JOURNAL

THE CAPCO INSTITUTE JOURNAL OF FINANCIAL TRANSFORMATION 04.2017 N°45

Investments

The Power of "Negative Beta": Why Every Portfolio Should Include Private Equity

Andrew Freeman, Iordanis Karagiannidis, D. Sykes Wilford

RECIPIENT OF THE APEX AWARD FOR PUBLICATION EXCELLENCE

TRANSFORMATION

CAPCO

EMPOWERING THE FINANCIAL WORLD

Pushing the pace of financial technology, we help our clients solve technology challenges for their business – whether it's capital markets in Mumbai or community banking in Macon.

We leverage knowledge and insights from our clients around the world:

20,000 clients in towns everywhere are becoming more efficient, modern and scalable.

27 billion transactions processed help solve clients' challenges – big and small.

$9 trillion moved across the globe in a single year empowers our clients' communities to build storefronts, homes and careers.

55,000 hearts and minds have joined forces to bring you greater capabilities in even the smallest places.

Empowering the Financial World

FISGLOBAL.COM

FIS

JOURNAL

THE CAPCO INSTITUTE JOURNAL OF FINANCIAL TRANSFORMATION

Recipient of the Apex Award for Publication Excellence

Transformation

FinTech/RegTech

Investments

Banking

The Power of "Negative Beta": Why Every Portfolio Should Include Private Equity

Andrew Freeman – Risk fellow, Centre for Risk Studies, Judge Business School, University of Cambridge
Iordanis Karagiannidis – Assistant Professor of Finance, The Citadel
D. Sykes Wilford – Hipp Chair Professor of Business and Finance, the Citadel and Senior Advisor to Access Corporate Finance[1]

Abstract
Building on previous work we analyze the option-like characteristics of investment in private equity. While the main academic focus has been on the disputed ability of this asset class to produce above-average risk-adjusted returns, our focus is on the underappreciated role played by volatility in private equity (PE) performance. Our conclusion is that PE is a much more attractive asset class (lower risk) than commonly believed. In contrast to most approaches, we focus on the influence of the options built into the private-equity investment business model primarily from the perspective of the fund manager, or General Partner (GP). As the owner of call options on the underlying investors' capital commitments initially and later of complex put options as assets can be retained before being sold, the GP is well placed to take advantage of market volatility, particularly during bad times. We posit the existence of "negative beta" as a reason for large investors to make much bigger allocations to PE than are typical.

1 The opinions presented in this paper reflect those of the authors and not necessarily those of the respective institutions noted. The authors can be contacted respectively at af489@cam.ac.uk, iordanis@citadel.edu, wsykes@dsykeswilford.com. The authors wish to thank Juan Montalvo Bressi for his comments.

INTRODUCTION

In recent decades, it has become fashionable to include an allocation to so-called alternative assets in most professionally managed investment portfolios. These assets have included hedge funds, venture capital, and private equity, as well as occasionally exotics such as art and real estate. The rationale for such investments has typically been that they provide the only free lunch in investing where assets in public markets represent the principal portfolio components – the benefits of non-correlated returns via diversification. While relatively small allocations in the 2-5% range have been the norm, there have been occasional examples of institutions, notably the Yale endowment, taking larger exposures, thereby attracting a mixture of approval and opprobrium [Provence (2008); Considine (2010); Kaplan and Sensoy (2014)].

Private equity (PE) is generally considered as the critical "alternative" asset class [Bain (2016)]. Like other alternative investments, PE offers benefits in the form of diversification from public markets, deploying information advantages and a variety of approaches to create value from corporate assets [Kaplan and Schoar (2005); French et al. (1987)]. Investors have traditionally paid high fees in return for access to the asset class, but this has not been without controversy. Investors remain doubtful about the efficacy of the PE model on a risk-adjusted basis.

However, we identify a much more intriguing argument for including PE in an asset allocation. Building on earlier work, it is suggested that PE offers unique option-like characteristics in the form of embedded "negative beta." Simply put, in a financial crisis PE will perform so differently from publicly traded assets that for a typical large investor it would constitute a buffer against a downturn that is in proportion to the relevant portfolio's allocation. As Bain (2016) stated: "The global financial crisis claimed many victims but barely put a dent in PE assets under management." We believe that this can be at least partly explained by the risk control that is intrinsic to the PE business model. However, we go further in arguing that even today's PE assets are underappreciated by most investors.

