95 in the 2000s.

2.3. Risk Measures

Fund-level data are poorly suited for estimating the risk of PE investing. Thus few, if any, academic studies attempt to use fund-level data. Instead, Ljungqvist and Richardson (2003) estimate the risk using publicly traded comparable companies, and report a beta of 1.08 for buyout investments and 1.12 for VC investments. These betas do not adjust for the leverage used in the buyout investments. Adjusting for systematic risk, they find a 5%-6% premium, which they interpret as the illiquidity premium of VC investments. Kaplan and Schoar (2005) regress IRRs on S&P 500 returns and find a coefficient of 1.23 for VC funds and 0.41 for buyout funds. A levered beta of 0.41 seems unreasonably low, however. More generally, Robinson and Sensoy (2011) argue that both investment and exit decisions correlate with the overall market, and that these correlations should be incorporated into a broader assessment of the overall risk of these investments.

2.4. Persistence and Predictability

Kaplan and Schoar (2005), Phalippou and Gottschalg (2009), Hochberg, Ljungqvist, and Vissing-Jorgensen (2010), along with other studies, find evidence of performance persistence for PE funds. The performance of an early fund predicts the performance of subsequent funds managed by the same GP. This persistence is interpreted as evidence that GPs vary in their skills and abilities to pick investments and manage the portfolio companies. Estimates suggest that a performance increase of 1.0% for a fund is associated with around 0.5% greater performance for the GP's next fund, measured either in terms of PME or IRR. For more distant funds, persistence declines.

3. Asset Allocations to Private Equity

Having discussed the measurement of PE returns, we now consider optimal allocations to PE. This requires, of course, a suitable risk-return trade-off for PE investments, as well as correlation of PE returns with other assets in the investor's opportunity set. As we point out in Section 2, measuring these inputs for PE for use in an optimization problem requires special considerations. We take as a given these inputs and focus on the illiquidity risk of PE and how to incorporate it into an optimal asset allocation framework. There have been several approaches to handling illiquidity risk in asset allocation, all of which have relevance. To put into context these contributions, we start with the case of asset allocation without frictions.

3.1. Frictionless Asset Allocation

The seminal contributions of Merton (1969, 1971) characterize the optimal asset allocation of an investor with constant relative risk aversion (CRRA) utility investing in a risk-free asset (with constant risk-free rate) and a set of risky assets. The constant relative risk aversion utility function with risk aversion *y* is given by:

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma}.$$

The constant relative risk aversion utility is homogeneous of degree one, which means that exactly the same portfolio weights arise whether $10 million of wealth is being managed or $1 billion. This makes the constant relative risk aversion utility function ideal for institutional asset management.

Assume the risky assets are jointly log-normally distributed. Under the case of independent and identically distributed *(i.i.d.)* returns, the vector of optimal holdings, w, of the risky assets is given by:

$$v = \frac{1}{\gamma}\Sigma^{-1}(\mu - r_f)$$

where Σ is the covariance matrix of the risky asset returns, μ is the vector of expected returns of the risky assets, and r_f is the risk-free rate. This is also the portfolio held by an investor with mean-variance utility optimizing over a discrete, one-period horizon.

There are two key features of this solution that bear further comment. First, the Merton model is dynamic and involves continuous rebalancing. That is, although the portfolio weights, w, are constant, the investor's optimal policy is always to continuously sell assets that have risen in value and to buy assets that have fallen in value in such a way as to maintain constant weights. Clearly, the discrete nature of PE investment and the inability to trade frequently mean that allocations to PE should not be evaluated with the standard Merton model.

Second, any other portfolio held by the investor other than the optimal portfolio results in lower utility. Thus, an investor holding a non-optimal portfolio needs to be compensated, as he can improve his utility by moving to the optimal portfolio. The cost of holding a non-optimal portfolio is called the utility certainty equivalent, and it is dependent on the investor's risk preferences and investment horizon. Formally, the certainty equivalent cost is how much an investor must be compensated in dollars per initial wealth to take a non-optimal strategy, but have the same utility as with the optimal strategy. Some particularly relevant costs, which the subsequent literature explores, are how much an investor should be compensated for the inability to trade assets like PE for certain periods of time or how much to be compensated for being forced to pay a cost whenever an asset is traded.

