EXPLAINING BUYOUT INDUSTRY RETURNS: NEW EVIDENCE

This version: February 12, 2018

Forthcoming in the Journal of Investment Management

David Turkington, CFA
State Street Associates
140 Mt Auburn St, Cambridge MA 02138
dturkington@statestreet.com

Abstract

Traditional equity factors such as the leveraged equity risk premium, the small-cap premium and the value premium have had high historical returns on average, as has the buyout fund industry in aggregate. Previous research has argued that these factors explain the excess performance of private equity. However, time series regression analysis reveals that these factors explain surprisingly little variation in buyout performance. In contrast, other factors such as the credit premium and dynamic sector selection are more effective at explaining variation in performance of the buyout industry over time.

EXPLAINING BUYOUT INDUSTRY RETURNS: NEW EVIDENCE

The strong historical returns of the private equity buyout industry have generated a substantial amount of interest in explaining its performance. Whether or not private equity investments outperform comparable public market investments is a matter of much debate and controversy. Even after one obtains broad and reliable data for private investments, there are at least two important challenges that make it difficult to compare public and private equity performance. The first challenge is to appropriately compare public market investments which trade continually, to private investments which consist of highly irregular cash flows. The second challenge is to identify a benchmark of publicly traded securities that properly matches the risks of private equity.

Academics and practitioners have widely coalesced around use of the Public Market Equivalent (PME) to solve the first challenge. Originally introduced by Kaplan and Schoar (2005), the PME essentially overlays public market returns onto cash flow patterns that match each private investment inflow and outflow. For a given choice of public market benchmark, such as the S&P 500, PME analysis provides a proper assessment of the relative performance of private equity. Unfortunately, due to the second challenge it is far from clear which benchmark is appropriate to use in the PME calculation, and conclusions about the value added – or lack thereof – associated with private equity hinges on this choice. For example, Harris et al (2016) find that buyout funds outperform the S&P 500 by 3 to 4 percent per year based on PME, but they fail to outperform the Russell 2000 small-cap value index by any significant margin over the same period. L'Her et al (2016) show that the PME-based performance advantage of buyout funds disappears after adding small-cap, leverage and sector tilts to the benchmark. Other papers have also shown that portfolios of stocks with various combinations of exposure to leverage, small-cap and value factors have experienced growth rates above the S&P 500 and similar to that of buyout funds (see for example, Phalippou (2012), Chingono and Rasmussen (2015) and Stafford (2015)). However, we must exercise caution in interpreting these results. While PME is useful to compare long-run returns of private and public investments, it does not reveal any meaningful information about the correlation between their returns. Thus, the fact that private equity's excess return decreases when leverage, small-cap and value are added to the benchmark proves only that these factors also outperform over long periods; it does not imply that these factors actually "explain" a portion of private equity return dynamics. In fact, I find that these particular factors explain little to no variation in private equity performance over time.

PME analysis measures the long-run performance multiple of private equity relative to a chosen public equity benchmark, controlling for the irregular cash flows of private equity. However, as I will illustrate later on, PMEs do not measure the degree to which various factors co-vary with performance over time. A completely uncorrelated random variable can reduce the PME of private equity simply because both have positive cumulative returns,[1] but in this scenario we must not conclude that the variables truly explain one another. Instead, explanatory relationships are reflected in the correlations between variables over time.

In this article, I apply standard regression analysis to time series of aggregate buyout returns. I argue that the factors that best explain buyout performance are those that co-move with it over time. Contrary to conventional wisdom, I find that the aggregate buyout industry has a beta of less than 1 to the overall public equity market, and a slight negative exposure to the small-cap premium. Regression results also suggest a modest positive relationship to the value premium, and a significant explanatory relationship with the credit premium and dynamic sector selection – two variables that are also intuitive but typically omitted from analysis of buyout funds.

Previous research applying time series analysis to private equity data has revealed interesting and sometimes counterintuitive results. However, there have been relatively few such studies, and they often rely on complex models applied to vast quantities of individual cash flows. This degree of complexity creates a challenge in reconciling divergent results from time series based analysis of cash flows versus PME-based analyses of cash flows. Ang et al (2013) use a Bayesian Markov Chain Monte Carlo procedure to fit a latent variable factor model based on observable and episodic cash flow data. Their results for buyout funds yield a beta of around 1.4 to the market risk premium, close to zero for the size premium, and around 0.6 for the value premium. Ang et al also provide a summary of the Fama-French 3 factor time series betas from 5 previous studies of buyout funds. The estimates vary dramatically, ranging from 0.8 to 1.7 for market beta, from –0.9 to 0.7 for the small-cap premium beta, and from 0.2 to 1.4 for the value premium beta. Despite the merits of studies that challenge some aspects of conventional wisdom, many academics and practitioners continue to promote the view that levered small-cap value stocks offer the most appropriate benchmark for private equity. These beliefs may be motivated in part by studies that have looked at the returns of publicly traded stocks with some degree of intuitive linkage to actual private equity investments. For example, Jegadeesh et al (2015) analyze the performance of publicly traded funds-of-funds that invest in private equity as well as publicly traded firms of general partners who manage private equity funds. One of their most striking results is a strong and significant positive regression coefficient on the small-cap factor. L'Her et al (2017) estimate the risk factor exposures of publicly traded companies after the buyout backed initial public offerings for those companies. Their regression results also highlight large positive and significant coefficients on the small-cap factor. These small-cap effects contrast with my findings in this paper and with the findings of some other prior literature. It is plausible that other differences between public and private investments may explain such divergent results. For instance, the stock prices of listed private equity firms may reflect information about the business model of the general partners and their ability to generate fee revenue, in addition to information about the underlying investments. It is also possible that the variables that determine entry and exit prices for a privately held company differ from the price the public market demands for bearing exposure to that same company. I will leave these questions open to conjecture and future research. Ultimately, my analysis in this article focuses on measuring the performance of private investments directly as viewed through the experience of investors who hold them.

The overall state of disagreement about the risk factors of buyout funds motivates a need for further analysis and a clear understanding of time series risk factor analysis. My contribution in this paper is threefold. First, I analyze private market returns using time series approaches that are standard in public markets. I apply multiple methods to correct in a transparent way for the challenges inherent in private market data. Second, I include credit and sector rotation factors which are rarely considered in the literature, and which regression reveals to be important. Third, I offer intuition and numerical examples to explain the apparent contradiction between the results of PME-based analysis and time series analysis with regard to the factors that explain private equity returns.

Data and Methodology

Data on the Performance of US Buyout Funds

I analyze the performance of US buyout funds within the State Street GX Private Equity Index (GXPEI) data set. The GXPEI data set derives from book-of-record cash flows and valuations obtained via State Street's role as custodian for more than $2.7 trillion of capital commitments and more than 2,800 unique partnerships as of Q3 2017, comprising in total more than half of all global private equity assets. In addition to its wide coverage, this data set is also less susceptible to reporting bias compared to data sets that rely instead on voluntary self-reporting by funds. The GXPEI contains data as early as the 1980s, but the data are most representative and reliable for our purposes since the early 2000s, when a greater number of funds are present. All analysis in this article spans Q3 2002 through Q2 2016 and pertains exclusively to funds focused on US buyout investments. This sample represents approximately half of the net asset value of the broader GXPEI universe across all types of funds and geographies.