The starting point for the argument can be found in Freeman and Wilford (2016). They identified and analyzed the options embedded in the fund structures of a hypothetical PE firm; during its investment period, a PE fund owns a series of call options that allow it to demand investment capital from investors who have signed up at the beginning of the life of the fund. For four or five years after a fund's creation, the PE firm can use guaranteed liquidity to buy assets, albeit highly idiosyncratic ones.[2] The option value of this liquidity rises sharply if financial markets are stressed, so a fund that owns calls can buy assets cheaply in bad times. In addition, a fund that is fully invested undergoes a neat reversal of its former position. From being long the right to call capital, it is now the owner of complex put options. Because the fund's objective is now to dispose of assets at or above some notional exit multiples, it can judge market conditions and decide when and how to exercise its right to sell. Crucially, the typical fund has the option of waiting for quite long periods in order to avoid selling if markets are distressed. It is in its, and its investors', interests to delay until the price is right. Ownership of this complex option constitutes a key element of the overall "negative beta" of PE.[3]

Freeman and Wilford (2016) provided a simple initial simulation of PE option values and their response to movements in the volatility of public markets. This supported the suggestion that using embedded options as a way to analyze PE risk from a risk management perspective might be beneficial. Tentative conclusions were offered about the possible positive influence of PE embedded options on systemic risk, as well as on the possibility that large investors might be underexposed to this asset class because of a failure to appreciate its portfolio risk-management potential.

This paper further explores the rationale for including PE in an asset allocation on the explicit grounds that it has unrivaled risk-return characteristics, crucially once the risk elements in the equation are better understood. Arguments against PE typically cite high management fees as unjustified given the level of performance once adjusted using conventional risk-return analysis [SEC (2015); Sorensen et al. (2014)]. This objection fades in light of our insight that the embedded option values, intrinsically difficult to model though they are, mean that PE assets are in reality very different from their typical characterization in the literature. In this paper, we extend the original analysis by undertaking simple simulations to show the effect of optionality on the volatility of PE assets compared to a standard equity market investment.

2 A typical buy-out fund, for example, will make between 6 and 12 investments over 5 years

3 Thomas Meyer recognized as long ago as 2007 that real options could offer new insights into private equity. See Mathonet and Meyer (2007) as well as Meyer (2014). However, this considers real options from the perspective of the investor. Our focus is on the GP and the interaction of optionality with both investors and the market. Separately, Chen et al. (2008) focused on the optionality of a GP's investment portfolio. Thanks to Thomas Meyer for the references to his work.

PE AS A PORTFOLIO RISK MITIGATION TOOL

A useful starting point for evaluating the value that PE can bring to a portfolio as a risk mitigation tool is a simple side-by-side comparison. Taking the S&P as our base case diversified portfolio, the analysis builds simplified portfolio allocations utilizing monthly data for the period January 2005 to December 2014. These data encompass the 2008 – 09 crisis and the volatility spikes that occurred. Further, initial modeling controls for performance by assuming that it is equal to the S&P for each period except for the implicit value of the options in the PE fund. In this way, the focus is placed directly on volatility. For example, if the S&P goes up from a base of 100 to 110, but the fund has not yet invested, then the value of that investment remains 100 plus the value of the option to time the investment by calling capital. If the funds were already invested and the S&P goes up by 10% then the value of the investment is assumed to rise by 10%; however, this value is not the overall value of the PE investment. As shown in Freeman and Wilford (2016), the actual value includes the value of the right to put the component parts of the PE investment portfolio to the market at a chosen time, or opportunistically if the price is right. Thus, the value of the PE investment is the combination of the call and put options and the invested capital as indicated by the S&P. Simply using a standard option pricing model,[4]

PE fund value = call option value (C) + investment growing at the market rate + put option value (P), where

Call option value $(C)=S\times N(d_1) - X\times e^{-rT}\times N(d_2)$

where $d_1=\frac{\ln\left(\frac{S}{X}\right)+\left(r+\frac{\sigma^2}{2}\right)\times T}{\sigma\times\sqrt{T}}$ and $d_2 = d_1 - \sigma\times\sqrt{T}$

and C = call option value, S = the amount still to be called, X = the amount still to be called, R = 2%, T = time to maturity (time to reach month 60) in years, and sigma = the monthly volatility (VIX)

And,

Put option value $(P) = C + X \times e^{-rt} - S$, however now

S = current investment value, and X = forward value of that investment at time 120 using a 2% rate or a 25% of the original investment as a cap.[5]

Put option value remains zero if no investment has been called.

As one can observe from the calculation process above for P one must be careful in understanding the concept of valuing the put. This analysis makes the simplifying assumption that the put option's strike price is a forward value at each date of the option's price calculation unless the forward is greater than 25% of the original investment (this will be modified for examples where the takeout target is greater).[6] In reality, there is not a true strike price to calculate against at each calculation date, suggesting that our measure will have errors. Also, the option value will tend to move with the spot price (as the S&P moves) not just from spot movements but also from changes in the strike price implied by our assumption of how the forward is calculated. Such an analysis does take away the subjectivity implicit in alternative measures of the strike that may be considered, and will be altered if the takeout target is altered for the experiment. More importantly, the impacts of time and volatility on the value of the right to time taking an investment to market are consistently considered along the path of pricing.