3.2. Asset Allocation with Transactions Costs

Investing in PE incurs large transactions costs in finding an appropriate PE manager and conducting appropriate due diligence. Then, there are potentially large discounts to the recorded asset values that may be taken in transferring ownership of a PE stake in illiquid secondary markets.

Constantinides (1986) considers the case of one risk-free and one risky asset. When there are proportional transaction costs, so that whenever the holdings of the risky asset increase (or decrease) by v, the holding of the riskless asset decreases by *(1+k)v*. When there are trading costs, the investor trades infrequently. Constantinides shows that the optimal trading strategy is to trade whenever the risky asset position hits upper and lower bounds, $\underline{w}$ and $\overline{w}$, respectively. These bounds straddle the optimal Merton model where there are no frictions. The holdings of risky to risk-free assets, *y/x*, satisfy:

$$\underline{w} \le \frac{y}{x} \le \overline{w}$$

so that when *y/x* lies within the interval $[\underline{w}, \overline{w}]$ there is no trade and when *y/x* hits the boundaries on either side, the investor buys and sells appropriate amounts of the risky asset to bring the portfolio back to the optimum Merton model.

The no-trade interval, $\overline{w} - \underline{w}$, increases with the transactions costs, *k*, and the volatility of the risky asset. The transactions costs to sell PE portfolios in secondary markets can be extremely steep. When the Harvard endowment tried to sell its PE investments in 2008, potential buyers were requiring discounts to book value of more than 50%. Even for transactions costs of 10%, Constantinides (1986) computes no-trade intervals greater than 0.25 around an optimal holding of 0.26 for a risky asset with a volatility of 35% per annum. Thus, PE investors should expect to rebalance PE holdings very infrequently.

The certainty equivalent cost to holding a risky asset with large transaction costs is small for modest transaction costs (approximately 0.2% for proportional transaction costs of 1%), but can be substantial for large transaction costs, which is the more relevant range for PE investments. For transaction costs of 15% or more, the required premium to bring the investor to the same level of utility as the frictionless Merton model is more than 5% per annum.

A major shortcoming of this analysis is that it assumes that trade in assets is always possible, albeit at a cost, which is not true for PE; over a short horizon, there may be no opportunity to find a buyer and even if a buyer is found, there is not enough time, relative to the investor's desired short horizon, to raise capital to go through legal and accounting procedures to transfer ownership. There may be no opportunity to trade, even if desired, at considerable discounts. This case is what the next literature considers.

3.3. Asset Allocation with Search Frictions

An important part of rebalancing a PE portfolio is searching for a counterparty for the transaction. Since Diamond (1982), search-based frictions have been modeled by Poisson arrival processes. Agents find counterparties with an intensity λ, and conditional on the arrival of the Poisson process, agents can trade and rebalance. This produces intervals where no rebalancing is possible for illiquid assets and the times when rebalancing are possible are stochastic. This notion of illiquidity is that there are times where it is not possible to trade, at any price, an illiquid asset. These particular types of stochastic rebalancing opportunities are attractive for modeling PE in another way: the exit in PE vehicles is often uncertain. Although a PE vehicle may have a stated horizon, say of 10 years, the return of cash from the underlying deals may cause large amounts of capital to be returned before the stated horizon, or in many cases the horizon is extended to maximize the profitability of the underlying investments (or to maximize the collection of fees by GPs).

A number of authors have used this search technology to consider the impact of illiquidity (search) frictions in various over-the-counter markets, such as Duffie, Garleanu, and Pedersen (2005, 2007). While these are important advances for showing the effect of illiquidity risk on asset prices, they are less useful for deriving asset allocation advice on optimal PE holdings. Duffie, Garleanu, and Pedersen (2005, 2007) consider only risk-neutral and CARA utility cases and restrict asset holdings to be 0 or 1. Garleanu (2009) and Lagos and Rocheteau (2009) allow for unrestricted portfolio choice, but Garleanu considers only CARA utility and Lagos and Rocheteau focus on showing the existence of equilibrium with search frictions rather than on any practical calibrations. Neither study considers asset allocation with both liquid and illiquid assets.