Analytical Challenges Posed by Illiquid Assets

There are multiple features of private equity which make performance analysis more complicated than for publicly traded securities. The first is cash flow timing. Cash contributions and subsequent distributions to limited partners occur at the general partners' discretion. They are irregular in their size and timing throughout the lifecycle of each fund. Researchers most commonly address this issue by computing dollar-weighted performance metrics rather that time-weighted performance metrics. For example, the Internal Rate of Return (IRR) for a private equity investment is equal to the discount rate that would render the net present value of all cash inflows and outflows equal to zero. Analysts typically compare private and public investment performance by applying the same historical cash flows to a public market investment and observing the performance multiple, which is called the Public Market Equivalent (PME). Both the IRR and PME serve to summarize average performance over a long sample of data. A PME of 1, for example, indicates that a private investment and a chosen public market benchmark deliver the same cumulative cash-flow-weighted growth, but it does not reveal whether periods of above-average and below-average growth for the two investments align with the same time periods within the sample (see Appendix A for an illustrative example). Therefore, PMEs are not amenable to time series regression analysis. To facilitate time series analysis, I measure instead the quarterly returns of the buyout industry in aggregate, which we can compare directly to the quarterly returns of public market investments. I compute quarterly IRRs which account for all pooled cash flows as well as private asset valuations at the beginning and end of each quarter. In practice, most buyout investments are not bought or sold during a given quarter, so the quarterly IRRs mostly reflect changes in company valuations over the period. It is worth noting that the components of return which correspond to investments that have no interim cash flows are no different from simple time-weighted returns. Indeed, the time series of quarterly IRRs for the pooled buyout universe is very similar to a time series of quarterly valuation changes for the investments which are not transacted. This time series represents full investment in the full range of US buyout funds available each quarter.

The quarterly valuations present a second important challenge, though. As is commonly known, the appraisal-based company valuations reported by private equity funds generally lag true economic values. Stale valuations artificially smooth returns and obscure the true relationships with other variables. As a result, volatility and correlation estimates – and by extension regression betas – based on quarterly returns are usually misleading. When performance is measured over longer periods, however, there is a much higher likelihood that private market valuations will reflect economic outcomes over that period. It also makes intuitive sense to evaluate performance over longer horizons because the private equity asset class is inherently long term in nature, as most funds require capital commitments for 10 years and do not allow for frequent entry and exit. Prior research including and Kinlaw et al (2013) "desmooth" a quarterly time series of private equity by reverse engineering its autocorrelations. This approach renders more realistic estimates of private equity volatility, but it will not necessarily capture time-lagged correlations with other market variables. Thus, in this analysis I analyze returns over annual intervals. Long-horizon correlations implicitly capture all of the serial dependence present in the higher-frequency quarterly time series.[2] Annual return intervals also mitigate seasonality effects which may be present in returns, such as the greater tendency of buyout valuations to fall in the fourth quarter as documented by Czasonis et al (2017). The main disadvantage of using longer-horizon returns is that there are fewer independent annual data points than there are independent quarterly data points. Nevertheless, annual returns can still support meaningful conclusions.

Analyzing pooled quarterly IRRs as a time series also avoids some important challenges associated with regressions performed at the level of individual buyout deals. As an example, Axelrod et al (2014) note that in the data sample they analyzed, 479 out of 2,075 buyout deals had returns of -100%. The authors apply a jump-CAPM model to deal with the substantial non-normality of deal level returns, and they show that failing to account for these features of the data may lead to mismeasurement problems.

Public Market Factors

I consider the following five public market factors, each of which has a plausible connection to private equity performance. This list is not necessarily exhaustive, but I believe it is a reasonable starting point based on results and hypotheses from previous literature.

Equity risk premium. It is natural to compare the performance of private equity investments to public equity investments. After all, companies are often public before they are brought private in a leveraged buyout, and many become public via IPOs after they are privately held. I measure the equity risk premium using the S&P 500 total return index.

Small cap premium. Many companies held within buyout funds have market capitalizations below those of the firms in the S&P 500. Most companies are perhaps better classified as small cap firms (see, for example, L'Her et al. (2016)). It is possible that the small cap premium explains a portion of buyout performance. I measure the small cap premium as the total return of the S&P 600 small cap index minus the S&P 500.

Value premium. Research has also documented that buyout managers tend to favor value companies rather than growth companies (see, for example, Chingono and Rasmussen (2015) and Phalippou (2012)). It is possible, and has been argued often, that the value premium explains a portion of buyout performance. I measure the value premium as the total return of the S&P 500 value index minus the S&P 500 growth index.

Credit premium. Given that debt provides a substantial portion of the funding for leveraged buyout deals, it seems plausible that changes in credit conditions would impact the economic value of buyout funds through the liabilities side of the balance sheet. Funds issue debt at or around corporate yields, which creates an implicit short exposure to the credit premium: returns of corporate debt in excess of treasury bonds. In other words, after issuing debt at a particular yield, funds might expect to suffer economic losses if credit spreads tighten and benefit if credit spreads widen. Kaplan and Stromberg (2008) describe a theory for how debt market pricing may offer a source of arbitrage-based return to buyout managers when debt and equity markets are mispriced. I measure the credit premium as the total return of the Barclays U.S. credit index minus that of the Barclays U.S. treasuries index.

Sector rotation. In recent research, Kinlaw et al (2015) find that dynamic sector selection, inferred via rolling regression analysis of private equity returns on sectors, produces a positive excess return above a large cap public benchmark. They also show that the returns attributable to sector selection are distinct from other factor premiums. This factor has been introduced into the literature more recently than the others I consider, and it requires additional analysis to model because the sector exposures of buyout companies vary through time. I compute a private equity sector rotation strategy in a manner very similar to Kinlaw et al. The authors suggest that economic sector exposures inferred through regressions are more representative than reported sectors, because many companies have exposure to more than one economic sector. To construct an out-of-sample sector rotation strategy, I regress quarterly returns for the GXPEI US Buyout index on contemporaneous and one-quarter lagged returns of the 11 GICS sectors, as well as contemporaneous and lagged returns of the S&P 500 to control for overall market risk, and one-quarter lagged returns of the buyout index to account for the most substantial short-term effects of lagged valuations. Each month, based on the past 5 years of quarterly returns, I use a LASSO regression to identify the ten most significant sector coefficients (potentially including their lags), while shrinking the rest of the coefficients to zero. LASSO stands for Least Absolute Shrinkage and Selection Operator (introduced by Tibshirani (1996)). It is an Ordinary Least Squares regression with an added penalty on the total sum of the coefficients' absolute values. Therefore, it is effective at identifying the most statistically significant coefficients while shrinking the others to zero. It is useful as a "selection" model in which we do not necessarily know in advance which variables (in this case, sectors) are most important in explaining the dependent variable. Kinlaw et al use stepwise regression, which is a related method for selecting a subset of coefficients based on statistical significance. My results for the private equity tracking portfolio based on rolling regression are qualitatively similar to theirs. I sum the resulting contemporaneous and lagged coefficients for each sector and form a long-only portfolio of these sectors using mean-variance optimization with a maximum allowable weight of 25 percent per sector. The optimization minimizes portfolio volatility, based on a covariance matrix of monthly returns over the prior 5 years, subject to a constraint that the weighted average sector coefficients from the private equity regression are greater than or equal to the largest positive sector beta divided by 2. Thus, the optimizer will favor the sectors most correlated to private equity, but may choose some sectors that are slightly less correlated if they diversify portfolio risk effectively. The sector rotation factor equals the out-of-sample quarterly returns of the sector portfolios minus the S&P 500.

Regression Analysis

Due to the valuation smoothing of private equity, it is important to analyze performance over periods longer than one quarter, even though doing so is somewhat inconvenient from a statistical perspective. I present results for three different approaches to this challenge, as summarized below. Each method has benefits and drawbacks. However, they all point to conclusions that are meaningfully different from those of a baseline quarterly returns regression, and also different from some commonly held beliefs about private equity.