In making the assumptions concerning the strike of P, the goal is not to bias the measures of risk that will be estimated using this options approach. Ultimately, due to the methodology used, all of the investments will have a return similar to that of the S&P (in some cases leveraged) or less, depending on the timing of investment, which is generated randomly to avoid unintended bias. Over time, however, the PE fund will take a significantly different path to that end. The call options will go from valuable to worthless once exercised. The put options will first fade and finally cease as investments are placed in the market. During the lives of the options the value of the fund and the S&P will differ significantly.

4 The standard references are Black and Scholes (1973) and Merton (1973).

5 For example: If S = $1,200,000 at time 20, then X= $1,200,000 × $(1+0.02)^{8.3}$, where 8.33 is the time remaining in years (100 months /12). If S = $1,300,000 at time 20, then x = 1,250,000. If S = 900,000 at time 20, then x = $900{,}000*(1+0.02)^{8.3}$, where 8.33 is the time remaining in years.

6 We acknowledge that the put option takes the form of a complex set of real options rather than a pure financial option, hence the need to simplify in the way we propose to calculate the strike price at each point in time. The 25% cap is arbitrary and based upon our modeling the date of when an investment goes to the market, which is set at a 25% gain (no leverage version) in the value of the investment. One can correctly argue that without a set strike for the put option it cannot be truly valued as a typical financial option. We attempt to deal with this problem by biasing its "value" downward, in a sense, through the use of the forward and cap technique; this is not a well-defined "closed form" solution to the problem of valuing a set of real options, however it highlights the critical issues underlying any option value – time until expiry and volatility – because without volatility the exercise of pricing any option is meaningless.

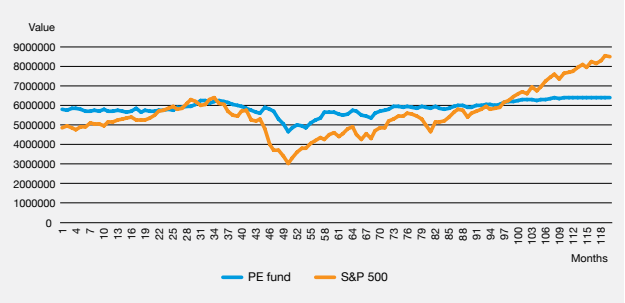

Figure 1 – Path of $5 million investment

	Call 1	Call 2	Call 3	Call 4	Call 5
Time	9	12	32	39	53
	Put 1	Put 2	Put 3	Put 4	Put 5
Time	98	99	109	104	63

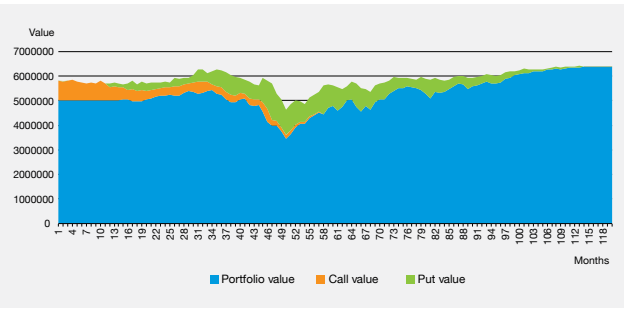

Figure 1a – Value of each component

Typical PE allocations are intended to produce returns superior to those available from public markets, but our analysis begins with the assumption that outperformance is not the goal. The focus is on volatility, and this allows us also to generate the correlation of returns to the S&P. During a long period of low volatility and good performance of the S&P, the path of returns for the fund will differ based upon time of entry into and exit from an investment plus how the option value changes as volatility (even if small) moves.

The interesting case from a risk diversification perspective is what happens to that path of returns when there is a volatile market, a crisis, a rise in fear and consequent plunge in value. In these instances, the value of the embedded options rises significantly, offsetting the mark-to-market plunge in value that would normally (directly in our simulations) be assumed to occur if the embedded options in the fund are ignored.[7] The path of the value of the fund will be different from that of the S&P, even if in the end the overall return is the same and, therefore, the correlation of the value of the fund to the S&P should not be perfect. The degree to which it will vary will be sensitive to entry and exit of investments (creation and extension of the options) and the movement of market volatility. Aware of its limitations, we use the VIX as the best available representative measure for expected volatility.