3.4. Asset Allocation with Stochastic Non-Traded Periods

Ang, Papanikolaou, and Westerfield (2011) [APW] solve an asset allocation problem with liquid securities, corresponding to traded equity markets that can be traded at any time, and illiquid securities, which can be interpreted as a PE portfolio. The investor has CRRA utility with an infinite horizon and can only trade the illiquid security when a liquidity event occurs, which is the arrival of a Poisson process with intensity λ. In this framework, the Merton model with continuous rebalancing is given by $\lambda \rightarrow \infty$. As λ decreases to zero, the opportunities to rebalance the illiquid asset become more and more infrequent. The mean time between rebalancing opportunities is $1/\lambda$. Thus, λ indexes a range of illiquidity outcomes.

The inability to trade for stochastic periods introduces a new source of risk that the investor cannot hedge. This illiquidity risk induces large effects on optimal allocation relative to the Merton model. APW show that illiquidity risk affects the mix of liquid and illiquid securities even when the liquid and illiquid returns are uncorrelated and the investor has log utility.

The most important result that APW derive is that the presence of illiquidity risk induces time-varying, endogenous risk aversion. The intuition is that there are two levels of wealth that are relevant for the investor: (1) total wealth, which is the same effect as the standard Merton problem where the risk is that if total wealth goes to zero, the agent cannot consume, and (2) liquid wealth. The agent can only consume liquid wealth. Thus, with illiquid and liquid assets, the investor also cares about the risk of liquid wealth going to zero. This can be interpreted as a solvency condition: an agent could be wealthy but if this wealth is tied up all in illiquid assets, the agent cannot consume. Although the CRRA agent has constant relative risk aversion, the effective risk aversion—the local curvature of how the agent trades off liquid and illiquid risk in the portfolio—is affected by the solvency ratio of the ratio of liquid to illiquid wealth. This solvency ratio also becomes a state variable that determines optimal asset allocation and consumption. This illiquidity risk causes the optimal holdings of even the liquid asset to be lower than the optimal holding of liquid assets in a pure Merton setting.

APW derive five findings that are important considerations for investing in PE:

1. Illiquidity risk induces marked reductions in the optimal holdings of assets compared to the Merton model. Their calculations for the same risk aversion as a 60% risky asset holding (and 40% risk-free holding) in the Merton model, introducing an average rebalancing period of once a year, reduces the risky asset holding to 37%. When the average rebalancing period is once every five years, the optimal allocation is just 11%. Thus, PE, which is highly illiquid, should be held in modest amounts.
2. In the presence of infrequent trading, the fraction of wealth held in the illiquid asset can vary substantially and is very right skewed. For example, if the optimal holding to illiquid assets is 0.2 when rebalancing can take place, then, the investor should expect the range of illiquid holdings to vary from 0.15 to 0.35 during non-rebalancing periods. Because of the skew, the average holdings to the illiquid asset will be higher than the optimal rebalancing point, at say 0.25. Thus, when an illiquid PE portfolio is rebalanced, the optimal rebalancing point is much lower than for an average holding.
3. The consumption policy (or payout policy) with illiquid assets must be lower than the Merton payout policy with only liquid assets. Intuitively, holding illiquid assets means that there is additional solvency risk that liquid wealth goes to zero and consumption cannot be funded. Thus, payouts of funds holding illiquid assets should be lower than the case when these assets are all fully traded.
4. The presence of illiquidity risk means that an investor will not fully take advantage of opportunities that might look like close to an arbitrage. For example, where correlations to the liquid and illiquid returns are nearly plus or minus one. Traditional mean-variance optimizers without constraints would produce weights close to plus or minus infinity in these two assets. This does not happen when one asset is illiquid because taking advantage of this apparent arbitrage involves a strategy that causes the investor's liquid wealth to drop to zero with positive probability. Thus, near-arbitrage conditions when there is illiquidity risk are not exploited like in the Merton model.
5. Finally, the certainty equivalent reward required for bearing illiquidity risk is large. They report that when the liquid and illiquid returns are poorly correlated and the illiquid portfolio can be rebalanced, on average, once every five years (which is a typical turnover of many PE portfolios), the liquidity premium is over 4%. For rebalancing once a year, on average, the illiquidity premium is approximately 1%. These numbers can be used as hurdle rates for investors considering investing in PE.

APW still miss a number of practical considerations that the future literature should address. The most important one is that in the Merton setting into which APW introduce illiquidity, there are no cash distributions; all risky asset returns (both liquid and illiquid) are capital gains. PE investments require cash flow management of capital calls and distributions. Some ad hoc simulations on this issue have been conducted by Siegel (2008), Leibowitz and Bova (2009), and De Zwart, Friesser, and van Dijk (2012). An extension of APW to incorporate cash flow streams could address this.