<u>Method #1: Quarterly returns with lags</u>. In addition to the contemporaneous quarterly returns, include lags of public market variables to allow for the possibility that private equity returns may reflect these factors after some time delay. This approach has the benefit of allowing for standard statistical inference regarding significance tests. Its disadvantages include the possibility that the regression will be over-fit and subject to noise between each lag of public market variables as well as the inability to easily evaluate the joint significance of each variable while accounting for all of its lags.

<u>Method #2: Non-overlapping annual returns</u>. Condense returns into distinct calendar-year cumulative returns. This approach has the benefit of measuring annual-horizon relationships

while still allowing for standard statistical inference regarding significance tests. Its disadvantages include the fact that the number of data points decreases substantially, and the results could differ if annual returns are aligned to the end of Q1, Q2 or Q3, rather than calendar year end (Q4).

Method #3: Overlapping annual returns. Perform a regression using all overlapping annual interval returns. This approach has the benefit of capturing annual horizon relationships while still reflecting variation that results from different quarterly starting points. The use of overlapping intervals does not impose any bias on coefficient estimates. The primary disadvantage of this approach is that statistical inference regarding significance tests must be performed with caution and based on standard errors that are adjusted for the impact of highly correlated regression residuals.

Baseline Quarterly Regression

Exhibit 1 shows the results of a regression of quarterly buyout returns versus contemporaneous quarterly returns for the five public market factors. This regression shows, rather implausibly, that the beta of private equity to the overall public equity market is only 0.46. It shows a positive relationship with the small-cap premium of 0.20, which is not statistically significant. None of the other factors are statistically significant in this quarterly regression, and the value premium has a negative coefficient. In sum, I argue that this quarterly regression is misleading because it ignores meaningful relationships that occur over longer intervals.

Exhibit 1: Baseline Quarterly Regression Results

QUARTERLY

	Intercept	Public Equity	Small Cap Premium	Value Premium	Credit Premium	Sector Premium
Coefficient	1.98%	0.46	0.20	-0.10	-0.09	0.04
t-Statistic	3.94	4.67	1.34	-0.59	-0.34	0.25
R-squared	0.56		Number of observations		56	
Adj. R-squared	0.52					

Notes: Buyout returns come from the State Street Global Exchange Private Equity Index (GXPEI). All public market factors are derived from total return indices. Public equity is S&P 500, the small-cap premium is S&P 600 minus S&P 500, the value premium is S&P 500 value minus S&P 500 growth, the credit premium is Barclays government bonds minus Barclays corporate bonds, and the sector premium is a dynamic portfolio of S&P 500 sectors created on an out-of-sample rolling basis to replicate the implicit sector exposures of the buyout returns (as described in detail in the methodology section) minus the S&P 500. All data are sourced from Datastream, except for the GXPEI which is sourced from State Street Corporation.

Quarterly Regression with Lags

Exhibit 2 shows the results of a quarterly regression which is extended to include 1-quarter, 2-quarter and 3-quarter lags of each of the five public market variables, in addition to their contemporaneous quarterly returns. This regression reveals that many of the lagged coefficients are large in size and statistically significant, even with the relatively large number of variables included which reduces the degrees of freedom and hence increases the standard errors of the coefficient estimates. Interestingly, the small-cap premium has a negative lagged impact, whereas the value premium has a positive lagged impact, which is opposite to the findings of the contemporaneous-only quarterly regression. Exhibit 2 also reports the sum of all coefficients corresponding to each factor; a method which Asness et al. (2001) used in a similar context to analyze lagged S&P 500 exposures of hedge fund returns. The sum of betas shows a public equity exposure of 0.80, value premium exposure of 0.56, and sector premium exposure of 0.73.

Exhibit 2: Quarterly Regression Results with Lags

QUARTERLY: WITH LAGS

	Intercept	Public Equity	Small Cap Premium	Value Premium	Credit Premium	Sector Premium
Contemporaneous quarterly returns						
Coefficient	1.14%	0.39	0.12	0.02	0.23	0.02
t-Statistic	1.55	3.70	0.76	0.13	0.79	0.12
1 lag quarterly returns						
Coefficient		-0.02	-0.06	0.34	0.44	0.12
t-Statistic		-0.15	-0.43	1.96	1.38	0.73
2 lag quarterly returns						
Coefficient		0.27	-0.04	0.08	-0.62	0.39
t-Statistic		2.77	-0.32	0.43	-2.30	2.19
3 lag quarterly returns						
Coefficient		0.15	-0.04	0.11	0.02	0.20
t-Statistic		1.41	-0.32	0.53	0.08	1.24
Sum of contemporaneous and all lagged coefficients						
		0.80	**-0.03**	**0.56**	**0.07**	**0.73**
R-squared	0.78		Number of observations		56	
Adj. R-squared	0.65					

Notes: Buyout returns come from the State Street Global Exchange Private Equity Index (GXPEI). All public market factors are derived from total return indices. Public equity is S&P 500, the small-cap premium is S&P 600 minus S&P 500, the value premium is S&P 500 value minus S&P 500 growth, the credit premium is Barclays government bonds minus Barclays corporate bonds, and the sector premium is a dynamic portfolio of S&P 500 sectors created on an out-of-sample rolling basis to replicate the implicit sector exposures of the buyout returns (as described in detail in the methodology section) minus the S&P 500. All data are sourced from Datastream, except for the GXPEI which is sourced from State Street Corporation.

Annual Regression with Non-overlapping Calendar Years

Exhibit 3 shows the results of a regression in which the 56 available quarterly data points are aggregated into returns corresponding to 13 distinct calendar years (2003 through 2015). These results show relationships that are generally consistent with the summed betas from the quarterly lagged regression, but the annual results are easier to interpret because the betas themselves are more meaningful and they are associated with straightforward t-statistics. Despite a low number of data points, which reduces degrees of freedom and increases standard errors of the coefficients, both the public equity and sector premium coefficients are statistically significant. The value premium and credit premium are both intuitive in their signs, but with low statistical significance. The small-cap premium is slightly negative and close to zero.

Exhibit 3: Annual Non-overlapping Regression Results

ANNUAL: NON-OVERLAPPING CALENDAR YEARS

	Intercept	Public Equity	Small Cap Premium	Value Premium	Credit Premium	Sector Premium
Coefficient	5.37%	0.84	-0.12	0.38	-0.34	0.83
t-Statistic	2.33	5.25	-0.47	1.28	-0.98	3.09
R-squared	0.92		Number of observations		13	
Adj. R-squared	0.86					

Notes: Buyout returns come from the State Street Global Exchange Private Equity Index (GXPEI). All public market factors are derived from total return indices. Public equity is S&P 500, the small-cap premium is S&P 600 minus S&P 500, the value premium is S&P 500 value minus S&P 500 growth, the credit premium is Barclays government bonds minus Barclays corporate bonds, and the sector premium is a dynamic portfolio of S&P 500 sectors created on an out-of-sample rolling basis to replicate the implicit sector exposures of the buyout returns (as described in detail in the methodology section) minus the S&P 500. All data are sourced from Datastream, except for the GXPEI which is sourced from State Street Corporation.

Annual Regression with Overlapping 4-Quarter Returns

Exhibit 4 shows the results of a regression in which the dependent variable consists of every overlapping 4-quarter annual cumulative return of the buyout index, and the independent variables consist of the corresponding 4-quarter annual cumulative returns of the public market variables. The coefficient estimates are qualitatively very similar to those of the non-overlapping annual regression. It is important to note that the presence of overlapping windows does not impose any bias on coefficient estimates. Standard errors computed in a conventional manner will indeed be biased, however. I report Newey-West adjusted t-statistics which correct for the bias in standard errors (see Appendix B for details).