Figure 1 depicts the path generated for the S&P and one of many possible simulation paths. The simulations assume that the call is made by the PE fund randomly with all calls made within five years; many possible paths can thus be generated. Since all the calls are made randomly during the first five years of the life of the fund, the value of the calls has been dissipated before the depth of the recession. In these simulations, much of the call value expired before the worst period during the life of the data. Calls are made between the 9th and 53rd months. Interestingly, the last call occurs at one of the low points of the recession and will be the first investment to be monetized just 10 months later.[8]

Offsetting the loss of the value of the call, the put option value commences after a call is made. As the crisis volatility increases, we begin to observe the impact of volatility on the value of the PE fund. Through the recession, as expressed in lower stock market prices, the fund loses money at a much slower pace, even as the shortening time to maturity impacts the option value negatively. Finally, our decision to cap the value of the put as a measure of the forward rate comes into play during the deep downturn by limiting the value of the put due to the treatment of the strike price for calls made at much higher levels of the S&P. In general, however, the put values impact the overall volatility of the fund to offset the slide in the S&P (the underlying).[9]

The valuation impact of the options is more easily observed in Figure 1a. Initially, all of the value is in the call, then at 9 months and 12 months two of the calls disappear (the investor commitments are called) while the puts now become valuable.

7 A mark-to-market mentality is, of course, the antithesis of the PE approach and is resisted fiercely by the industry. Our analysis is a rare and admittedly stylized example of where the direct comparison in fact does not harm the underlying justification for PE investing.

8 See bottom of figure for a list of put and call timings. For example, the first call is made month 9 of the 5-year period, chosen randomly, and that investment is not taken to the market until month 98, when its value is now 1.25% of the original investment.

9 Since the "correct" value for the put is always in doubt due to the methodology of estimating the financial put, one may wonder in what ways the volatility of the PE fund may be biased if alternative estimation procedures were used. We believe this methodology does not bias the volatility of the fund downward, and in some instances perhaps the opposite is true.

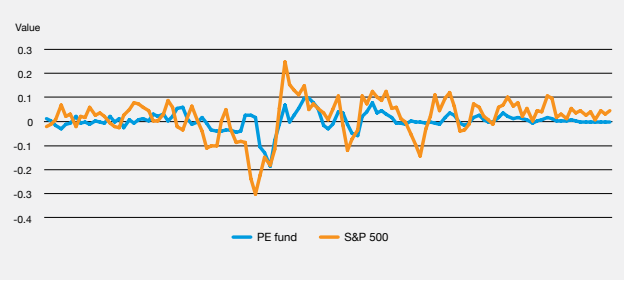

Figure 2 – 3-month rolling returns

	PE fund		
	Mean	**Min**	**Max**
PE fund risk – standard deviation (returns)	9.01%	6.71%	15.78%
S&P risk – standard deviation (returns)	14.62%	14.62%	14.62%
Correlation (PE fund, S&P)	0.39	0.15	0.57
Ratio S&P risk/PF fund risk	1.62	2.18	0.93

Table 1 – Average volatilities

The bulk of the value remains in the underlying that varies depending on the movement of the S&P. While the S&P is losing value quickly as the crisis occurs, the path for the PE fund is smoothing out the impact resulting from the rise in the value of the puts; volatility is rising. As noted, the first put value drops to zero as it is cashed out on month 63, after which 4 puts remain, slowly losing value until the underlying investments are cashed out in months 98, 99, 104, and 109. At this juncture, all of the put values are gone from the estimates and the cash is left sitting in the fund ready for distribution to investors.[10]

With these caveats in mind, we observe biases in returns and volatility that may be introduced due to the methodology chosen for simulations. Still, the positive impact of the options during the crisis is evident in the form of obviously lower volatility, with any bias occurring toward the tail of the 10-year period.

In Figure 2, we look at the volatilities of the fund and S&P more closely. To take advantage of the full dataset while making the numbers a bit more meaningful the data are smoothed by using 3-month rolling average index levels of value for both the PE fund and the S&P. Specifically, the formula is:

$$\text{PE fund return}_t = \frac{\text{PE fund value}_t - \text{PE fund value}_{t-3}}{\text{PE fund value}_{t-3}}$$

$$\text{S\&P return}_t = \frac{\text{S\&P index value}_t - \text{S\&P index value}_{t-3}}{\text{S\&P index value}_{t-3}}$$

It is clear that the level of the volatility for the fund is systematically lower than that of the S&P. However, these observations are for one path (one we will hold in subsequent figures for comparison purposes). Because the calls are generated randomly, other paths could look entirely different. For example, if the random generator selected all the calls in months 50 through 60 the cash out would occur quickly, thereby implying an even lower volatility.