4. Conclusion

Our findings and recommendations for investments in PE may be summarized as follows:

- Empirical approaches commonly used to estimate the risk and return of standard publicly traded securities are difficult to apply. Complicating features of PE investments include the limited data, the irregular nature of such investments, and the sample selection problems that typically arise in reported PE data. Adjusting for these difficulties requires sophisticated econometric techniques. Without appropriate adjustment, naïve analysis tend to understate the risk and volatility and may exaggerate performance estimates.

  Recommendation: Reported estimates of PE risk and return should be interpreted with caution. Simple standard methodologies fail to take into consideration all the nuances that a thorough and accurate evaluation of a PE investment requires. Studies that develop methodologies to perform these adjustments are still in a preliminary phase, and a consensus on the appropriate adjustments has yet to emerge.

- Commonly used fund performance measures, such as the internal rate of return (IRR), the total value to paid-in (TVPI) multiple, and public market equivalent (PME), are problematic. There is substantial variation in the estimates of these measures across studies and data sources. The measures can, to some extent, be manipulated by the timing and magnitude of the individual investments. These fund performance measures use only rough risk adjustments too. Fundamentally, these measures are not derived from underlying financial theories of risk and return, making them difficult to interpret consistently.

  Recommendation: Commonly reported performance measures should be interpreted with caution. They are not return measures as commonly understood.

- Asset allocation models that account for transaction costs (which are high for PE) and illiquidity risk (which is substantial for PE) recommend modest holdings of PE. In these models, rebalancing will be infrequent, so wide swings in the holdings of PE can be expected. Also, the holdings of illiquid PE will be much lower than predicted by asset allocation models, assuming that all assets can be rebalanced when desired.

  Recommendation: When determining optimal PE allocations, asset allocation models must account for the inability to rebalance PE positions. Allocations to illiquid PE investments should generally be modest.

---

[*] We thank the editor Edward Szado. We thank Ludovic Phalippou and Fernando Zapatero for comments and Jeffrey Lu for excellent research assistance. Portions of this paper were presented at the 2011 CFA Institute Conference in Chicago. This article is a summary version of an original article first published in the *Quarterly Journal of Finance*, Volume 2, Issue 3, 2012, DOI: 10.1142/S2010269012500093.

Ang and Morten acknowledge funding from Netspar and the Program for Financial Studies.
Corresponding author: Morten Sorensen, email: ms3814@columbia.edu

[1] See "Liquidating Harvard" Columbia CaseWorks ID#100312, 2010.

References

Ang, A., D. Papanikolaou, and M. Westerfield. "Portfolio Choice with Illiquid Asset." Working paper, Columbia University, 2011.

Cochrane, J. "The Risk and Return of Venture Capital." *Journal of Financial Economics*, 2005, Vol. 75, No. 1, pp. 3-52.

Constantinides, G. M. "Capital Market Equilibrium with Transaction Costs." *Journal of Political Economy*, 1986, Vol. 94, No. 4, pp. 842-62.

De Zwart, G., B. Frieser, and D. van Dijk. "Private Equity Recommitment Strategies for Institutional Investors." *Financial Analysts Journal*, 2012, Vol. 68, No. 3, pp. 81-99.

Diamond, P. A. "Aggregate Demand in Search Equilibrium." *Journal of Political Economy*, Vol. 1992, Vol., 90, No. 5, pp. 891-894.

Duffie, D., N. Garleanu, and L. H. Pedersen. "Over-the-Counter Markets." *Econometrica*, 2005, Vol. 73, No. 6, pp. 1815-1847.

Duffie, D., N. Garleanu, and L. H. Pedersen. "Valuation in Over-the-Counter Markets." *Review of Financial Studies*, 2007, Vol. 20, No. 6, pp. 1865-1900.

Franzoni, F., E. Nowak, and L. Phalippou. "Private Equity Performance and Liquidity Risk." *Journal of Finance*, 2012, Vol. 67, No. 6, pp. 2341–2373.

Garleanu, N. "Pricing and Portfolio Choice in Illiquid Markets." Journal of Economic Theory, 2009, Vol. 144, No. 6, pp. 532-564.

Gompers, P. and J. Lerner. "Risk and Reward in Private Equity Investments: The Challenge of Performance Assessment." *Journal of Private Equity*, 1997, Vol. 1, No. 2, pp. 5-12.