Exhibit 4: Annual Overlapping Regression Results

ANNUAL: ALL OVERLAPPING 4-QUARTER INTERVALS

	Intercept	Public Equity	Small Cap Premium	Value Premium	Credit Premium	Sector Premium
Coefficient	5.06%	0.91	-0.10	0.25	-0.75	0.70
t-Statistic (adjusted)	2.38	6.85	-0.51	0.99	-3.27	4.02
R-squared	0.80		Number of observations		53	
Adj. R-squared	0.78					

Notes: Buyout returns come from the State Street Global Exchange Private Equity Index (GXPEI). All public market factors are derived from total return indices. Public equity is S&P 500, the small-cap premium is S&P 600 minus S&P 500, the value premium is S&P 500 value minus S&P 500 growth, the credit premium is Barclays government bonds minus Barclays corporate bonds, and the sector premium is a dynamic portfolio of S&P 500 sectors created on an out-of-sample rolling basis to replicate the implicit sector exposures of the buyout returns (as described in detail in the methodology section) minus the S&P 500. All data are sourced from Datastream, except for the GXPEI which is sourced from State Street Corporation.

Statistical inference with overlapping data points can be complicated, and while the Newey-West adjustment approach is fairly standard in these scenarios, it may not be perfect. In Exhibit 5, I present a range of standard errors and associated t-statistics using a range of simple methods. The unadjusted standard errors shown in the first panel of Exhibit 5 are clearly flawed because they are derived from a traditional OLS methodology which assumes uncorrelated residuals. The bottom panel of Exhibit 5 shows the standard errors from the non-overlapping annual regression (originally presented in Exhibit 3). These standard errors are conservative in that they assume we have a sample of 13 non-overlapping data points. In this scenario, we can interpret the regression with overlapping data as providing an average of the coefficients with annual horizons that span not only Q1 through Q1, but all other quarterly starting points. Because the coefficient estimates are unbiased, we can think of this test as a slight refinement to the non-overlapping case but with roughly equal statistical power. Exhibit 5 also shows standard errors that are adjusted up to account for serial correlation. As explained in Appendix B, this adjustment factor is based on an assessment of the autocorrelation of the residuals and explanatory variables. The adjustment raises the standard errors by varying degrees for each independent variable based on its characteristics in the sample. The adjusted standard errors are nonetheless still lower than those of the non-overlapping annual regression. Intuitively, this occurs because the information contained in the additional overlapping data is correlated – though not completely redundant – to the original annual observations. The addition of intra-year information offers marginally more reliable evidence from a statistical perspective. Ultimately, it is important to acknowledge that statistical inference for overlapping data is a complex topic, so these results can be considered alongside the other regression methods I present.

Exhibit 5: Comparison of Standard Error (SE) and t-Statistic Calculation

COMPARISON OF STANDARD ERROR (SE) and T-STATISTIC CALCULATION

	Intercept	Public Equity	Small Cap Premium	Value Premium	Credit Premium	Sector Premium
Coefficient	5.06%	0.91	-0.10	0.25	-0.75	0.70
Unadjusted SE	1.39%	0.0957	0.1506	0.1797	0.2183	0.1648
t-Statistic	3.66	9.46	-0.65	1.39	-3.45	4.23
Adjusted SE	2.12%	0.1323	0.1905	0.2525	0.2303	0.1736
t-Statistic	2.38	6.85	-0.51	0.99	-3.27	4.02
SE from non-overlapping annual	2.31%	0.1607	0.2616	0.2938	0.3420	0.2700
t-Statistic	2.19	5.63	-0.37	0.85	-2.20	2.58

Notes: Buyout returns come from the State Street Global Exchange Private Equity Index (GXPEI). All public market factors are derived from total return indices. Public equity is S&P 500, the small-cap premium is S&P 600 minus S&P 500, the value premium is S&P 500 value minus S&P 500 growth, the credit premium is Barclays government bonds minus Barclays corporate bonds, and the sector premium is a dynamic portfolio of S&P 500 sectors created on an out-of-sample rolling basis to replicate the implicit sector exposures of the buyout returns (as described in detail in the methodology section) minus the S&P 500. All data are sourced from Datastream, except for the GXPEI which is sourced from State Street Corporation.

Connection to Prior Research

The regression results presented above suggest that the buyout industry, in aggregate, has:

- a beta slightly less than one with respect to the public market equity risk premium,
- a meaningful exposure to dynamic sector selection,
- short (i.e. negative) exposure to the credit premium,
- a positive, but small, exposure to the value premium,
- a slightly negative, but close to zero, exposure to the small-cap premium, and
- an annualized alpha of approximately 500 basis points.

These findings are contrary to recent research which has argued that the perceived outperformance of the buyout industry can be explained away by adding leverage, small-cap, and value tilts to a public market benchmark. How can we reconcile both sets of results? The answer lies in the methodology that is used to compare public and private performance. Most prior research relies on the useful and well established performance measure of Public Market Equivalent (PME), first introduced by Kaplan and Schoar (2005). PME analysis is appropriate to show that certain public equity investments can achieve a cumulative long-run growth rate comparable to that of private equity (after accounting for cash flow timing), but it cannot offer evidence that these returns occur at the same time or for the same reason. As noted earlier, time series regression identifies factors that co-vary with private equity, while PMEs identify factors that achieve similar levels of return over long periods.

Much research has documented that PMEs for the overall buyout industry decrease as factor adjustments are introduced into the public market benchmark against which the performance multiple is computed. Phalippou (2012) analyzed data from Prequin and incrementally introduced value, size, and leverage tilts to the public market benchmark. Chingono and Rasmussen (2015) form small-cap value portfolios with an emphasis on listed companies with low EBITDA multiples and which they also believe are candidates for balance sheet improvement through debt deleveraging. They show that the cumulative returns of portfolios consisting of these publicly traded stocks outperform a large cap public equity benchmark, and that this outperformance is similar in size to the outperformance of the Cambridge Associates private equity index over the same public equity benchmark. Stafford (2015) creates levered small cap equity portfolios based on the premise that private equity holdings share these same criteria. L'Her et al. (2016) analyze aggregate buyout performance using data from Burgiss and incrementally introduce size, leverage, and sector factors into the PME calculation in a manner similar to Phalippou (2012).

L'Her et al. also carefully document fundamental characteristics of the companies held within buyout funds, which differ on average from the characteristics of companies in the S&P 500. Most notably, buyout investments do indeed have smaller market capitalization, more balance sheet leverage, and different reported sector composition. The hypothesis that these factors explain performance is intuitive, but it is not guaranteed to be true. The weak explanatory power of the small cap and value factors in the regressions shown earlier may arise from other factors that exert a more powerful influence on performance. It is conceivable that the entry prices for companies taken private, as well as their eventual IPO or sale prices upon exit, are determined by other variables that are more influential than the size and value attributes of the underlying companies. For example, Gompers and Lerner (2000) note that the amount of money entering the industry and competing for deals is likely to affect entry pricing and overall profitability. Likewise, Kaplan and Stromberg (2008) and Harris et al (2016) illustrate that the amount of capital committed to private equity each year in proportion to the size of the public equity market value is significantly negatively related to the performance of funds in those vintages. It is also possible that general partners' ability to extract excess returns through governance, management, and financing techniques varies through time. I leave it to future research to investigate further the performance impact of these variables.

The Phalippou and L'Her et al papers report PMEs by vintage year. We can use these results to help reconcile with the time series regression results. Recall that PMEs are performance multiples above a chosen public benchmark. A PME equal to 1 indicates that the since-inception returns of that vintage year are equivalent to the same cash flows invested in the public benchmark. To the extent the PME is below or above 1, it indicates a private market return for that vintage year that deviates from the comparable return of the benchmark. The PMEs therefore represent the unexplained returns of private equity for each vintage year. This quantity is akin to the intercept of a regression plus its residuals. If we add factors to our benchmark that explain variation in performance across vintages, we should expect to see smaller average return deviations for each vintage year. Interestingly, computing the standard deviation of PMEs across vintage years from these two studies suggests that while average PMEs decrease as a result of higher average returns for the adjusted benchmarks, the variation in PMEs does not decrease. This result raises the possibility that other factors play an important role in explaining performance variation across vintages. Exhibit 6 shows the difference of the PMEs from 1 for each vintage year, as reported in the two papers cited, along with simple averages and standard deviations of these series.