We next ran multiple simulations to calculate the average volatility of the simulations, as well as a maximum and a minimum volatility paths. Just as with the "average" path volatility pattern, we can calculate the correlation of the PE fund to the S&P. Table 1 clearly shows that the volatility of the PE fund is much lower than that of the S&P in the maximum, minimum, and average case simulations. The implication that PE is a highly attractive asset class should be clear even from this simple analysis and its assumptions that the call and put options have value for the PE fund – despite the fact that this is not a typical way to view the asset class today.

Table 1 presents the data needed to create an optimal portfolio based upon the noted paths, assuming the same expected return for both the S&P and the PE Fund. We can observe that the volatility ratios show that the embedded PE options dampen the volatility significantly, with the random call process implying a correlation, in all three cases, significantly less than one. In the maximum path with the highest standard deviation the correlation is still significantly less than one, while in the average path the correlation is only 0.39. This is important because the base for the investment is the S&P. We could be satisfied at this juncture that PE fund investments do indeed provide diversification, even in this simplistic form.

10 In effect, we can ignore the tail of the figures. In reality, cash is distributed to investors as the puts are realized, so the fund's assets naturally shrink. Occasionally funds for distribution are held in escrow for future investment by the PE fund. The eventual returns received by investors are a function of how much, and when, their cash is put to work. More analysis needs to be done on the effects of cash flow movements in PE investment cycles, but this is not necessary for our purposes in this article.

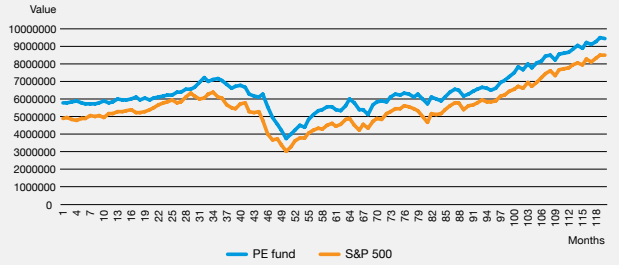

Figure 3 – Path of $5 million investment

	Call 1	Call 2	Call 3	Call 4	Call 5
Time	**9**	12	32	39	53
	Put 1	Put 2	Put 3	Put 4	Put 5
Time	106	106	119	109	72

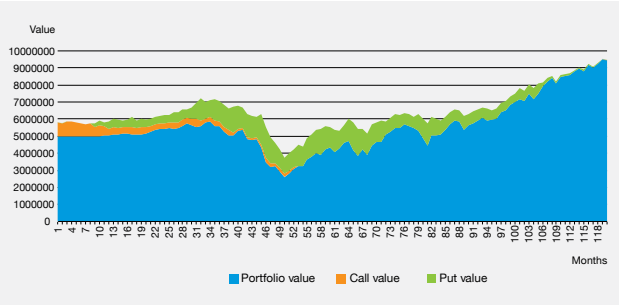

Figure 3a – Value of each component

	PE fund		
	Mean	**Min**	**Max**
PE fund risk – standard deviation (returns)	14.12%	12.25%	15.60%
S&P risk – standard deviation (returns)	14.62%	14.62%	14.62%
Correlation (PE fund, S&P)	0.69	0.48	0.75
Ratio S&P risk/PF fund risk	1.04	1.19	0.94

Table 2 – With leverage average volatilities

ADDING REALISM TO THE ANALYSIS

A knowledgeable PE investor might scoff at our simple model above. After all, as noted above, what happens to the money once any cash out occurs could bias the risk, albeit only at the tail in the example. Clearly it matters when funds are returned to the investor. Obviously there would be leverage. Typically, the reason to invest with a PE fund is to take advantage of the manager's knowledge to choose a segment of the market in which to invest, determine what company might be undervalued in that segment, and to make efficient use of leverage. To address some of these challenges, we now assume the following:

1. The fund is leveraged by 50%. The borrowing rate is 2% above Libor, and the borrowing occurs at the time of investment.
2. Once the PE fund returns 150% of the investment it is cashed out.
3. Any cash in this example is now invested in the S&P, which likely biases the correlation up and increases the volatility.

With these assumptions, the PE fund's performance versus the S&P is presented in Figures 3 and 3a. The volatility of the PE fund is higher due to leverage.[11] Summary statistics are provided in Table 2. In this case, the volatility of the PE fund is higher, although it is similar to that of the unleveraged S&P.

From Figures 3 and 3a, which have the same call dates as Figures 1 and 1a, we see a different alignment and can easily observe the volatility due to leverage affecting the put values. The call values are not impacted since the leverage kicks in when the call is made and shows up in the put. Leverage also shows up in the return pattern influenced by the S&P's movement because the size of the investment has now increased by the leverage factor.