Harris, R., T. Jenkinson, and S. N. Kaplan. "Private Equity Performance: What Do We Know?" Working paper, University of Chicago, 2011.

Heckman, J. "Sample Selection Bias as a Specification Error." *Econometrica*, 1979, Vol. 47, No. 1, pp. 153-62.

Hochberg, Y., A. Ljungvist, and A. Vissing-Jorgensen. "Informational Hold-Up and Performance Persistence in Venture Capital." Working paper, Northwestern University, 2010.

Kaplan, S. N. and Antoinette Schoar. "Private Equity Performance: Returns, Persistence, and Capital Flows." *Journal of Finance*, 2005, Vol. 60, No. 4, pp. 1791-1823.

Kaplan, S. N. and P. Stromberg. "Leverage Buyouts and Private Equity." *Journal of Economic Perspectives*, 2009, Vol. 23, No. 1, pp. 121-46.

Korteweg, A. and M. Sorensen. "Risk and Return Characteristics of Venture Capital-Backed Entrepreneurial Companies." *Review of Financial Studies*, 2010, Vol. 23, pp. 3738-3772.

Lagos, R. and G. Rocheteau. "Liquidity in Asset Markets with Search Frictions." *Econometrica*, 2009, Vol. 77, No. 2, pp. 403-426.

Leibowitz, M. and A. Bova. "Portfolio Liquidity." Morgan Stanley Research, 2009.

Ljungqvist, A. and M. Richardson. "The Cash Flow, Return and Risk Characteristics of Private Equity." Working paper, NYU, 2003.

Merton, R. C. "Lifetime Portfolio Selection under Uncertainty: The Continuous Time Model." *Review of Economics and Statistics*, 1969, Vol. 51, No. 3, pp. 247-257.

Merton, R. C. "Optimum Consumption and Portfolio Rules in a Continuous Time Model." *Journal of Economic Theory*, 1971, Vol. 3, No. 4, pp. 373-413.

Metrick, A. and A. Yasuda. "The Economics of Private Equity Funds." *Review of Financial Studies*, 2010, Vol. 23, No. 6, pp. 2303-2341.

Phalippou, L. "Beware When Venturing into Private Equity." *Journal of Economic Perspectives*, 2009, Vol. 23, No. 1, pp. 147-166.

Phalippou, L., and O. Gottschalg. "The Performance of Private Equity Funds." *Review of Financial Studies*, 2009, Vol. 22, No. 4, pp. 1747-1776.

Robinson, D. T. and B. A. Sensoy. "Cyclicality, Performance Measurement, and Cash Flow Liquidity in Private Equity." Working paper, Duke University, 2011.

Siegel, L. B. "Alternatives and Liquidity: Will Spending and Capital Calls Eat Your "Modern" Portfolio." *Journal of Portfolio Management*, 2008, Vol. 35, No. 1, pp. 103-114.

Author Bios

Andrew Ang is the Ann F. Kaplan Professor of Business at Columbia Business School and Chair of the Finance and Economics Division. He is a financial economist whose work centers on understanding the nature of risk and return in asset prices. His work spans bond markets, equities, asset management and portfolio allocation, and alternative investments. Ang is the recipient of several grants from government and major industry organizations, and has served as associate editor for several leading journals. Ang has consulted for several financial institutions and currently serves as advisor to Martingale Asset Management and the Norwegian sovereign wealth fund.

Morten Sorenson is the Daniel W. Stanton Associate Professor of Finance and Economics at Columbia Business School where he is also the academic director of the Private Equity Center. He is a research fellow at the National Bureau of Economic Research (NBER) and the Stockholm Institute of Financial Research (SIFR). Before joining Columbia Business School, he was the William S. Fishman Scholar at the University of Chicago Graduate School of Business (now Booth School of Business). He received a PhD in economics from Stanford University in 2005, where he was a Kapnick Fellow at the Stanford Institute for Economic Policy Research (SIEPR).

His teaching and research is focusing on entrepreneurial finance, venture capital, and private equity and understanding how these alternative investments are different from regular investments in traditional financial securities. In particular, he has studied the impact of private equity and venture capital investors on the companies they finance and the industries in which they operate, the determinants of their investment decisions, the importance of portfolio company management for investment outcomes, the risk and return of PE and VC investments, and the impact of private equity and venture capital on the broader economy.