Prior research that applies regression analysis to buyout returns has yielded a wide range of estimates for factor exposures, especially for market beta and the small-cap factor. Of seven papers recently summarized by L'Her et al (2017), three reported market betas less than 1, three reported betas greater than 1, and one paper reported a beta equal to 1. High market betas (such as those from Franzoni et al (2012) and Ang et al (2013)) often derive from time series of deal level cash flows. This may reflect the reality that deal level returns are more likely to include extreme positive and negative outcomes as compared to pooled time series returns across the aggregate buyout industry. Interestingly, these same studies report slightly negative exposures to the small-cap factor. As noted previously, two recent papers that document large positive exposure to the small-cap factor derive their estimates from the performance of publicly traded securities that are proximate to private equity as opposed to return outcomes from privately held companies (Jegadeesh et al (2015) and L'Her et al (2017)).

Exhibit 6: PMEs by vintage year (as reported in prior research) in excess of a baseline value of 1

Data as reported by Phalippou (2012)				
Benchmark:	Large cap	Adjusted for value	Adjusted for value and size	Adjusted for value, size, and leverage
Average	0.20	0.02	-0.03	-0.10
Std. Dev.	0.23	0.24	0.23	0.24
PMEs by vintage year, in excess of a baseline value of 1				
1993	0.22	0.29	0.15	-0.02
1994	0.18	0.16	0.06	-0.07
1995	-0.05	-0.18	-0.27	-0.37
1996	0.03	-0.20	-0.28	-0.35
1997	0.18	-0.32	-0.38	-0.47
1998	0.34	-0.32	-0.38	-0.48
1999	0.28	-0.20	-0.25	-0.35
2000	0.58	0.15	0.10	-0.02
2001	0.63	0.37	0.33	0.22
2002	0.30	0.16	0.11	0.05
2003	0.35	0.29	0.22	0.21
2004	0.52	0.48	0.40	0.41
2005	0.18	0.13	0.07	0.08
2006	0.02	-0.04	-0.07	-0.08
2007	0.05	-0.04	-0.06	-0.10
2008	-0.03	-0.10	-0.11	-0.16
2009	-0.06	-0.13	-0.13	-0.18
2010	-0.13	-0.11	-0.12	-0.12

Data as reported by L'Her et al. (2016)				
Benchmark:	Large cap	Adjusted for size	Adjusted for size and leverage	Adjusted for size, leverage, and sector
Average	0.22	0.16	0.06	0.10
Std. Dev.	0.19	0.23	0.22	0.25
PMEs by vintage year, in excess of a baseline value of 1				
1986	0.33	0.47	0.28	0.41
1987	-0.10	-0.09	-0.19	-0.13
1988	0.04	0.03	-0.10	-0.06
1989	0.31	0.35	0.20	0.27
1990	0.14	0.16	0.05	0.11
1991	0.47	0.63	0.48	0.61
1992	0.09	0.23	0.06	0.23
1993	0.04	0.16	0.04	0.15
1994	0.46	0.59	0.45	0.63
1995	0.23	0.18	0.12	0.21
1996	0.03	-0.07	-0.09	-0.04
1997	0.32	-0.04	-0.11	-0.08
1998	0.21	-0.20	-0.27	-0.26
1999	0.23	-0.13	-0.21	-0.21
2000	0.37	0.14	0.05	0.04
2001	0.45	0.27	0.18	0.18
2002	0.44	0.31	0.24	0.23
2003	0.54	0.42	0.37	0.35
2004	0.34	0.26	0.25	0.22
2005	0.22	0.13	0.14	0.07
2006	0.01	-0.07	-0.12	-0.16
2007	-0.01	-0.07	-0.17	-0.20
2008	-0.01	-0.05	-0.16	-0.18

Notes: Based on data from Table 6, Panel B in Phalippou (2012), which reports size-weighted since-inception PMEs for vintages from 1993 to 2010 using Prequin data, and from Table 3 in L'Her et al. (2016), which reports since-inception PMEs for vintages from 1986 to 2008 using Burgiss data.

A simple view of factor correlations

Some readers may wonder if the slightly counterintuitive results shown earlier are a byproduct of a poorly specified regression, or an overcomplicated analysis. For purposes of transparency, it may be helpful to look at simple correlations of each factor with the portion of private equity returns that are not explained by large cap public equities. For this (admittedly simple) analysis, I subtract the non-overlapping calendar year annual returns of large cap public equities from the corresponding annual returns of the buyout index. Doing so implicitly assumes a public market beta of 1 for the buyout industry. I make this assumption purely for simplicity, though it is not too far off from the beta of 0.84 estimated from the same annual data in Exhibit 3. Exhibit 7 presents scatterplots of the excess buyout return versus the small cap, value, credit and sector rotation premiums. While this analysis is not as conclusive as multi-variable regression, it shows that factor relationships are clearly visible on an annual horizon, without requiring a complicated model.

Exhibit 7: Buyout minus Large Cap Equity versus Individual Factors (Calendar Annual Returns)