We can compare the ratios in Tables 1 and 2. Volatilities are higher for the PE fund as would be expected due to leverage. However, the reinvestment into the S&P does not alter the correlations as much as one might expect, ex-ante, given that for much of the period the bulk of the PE fund takes exactly the

11 Clearly, these paths depend on the various assumptions we have made. In future work we hope to use actual cashflows of a PE fund to determine whether the theoretical conclusions are indeed robust.

same pattern as the S&P, due to the reinvestment into the S&P as investments are taken to market. If we were comparing for performance we believe the returns would differ, probably in favor of the PE fund, but for our purposes this is not relevant.

Given the correlation and volatilities suggested by this new methodology, we can argue that an optimal portfolio should have a considerable amount allocated to the PE fund.[12] Even with the restrictions imposed, PE allocations clearly provide ample diversification. Given that many institutional investors and foundations have PE allocations of 5% or less, it is highly likely that, if our analysis is correct, most of them are sorely underallocated to PE [Ennis and Sebastian (2005) and CFA Digest (2005)].

ASSUMING SUBJECTIVE INVESTMENT DECISIONS

We have taken care not to bias our results significantly in favor of the PE fund's lower volatility. In reality, investors move assets into a PE fund if they believe in the manager's superior ability in some way. Superiority could take the form of asset selection, financial expertise and knowledge of capital structures, or ability to influence management behaviors at the level of portfolio companies. Again, we choose to simplify by giving to PE managers the ability to make investments when they observe an opportunity and decide to exercise the option to call investors' funds. Our intention is to highlight the call and put option elements that we believe are fundamental to a PE fund. To do so, we move to a different 10-year period, commencing in January 2007.

More of the calls now occur during the crisis (call dates are again held the same at month 9, 12, etc.). The impact on the volatility measure shows up differently. The timing on the cash out for the investments – the puts – will differ accordingly (the 150% return is met at very different times than in the earlier analysis given the starting points of the investment).

Again, due to the nature of the investment, the values end up in similar positions. Since calls are partially made during, rather than before, the crisis, it is interesting to see how the put values adjust and how quickly a cash out occurs.

The most interesting results, however, show up in the ratios of risk for the PE fund and the S&P. The reduction in risk relative to the S&P is apparent in Table 3. In all cases, the PE fund is much less volatile, even with 50% leverage. Moreover, the correlations are extremely low, but the difference here from Table 2 is that the general methodology favors a higher correlation,

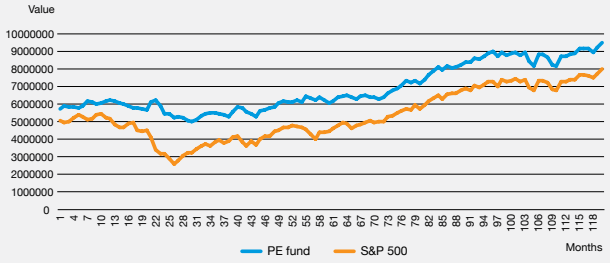

Figure 4 – Path of $5 million investment

	Call 1	Call 2	Call 3	Call 4	Call 5
Time	9	12	32	39	53
	Put 1	Put 2	Put 3	Put 4	Put 5
Time	119	100	71	79	87

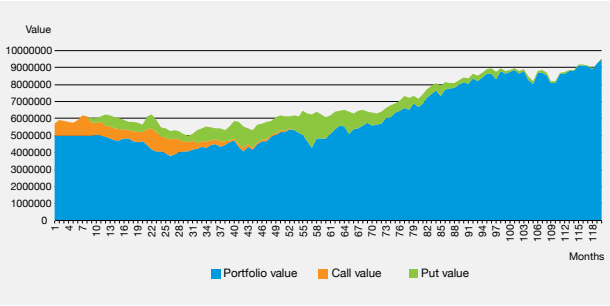

Figure 3a – Value of each component

	PE fund		
	Mean	**Min**	**Max**
PE fund risk – standard deviation (returns)	12.11%	11.30%	13.39%
S&P risk – standard deviation (returns)	15.24%	15.24%	15.24%
Correlation (PE fund, S&P)	0.37	0.09	0.50
Ratio S&P risk/PF fund risk	1.26	1.35	1.14

Table 3 – Actual crisis with leverage average volatilities

12 Our approach here would violate the theoretically correct way of determining an optimal portfolio [see Wilford (2012)]. However, our point is not to show a suggested level of allocation to PE, but rather to establish that for most large investors it would be significantly higher than allocations based on current understanding of risk-adjusted returns from PE.

as noted above. Still, the mean correlation is 37% and the minimum path is only 9%.

Freeman and Wilford (2016) noted that many believe that crises are the life-blood of the PE firm (fund). This view is supported by Table 3, not only in the much lower volatilities of the leveraged PE fund, but also in the very low correlation measures.