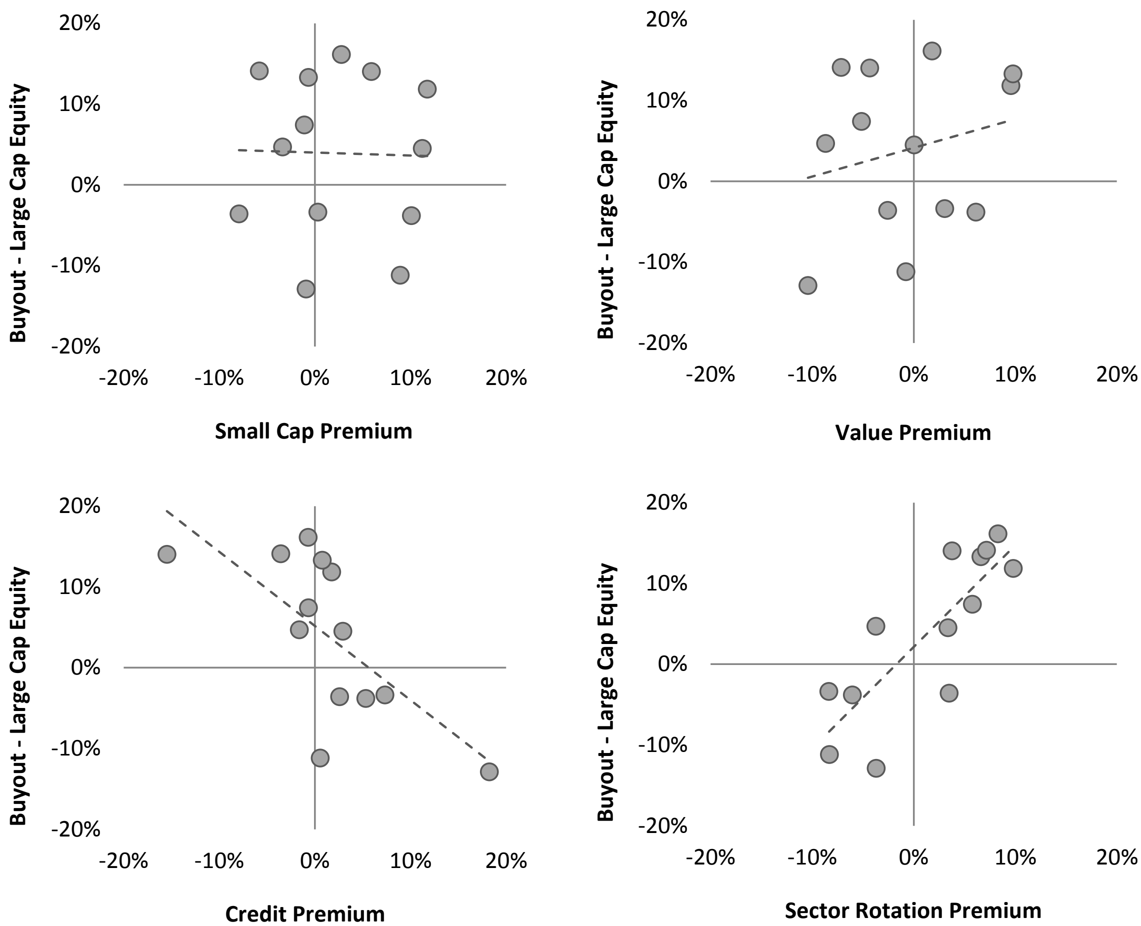

Notes: Based on the same buyout and market returns data as the regression analysis shown previously.

Conclusion

Whereas most prior studies of the aggregate buyout industry focus on Public Market Equivalents (PMEs) and other long-term performance measures, I present evidence of factor relationships based on standard regression analysis. The results are somewhat surprising and stand in contrast to many other studies. I find that the overall beta of private equity to public equity is less than 1, and that small-cap and value factors explain relatively little variation in buyout performance over time. The (short) credit premium and a recently proposed sector rotation premium appear to explain more. It is also possible that private equity adds value through improved governance, financial engineering, long term focus, and other management advantages available only to non-listed companies. Investors would receive these particular benefits in the form of an illiquidity premium, which we should not expect to be able to replicate with liquid securities. In summary, regression analysis is helpful to build intuition and also to offer practical investment solutions that track private equity performance over time. However, it must be approached with care given the smoothing bias inherent in private equity valuations which obscures true relationships on a quarterly basis. While I find evidence that some public market factors can explain a meaningful portion of buyout performance variation, the notion that buyout funds are not redundant to traditional factor premiums should be welcome news for investors who seek asset class diversification or enhanced returns from the private equity buyout asset class.

Appendix A: Review and Illustration of the Standard PME Metric

Public Market Equivalents (PMEs) are dollar-weighted return metrics, which means that they explicitly account for the timing of cash inflows and outflows. Indeed, this is potentially very important for evaluating private equity performance. PME is the multiple of wealth created by the actual private equity investment beyond the wealth that would result from investing the same cash flows with the same timing in a chosen benchmark index. A PME above 1 indicates performance superior to the benchmark, and below 1 is inferior. However, to reconcile bottom-up and top-down analyses, it is also important to note that PMEs offer a measure of average or cumulative return over an entire horizon. They can tell us whether two different investments have the same total return, but they do not reflect the timing of when those returns occur during the horizon. For a given set of cash flows, it is possible to achieve the exact same IRR or PME with a fund that outperforms the S&P 500 early in its life and one that outperforms the S&P 500 late in its life. These two funds would not have correlated performance, but their IRRs and PMEs would be identical.

PMEs generally decrease when the performance benchmark (discount rate) that is used has a long run positive return, even if it is uncorrelated to the actual private equity returns. The following is a contrived example in which private-equity-style contributions are made to the S&P 500 index and are subsequently distributed as profits. Exhibit A1 shows the annual performance of the S&P 500, sample assumed contributions, and the resulting distributions made in subsequent years. Exhibit A2 shows the Internal Rate of Return (IRR) for this investment over its lifespan, which is 4.6% annualized. The internal rate of return is the fixed discount rate that renders the present value of all cash inflows and outflows equal to zero.

Exhibit A2 also shows this investment's PME, which assumes similarly timed investments in a benchmark reference index. Because the benchmark is equal to the actual investment in this case, the PME is equal to 1 as expected.

Exhibit A1: Hypothetical cash flows into, and out of, S&P 500

	Investment return (S&P 500)	Contribution	Distribution	*Distribution is growth of…*
12/31/2005		-20.00	0.00	
12/31/2006	15.8%	-40.00	0.00	
12/31/2007	5.5%	-25.00	22.76	*10 from 2005, 10 from 2006*
12/31/2008	-37.0%	0.00	0.00	
12/31/2009	26.5%	-10.00	7.97	*10 from 2007*
12/31/2010	15.1%	0.00	0.00	
12/31/2011	2.1%	-5.00	29.11	*20 from 2006, 10 from 2007*
12/31/2012	16.0%	0.00	16.89	*10 from 2006, 5 from 2007*
12/31/2013	32.4%	0.00	26.58	*10 from 2005, 5 from 2009*
12/31/2014	13.7%	0.00	18.99	*5 from 2009, 5 from 2011*

Notes: S&P 500 data sourced from Datastream. Contributions and distributions are completely hypothetical.

Exhibit A2: IRR, PME and Direct Alpha for hypothetical investment in S&P 500

	Investment return (S&P 500)	Contribution	Distribution	Net cash flow	Benchmark return (S&P 500)	Contribution discounted by benchmark	Distribution discounted by benchmark
12/31/2005		-20.00	0.00	-20.00		-20.00	0.00
12/31/2006	15.8%	-40.00	0.00	-40.00	15.8%	-34.54	0.00
12/31/2007	5.5%	-25.00	22.76	-2.24	5.5%	-20.47	18.64
12/31/2008	-37.0%	0.00	0.00	0.00	-37.0%	0.00	0.00
12/31/2009	26.5%	-10.00	7.97	-2.03	26.5%	-10.27	8.19
12/31/2010	15.1%	0.00	0.00	0.00	15.1%	0.00	0.00
12/31/2011	2.1%	-5.00	29.11	24.11	2.1%	-4.37	25.46
12/31/2012	16.0%	0.00	16.89	16.89	16.0%	0.00	12.73
12/31/2013	32.4%	0.00	26.58	26.58	32.4%	0.00	15.14
12/31/2014	13.7%	0.00	18.99	18.99	13.7%	0.00	9.51
						-89.66	89.66
IRR	**4.6%**						
PME	**1.00**						

Notes: S&P 500 data sourced from Datastream.

Next, suppose we invest the same cash flows not in the S&P 500, but in a different "private equity" investment. For illustrative purposes, I simulated a series of completely random returns that will now represent the actual investment, as shown in the first column of Exhibit A3. This change to the investment return series affects the distributions and net cash flows in the example. The IRR increases to 9.9%. The PME of 1.30 indicates that the investment produced returns superior to the same cash flows to and from the S&P 500.

Exhibit A3: IRR and PME for hypothetical "private equity" discounted at S&P 500

	Investment return (random)	Contribution	Distribution	Net cash flow	Benchmark return (S&P 500)	Contribution discounted by benchmark	Distribution discounted by benchmark	Net cash flow discounted by benchmark
12/31/2005		-20.00	**0.00**	**-20.00**		-20.00	**0.00**	**-20.00**
12/31/2006	**0.1%**	-40.00	**0.00**	**-40.00**	15.8%	-34.54	**0.00**	**-34.54**
12/31/2007	**18.8%**	-25.00	**23.78**	**-1.22**	5.5%	-20.47	**19.47**	**-1.00**
12/31/2008	**-0.3%**	0.00	**0.00**	**0.00**	-37.0%	0.00	**0.00**	**0.00**
12/31/2009	**28.0%**	-10.00	**12.77**	**2.77**	26.5%	-10.27	**13.12**	**2.84**
12/31/2010	**11.9%**	0.00	**0.00**	**0.00**	15.1%	0.00	**0.00**	**0.00**
12/31/2011	**13.2%**	-5.00	**54.59**	**49.59**	2.1%	-4.37	**47.74**	**43.37**
12/31/2012	**-17.4%**	0.00	**22.54**	**22.54**	16.0%	0.00	**16.99**	**16.99**
12/31/2013	**4.0%**	0.00	**21.95**	**21.95**	32.4%	0.00	**12.50**	**12.50**
12/31/2014	**30.7%**	0.00	**12.72**	**12.72**	13.7%	0.00	**6.37**	**6.37**
						-89.66	**116.18**	
IRR	**9.9%**							
PME	**1.30**							

Notes: S&P 500 data sourced from Datastream.