REALISM AND MARGINAL RISK REDUCTION – "NEGATIVE BETA"

The set of options that exist for the PE fund manager are multiple and can become complex. As one adds realism to the process, more interesting aspects of decision-making can be considered. What type of investment is made and when (do we need a crisis to buy optimally, for example)? When does a PE firm take any particular investment to market? And for the PE fund investor, what is the marginal impact of any PE fund investment on the portfolio? Above, we have made simplifying assumptions and as the research progresses some of these can be modified to consider the complications that better reflect the reality of a PE fund.

EXTENDING THE LIFE OF A FUND – VALUING THE PUT

Consider one of the important elements in risk reduction available to the PE Fund/GP (management firm). Critical to the value of the investment is the ability of the PE fund to delay taking an investment to the market (we may call this action the "extension right"). To see the value of this option, which may be exercised due to a crisis or simply due to conditions related to the market's desire for the type of investment that is made by the fund, we create an exercise where there is a crisis that extends the implied value of the put. In this instance, the PE fund delays taking its remaining investment/s to market. In order to model this, we modify our method in several ways, while acknowledging again that this is a necessarily stylized version of reality. Here, the "put" option is in fact much closer to a set of real options than to a pure financial option.

1. Using the basic model a crisis is created in the data, allowing the S&P to fall. An artificial crisis in 2012 is assumed by allowing another 2007-type downturn, and then the dataset is extended from that now artificial date onward. All of the calls have happened as above, however instead of the PE fund being forced to close out straight away the GP can choose to extend the period for three years.
2. In this artificial path only one of the five investments has been returned to investors and cashed out (here again all proceeds are invested into the S&P as in Figure 4 once this occurs).
3. After the crisis occurs, the S&P is allowed to recover as it did in reality, however the dates are artificial, pushing out to 2019 (thus 2019 data are actually 2016, etc.).
4. The PE fund manager is able to extend the holding period of the investments that remain by exercising the fund's right to extend the investment period by three years after witnessing four successive months of negative returns while also recognizing that the normal life of the fund is about to end.

Extension increases the life of the put option, which creates a new path for the overall risk measure and valuation. Now the question becomes how valuable is this "crisis option"? From our modeling, we can see that the ability to extend the life of the put was very valuable. In reality, most PE firms will have

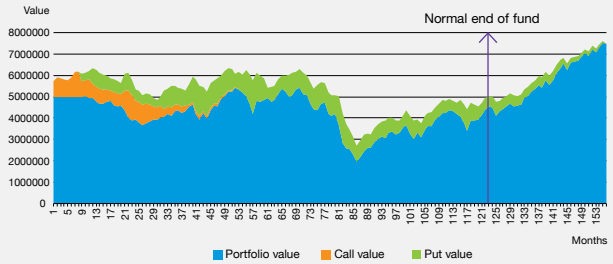

Figure 5 – Value of each component

	Call 1	Call 2	Call 3	Call 4	Call 5
Time	9	12	32	39	53
	Put 1	Put 2	Put 3	Put 4	Put 5
Time	156	156	69	145	156

	PE fund		
	Mean	**Min**	**Max**
PE fund risk – standard deviation (returns)	15.08%	14.14%	17.18%
S&P risk – standard deviation (returns)	17.13%	17.13%	17.13%
Correlation (PE fund, S&P)	0.62	0.39	0.73
Ratio S&P risk/PF fund risk	1.14	1.21	1.00

Table 4 – Simulated crisis average volatilities

the right to extend the life of a particular fund by one year in the first instance – this is typically written into the governing documents of the fund and may or may not be contingent upon a vote of approval from the fund's advisory board. Equally significant for our analysis, extensions can often be made in yearly increments for up to three years and depending on the fund there might even be an option to ask the Advisory Board for a further exceptional additional extension in circumstances where the interests of the manager and investors are clearly aligned. No one wants to a be a forced seller.[13] Many GPs will reduce any remaining management fees at this point in order to cement that alignment.

To compare the value of the put option we assume that the PE fund manager goes to market at whatever price is available at the end of the period of recovery. If this extension option had not existed the fund would have been extinguished at the end of year 10 with losses; this equates to the line drawn after 120 months on the figure. The implied losses relative to the return that was accepted after the extension is some $3.5 million or about half of what was returned to the client. The put option's value would be extinguished and its value reduced to zero, forcing the firm to take the loss not only in opportunity but in actuality as well. Of the original $5 million invested, $1.5 million or some 30% of the original investment would have been lost. Dates of when the investments are actually taken to the market are noted under Figure 4, with the extension during the artificial crisis. As such, the puts extinguish near to or at the end of the extra years, not during a crisis.