Next, for illustrative purposes, I generate a new independent series of completely random returns to add to the benchmark as a hypothetical risk premium. Exhibit A4 shows the relationship between this new uncorrelated variable and the actual investment. The two are essentially uncorrelated (correlation = -0.01). Clearly one does not explain the other.

Exhibit A4: Random noise relationship to actual investment (annual returns)

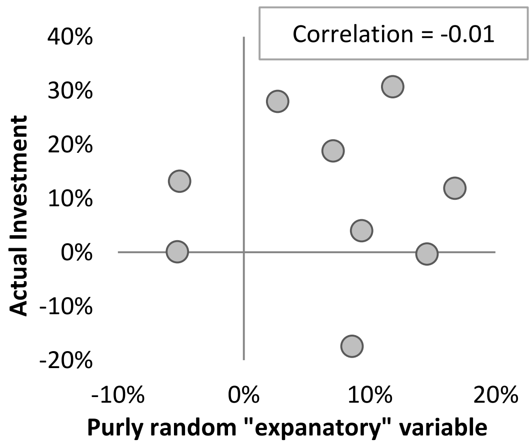

Notes: Data shown are randomly generated and purely hypothetical.

Exhibit A5 shows that PME decreases as a result of adding the uncorrelated random noise to the benchmark. Both of the random series (the actual investment and the premium added to the benchmark) contribute positive growth, but we know from construction that they are uncorrelated. This intentionally contrived example is merely meant to illustrate that PME analysis by itself does not guarantee an explanatory relationship between two variables.

Exhibit A5: IRR and PME for hypothetical "private equity" discounted at S&P 500 + random noise

	Investment return (random)	Contribution	Distribution	Net cash flow	**Benchmark return (S&P 500 + noise)**	Contribution discounted by benchmark	Distribution discounted by benchmark
12/31/2005		-20.00	0.00	-20.00		**-20.00**	**0.00**
12/31/2006	0.1%	-40.00	0.00	-40.00	**10.5%**	**-36.20**	**0.00**
12/31/2007	18.8%	-25.00	22.76	-2.24	**12.6%**	**-20.09**	**19.11**
12/31/2008	-0.3%	0.00	0.00	0.00	**-22.5%**	**0.00**	**0.00**
12/31/2009	28.0%	-10.00	7.97	-2.03	**29.1%**	**-8.03**	**10.25**
12/31/2010	11.9%	0.00	0.00	0.00	**31.8%**	**0.00**	**0.00**
12/31/2011	13.2%	-5.00	29.11	24.11	**-3.0%**	**-3.14**	**34.26**
12/31/2012	-17.4%	0.00	16.89	16.89	**24.6%**	**0.00**	**11.35**
12/31/2013	4.0%	0.00	26.58	26.58	**41.8%**	**0.00**	**7.80**
12/31/2014	30.7%	0.00	18.99	18.99	**25.5%**	**0.00**	**3.60**
						-87.45	**86.36**
IRR	**9.9%**						
PME	**0.99**						

Notes: S&P 500 data sourced from Datastream.

Appendix B: Unbiasedness and standard errors for coefficients from overlapping data

The use of overlapping data requires relaxing some of the assumptions that are made in standard Ordinary Least Squares (OLS) regression analysis. My goal in this section is to illustrate that coefficients derived from overlapping data are unbiased estimates, and also to motivate the use of an adjustment factor in-line with the now standard Newey-West approach to adjusted standard errors for serial correlation in residuals. My treatment of these subjects is not meant to be comprehensive nor is it meant to offer a rigorous or general proof. Rather, I hope to offer some intuition for these adjustments based on concepts from financial analysis with math that is simpler than that of a rigorous statistical proof.[3]

We begin by specifying a traditional linear regression model with a scalar dependent variable y_t, a row vector of N independent variables $\boldsymbol{x}_t$, a column vector of N coefficients $\boldsymbol{\beta}$, and a scalar error u_t:

$$y_t = \boldsymbol{x}_t\boldsymbol{\beta} + u_t$$

We estimate the model based on a time series of dependent variable observations (a $T \times 1$ vector $\boldsymbol{y}$) and a time series of observations for each independent variable (a $T \times N$ matrix $\boldsymbol{X}$). The standard estimate of $\boldsymbol{\beta}$ is given by:

$$\widehat{\boldsymbol{\beta}} = (\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\boldsymbol{y}$$

As is typical, we assume that the expected value of the residuals conditional on the dependent variables is zero: $E(u|\boldsymbol{X}) = 0$. In practice, this assumption means that the explanatory variables are exogenous, and it does not allow inclusion of lagged dependent variables. This is fine for our purposes.

Let us first examine the expected value of the coefficient estimates, to ensure they are unbiased. First we rewrite the estimated coefficients as:

$$\widehat{\boldsymbol{\beta}} = (\boldsymbol{X}'\boldsymbol{X})^{-1}\boldsymbol{X}'\boldsymbol{y} = (\boldsymbol{X}'\boldsymbol{X})^{-\mathbf{1}}\boldsymbol{X}'(\boldsymbol{X}\,\boldsymbol{\beta} + \boldsymbol{u}) = \boldsymbol{\beta} + (\boldsymbol{X}'\boldsymbol{X})^{-\mathbf{1}}\boldsymbol{X}'\boldsymbol{u}$$

Taking the expectation conditional on $\boldsymbol{X}$, we have:

$$E\left(\widehat{\boldsymbol{\beta}}\middle|\boldsymbol{X}\right) = \boldsymbol{\beta} + (\boldsymbol{X}'\boldsymbol{X})^{-\mathbf{1}}\boldsymbol{X}'E(\boldsymbol{u}|\boldsymbol{X}) = \boldsymbol{\beta}$$

To evaluate the statistical significance of $\widehat{\boldsymbol{\beta}}$, we must estimate its variance conditional on $\boldsymbol{X}$. Due to the fact that the expected value of the residuals conditional on $\boldsymbol{X}$ is zero, we can express the variance as:

$$Var\left(\widehat{\boldsymbol{\beta}}\middle|\boldsymbol{X}\right) = \boldsymbol{Var}((\boldsymbol{X}'\boldsymbol{X})^{-\mathbf{1}}\boldsymbol{X}'\boldsymbol{u}|\boldsymbol{X}) = \frac{\mathbf{1}}{\boldsymbol{T}}\boldsymbol{E}((\boldsymbol{X}'\boldsymbol{X})^{-\mathbf{1}}\boldsymbol{X}'\boldsymbol{u}\boldsymbol{u}'\boldsymbol{X}(\boldsymbol{X}'\boldsymbol{X})^{-\mathbf{1}}|\boldsymbol{X})$$