The value of the option to extend is very clear from our results as presented in Figure 4 (our focus here is upon the "put" area in particular).

In Table 4, the data are once again presented for comparison. The volatility ratios are again greater than one, suggesting the PE fund was able to enjoy lower volatility in this circumstance and, just as importantly, the correlations of returns are low.

In summary, our initial assumption that capital must be returned at year 10 implies in our modeling that the PE fund would be forced to sell into a down market. With the ability to extend the life of the fund, the manager is given the freedom to continue to wait for better conditions to take the investment to market. This right makes the fund much more valuable during a crisis. The risk mitigation during a crisis is expressed in large part by the value of that option to wait, providing a risk profile for the PE fund that is significantly less volatile than that of the market as a whole, but meanwhile the fund is providing the returns sought by the investor.

CONCLUSION

A significant conclusion of our work is that most investment funds are systematically underexposed to PE if standard risk parameters are utilized during asset allocation decisions. This holds in normal times, when well-managed PE should be expected to offer outperformance even after fees. However, the real justification is found in bad times, when the "negative beta" effect of PE assets is to mitigate the effects of any downturn in public markets. Depending on the specific PE assets, this mitigation can be significant.

- Private equity firms draw down their investor commitments via capital calls.
- The right to call money at any time, including during market distress, is a powerful option,
- PE firms can also delay selling assets during a crisis, which is a valuable, if complex, put option.
- This combination of options makes PE investments perform with powerful "negative beta" during bad times.
- Institutional investors, particularly those with long time horizons such as foundations, have traditionally focused mainly on PE fees and may, therefore, be systematically underexposed to this asset class

REFERENCES

- Bain, 2016, "Global private equity report 2016," Bain & Company, January, http://bit.ly/21hj9dz
- Black, F., and M. Scholes, 1973, "The pricing of options and corporate liabilities," Journal of Political Economy 81:3, 637–654
- CFA Digest, 2005
- Chen, A. J. Conover, and J. Kensinger, 2008, "Private equity arrangements as real options," Midwest Finance Association 2012 Annual Meetings Paper
- Considine, G., 2010, "Lessons from Yale's endowment model and the financial crisis," Advisor Perspectives, April 20, http://bit.ly/2lZQIGZ
- Ennis, R. M., and M. D. Sebastian, 2005, "Asset allocation with private equity," Journal of Private Equity 8:3, 81–87
- Freeman, A., and D. S. Wilford, 2016, "Private equity capital commitments: an options-theoretic risk management approach," Journal of Financial Transformation 43, 106-116
- French, K., W. Schwert, and R. Stambaugh, 1987, "Expected returns and volatility," Journal of Financial Economics, 19:1, 3-29
- Kaplan, S. N., and B. A. Sensoy, 2015, "Private equity performance: a survey," Annual Review of Financial Economics 7, 597-614
- Kaplan, S. N. and A. Schoar, 2005, "Private equity performance: returns, persistence, and capital flows," Journal of Finance 60, 1791-1823

13 A logical extension of this is for the PE manager to consider making certain funds into quasi-permanent vehicles where the nature of the underlying assets means that pay-offs are bond-like or of very long duration, infrastructure being the obvious example.

- Mathonet, P-Y., and T. Meyer, 2007, J-curve exposure: managing a portfolio of venture capital and private equity funds, Wiley
- Merton, R. C., 1973, "Theory of rational option pricing," The Bell Journal of Economics and Management Science 4:1, 141-183
- Meyer, T., 2014, Private equity unchained: strategy insights for the institutional investor, Palgrave-Macmillan
- Provence, L., 2008, "Missing billion: how UVA's investment strategy worked... until it didn't," The Hook, November 20, http://bit.ly/2l6kbQh
- SEC, 2015, "Blackstone charged with disclosure failures," Securities and Exchange Commission press release no. 2015-235, http://bit.ly/1L06uGk
- Sorensen, M., N. Wang, and J. Yang, 2014, "Valuing private equity," Review of Financial Studies 27:7, 1977-2021
- Wilford, D. S., 2012, "True Markowitz or assumptions we break and why it matters," Review of Financial Economics 21, 93 – 101

Layout, production and coordination:
Cypres – Daniel Brandt, Kris Van de Vijver and Pieter Vereertbrugghen

BANGALORE
BRATISLAVA
BRUSSELS
CHICAGO
DALLAS
DÜSSELDORF
EDINBURGH
FRANKFURT
GENEVA
HONG KONG
HOUSTON
KUALA LUMPUR
LONDON
NEW YORK
ORLANDO
PARIS
SINGAPORE
TORONTO
VIENNA
WASHINGTON D.C.
ZURICH

CAPCO.COM