For those less comfortable with matrix algebra, it can be helpful to express the variance for an individual element i of $\boldsymbol{\beta}$. First, let's introduce a simplification for expository purposes: assume that the expected returns of the independent variables are all zero, $E(\boldsymbol{X}) = \mathbf{0}$. This simplification is not essential to the proof, but it allows for $\boldsymbol{X}'\boldsymbol{X}$ to be written as $T\boldsymbol{\Omega}$ where $\boldsymbol{\Omega}$ is the covariance matrix of $\mathbf{X}$ (if the means were not zero, we would have to account for an additional constant term, which complicates the intuition). We now have:

$$\widehat{\boldsymbol{\beta}} = \frac{1}{T}\boldsymbol{\Omega}^{-1}\boldsymbol{X}'\boldsymbol{u}$$

We can express $\widehat{\boldsymbol{\beta}}_i$ in terms of the sums implied by matrix multiplication:

$$\widehat{\boldsymbol{\beta}}_i = \frac{1}{T}\sum_{t=1}^{T}\left(\sum_{k=1}^{N} c_{ik}x_{kt}\right)u_t = \frac{1}{T}\sum_{t=1}^{T} a_{it}u_t$$

In this expression, we define c_{ik} as the i-th row and k-th column of the inverse covariance matrix. Viewed in this way, we can think of each coefficient estimate as a weighted average of the errors, u_t. For now, we represent the terms that multiply u_t as a_{it}. The following expression describes the variance of the coefficient estimate:

$$Var(\widehat{\boldsymbol{\beta}}_i|\boldsymbol{X}) = E\left[\left(\frac{1}{T}\sum_{t=1}^{T} a_{it}u_t\right)^2\right] = \frac{1}{T^2}\left(\sigma^2_{a_{it}u_t}\left(T + 2\sum_{j=1}^{T-1}(T-j)\rho_{a_{it}u_t,a_{it+j}u_{t+j}}\right)\right)$$

$$Var(\widehat{\boldsymbol{\beta}}_i|\boldsymbol{X}) = \frac{\sigma^2_{a_{it}u_t}}{T}\left(1 + 2\sum_{j=1}^{T-1}\left(\frac{T-j}{T}\right)\rho_{a_{it}u_t,a_{it+j}u_{t+j}}\right)$$

In this expression, $a_{it}u_t$ represents the product of these two variables and $\rho_{a_{it}u_t,a_{it+j}u_{t+j}}$ is the correlation between the product of these variables and their j-th lag. Traditional regression analysis assumes that these autocorrelations are all zero. If they are not zero, the final term in the above equation represents an adjustment factor. We can define a variable $v_{it} = a_{it}u_t$ and write the adjustment factor as shown below. In practice, we can easily compute these adjustments from estimated residuals $\widehat{\boldsymbol{u}}$ and with a_{it} terms that account for non-zero means in the independent variables, which I do in the analysis presented earlier. I use $Q = 4$ for overlapping annual intervals of quarterly returns, as it is customary to truncate the number of lags incorporated to account for the expected degree of inherent autocorrelation due to the overlapping data.

$$Adjustment\ factor = 1 + 2\sum_{j=1}^{Q-1}\left(\frac{T-j}{T}\right)\rho_{v_{it},v_{it+j}}$$

Notes

I thank Megan Czasonis, Will Kinlaw, Mark Kritzman, Nan Zhang and an anonymous referee for helpful comments.

References

Ang, Andrew, Bingxu Chen, William N. Goetzmann and Ludovic Phalippou. 2013. "Estimating Private Equity Returns from Limited Partner Cash Flows." Working paper (November 25).

Appelbaum, Eileen and Rosemary Batt. 2016. "Are Low Private Equity Returns the New Normal?" Center for Economic and Policy Research (June).

Asness, Clifford, Robert Krail and John Liew. 2001. "Do Hedge Funds Hedge?" Journal of Portfolio Management, vol. 28, no. 1 (Fall).

Axelson, Ulf, Morten Sorensen and Per Stromberg. 2014. "Alpha and Beta of Buyout Deals: A Jump CAPM for long-term illiquid investments." Working paper, London School of Economics.

Chingono, Brian and Daniel Rasmussen. 2015. "Leveraged Small Value Equities." Working paper (August).

Czasonis, Megan, Mark Kritzman and David Turkington. 2017. "Private Equity Valuations and Public Equity Performance." MIT Sloan School of Management Working Paper 5237-17 (September 15).

Ewens, Michael, Charles M. Jones and Matthew Rhodes-Kropf. "The Price of Diversifiable Risk in Venture Capital and Private Equity." 2013. The Review of Financial Studies, vol. 26, no. 8.

Franzoni, Francesco A., Eric Nowak and Ludovic Phalippou. 2012. "Private Equity Performance and Liquidity Risk." Journal of Finance, vol. 67, no. 6 (December).

Gompers, Paul A. and Josh Lerner. 2000. “Money Chasing Deals? The Impact of Fund Inflows on Private Equity Valuations.” Journal of Financial Economics, vol 55.

Harris, Robert S., Tim Jenkinson and Steve N. Kaplan. 2016. “How do Private Equity Investments Perform Compared to Public Equity?” Journal of Investment Management, vol. 14, no. 3 (Third Quarter).

Jegadeesh, Narasimhan, Roman Kraussl and Joshua M. Pollet. 2015. “Risk and Expected Returns of Private Equity Investments: Evidence Based on Market Prices.” Review of Financial Studies, vol. 28, no. 12.

Kaplan, Steven N. and Antoinette Schoar. 2005. “Private Equity Performance: Returns, Persistence and Capital Flows.” Journal of Finance, vol. 60, no. 4 (August).

Kaplan, Steven N. and Per Stromberg. 2008. “Leverage Buyouts and Private Equity.” Journal of Economic Perspectives, vol. 22, no. 4.

Kinlaw, Will, Mark Kritzman and Jason Mao. 2015. “The Components of Private Equity Performance: Implications for Portfolio Choice.” Journal of Alternative Investments, vol. 18, no. 2 (Fall).

Kinlaw, Will, Mark Kritzman and David Turkington. 2013. “Liquidity and Portfolio Choice: A Unified Approach.” Journal of Portfolio Management, vol. 39, no. 2 (Winter).

Kinlaw, Will, Mark Kritzman and David Turkington. 2014. “The Divergence of High- and Low-Frequency Estimation: Causes and Consequences.” Journal of Portfolio Management, vol. 40, no. 5 (40th Anniversary Edition).

L’Her, Jean-Francois, Ram Karthik and Stephanie Desrosiers. 2017. “How to Calibrate the Risk of Buyout Investments? Through Butyout-Backed Initial Public Offerings.” The Journal of Investment Management, vol. 15, no. 4 (Fourth Quarter).

L’Her, Jean-Francois, Rossitsa Stoyanova, Kathryn Shaw, William Scott and Charissa Lai. 2016. “A Bottom-Up Approach to the Risk-Adjusted Performance of the Buyout Fund Market.” Financial Analysts Journal, vol. 72, no. 4 (July/August).

Newey, Whitney K. and Kenneth D. West. 1987. “A Simple, Positive-Semi-Definite Heteroskedasticity, and Autocorrelation Consistent Covariance Matrix.” Econometrica, vol. 55, no. 3.

Phalippou, Ludovic. 2012. “Performance of Buyout Funds Revisited?” Working paper, obtained via SSRN (November).

Stafford, Erik. 2015. “Replicating Private Equity with Value Investing, Homemade Leverage, and Hold-to-Maturity Accounting.” Working paper.

Tibshirani, Robert. 1996. “Regression Shrinkage and Selection via the Lasso.” Journal of the Royal Statistical Society, Series B (Methodological), vol. 58, no. 1.

[1] See Appendix A for an illustrative example.
[2] Kinlaw, Kritzman and Turkington (2014) provide formulas that use lagged correlations to directly relate single-period standard deviations and correlations to their multi-period counterparts.
[3] For more details, readers may refer to Newey and West (1987), for example.