BVCA
British
Private Equity &
Venture Capital
Association

Private Equity
Performance Measurement
BVCA Perspectives Series

Authored by the BVCA's Limited Partner Committee
and Investor Relations Advisory Group

Introduction

Private equity is an asset class with the potential to generate sustained, long-term outperformance for its investors, and is a key component of many investors' portfolios. However, it is still relatively immature as an asset class, with some unique characteristics that may not be familiar to investors. One of its defining features is the irregular timing of cash flows from investors to fund managers and vice versa. Once an investor has made a commitment to a fund, it may not be called upon for a period of months or even years, and when it is called, this may be at irregular intervals and sizes. As a consequence, returns are typically measured in a different manner than other asset classes, and it is not necessarily a straightforward exercise for an investor to benchmark their private equity holdings against other asset classes.

As a result, a number of different metrics have been adopted to give investors the greatest possible understanding of the performance of their private equity and venture capital investments and benchmark these against other types of asset.

This guide seeks to give a brief overview of the most common methods, highlighting where these metrics find their greatest use and giving some indication of the limitations involved with each.

It is important to stress that no single measurement represents the right or wrong way of measuring the performance of private equity and venture capital investments. Individual investors and fund managers will find that different combinations of these metrics will work best for them in assessing their private equity and venture capital investments and that they are best placed to take a view on this.

Table 1.1: Differing performance measures and the information they require

Degree of information	Absolute return measures			Relative performance measures	
	IRR	MIRR	Multiple	PME	PME+
Amount of PE cash flows	✓	✓	✓	✓	✓
Timing of PE cash flows	✓	✓		✓	✓
Assumption for reinvestments		✓		✓	✓
Index values				✓	✓

Source: Montana Capital Partners (2011)

Internal Rate of Return (IRR)

The internal rate of return (IRR) is a metric used to measure and compare returns on an investment. It calculates the return by looking at all of the cash flows from the investment over a given period, taking into account drawdowns, distributions such as capital gains and income through dividends, and a valuation if the fund still has residual value.

IRRs are often used by the PE industry to measure returns (and, indeed, are one of the metrics of performance that we publish in our annual Performance Measurement Survey), because they offer a means of comparing two investments with irregular timings and size of cash flows. They are, however, a measure that is not directly comparable to the buy and hold returns that can be found in the public markets.

When analysing IRRs, it's important to understand whether the returns are gross or net of fees and carried interest. There can be significant differences between the two metrics, particularly when there has been a long hold period (management fee drag) or when a portfolio has generated a disproportionately high and early "win". For these reasons, LPs tend to focus on net IRRs to get the most transparent picture of an actual return. However, many GPs use gross IRRs in marketing materials as they present the most compelling return.

Uses and strengths

- The IRR allows investments with irregular cash flows - one of the defining features of private equity and venture capital funds – to be analysed. In doing so, it offers a way to compare and rank different investments, clearly showing which investments offer the best rates of return.

- The IRR considers the time value of money. For example, having £100 today is (generally) worth more than an expected £100 in a year, and the IRR accounts for this in its calculations.

- Relatively easy to calculate in a standard fashion with the assistance of a computer, and can be straightforward to interpret.

Limitations

- The IRR assumes that cash flows are reinvested at the same rate of return. This can lead to the over or understatement of the performance of a given investment where the returns on reinvestment do not match those produced by the investment, and should be accounted for when comparing IRRs.

- The IRR is not an effective way of assessing mutually exclusive projects, as it does not take into account the scale of the projects – this can be difficult when two projects require a significantly different amount of capital, but the smaller project has a higher IRR.

- If returns are looked at only on an IRR basis, then there is the potential for performance to be artificially improved by changing the timings of distributions back to investors. Early wins (quick returns of significant amounts early in the life of the fund or investment) can disproportionately boost the IRR.

- Since the IRR represents the discount rate at which the value of all cash flows equals zero, it is possible that multiple IRRs, or no IRR at all, can be calculated in some cases.

- In order to compare the IRR to a time-weighted return, backward-looking IRRs over a given time horizon are often used. Caution should be used however, as when a given point in time (and therefore an attendant valuation) is taken as a start date, this can then give a misleading perspective on the performance of a fund. The example[1] below illustrates this – suggesting that on various measures (three, five and ten year returns), a fund can look like it is doing reasonably or exceptionally well. However, on a since inception basis, the fund performed poorly throughout its life.

- Another pitfall associated with "IRR snapshots" is that they are only accurate in gross form as it's virtually impossible to disaggregate management fees for a specific time period. To that end, LPs often use assumptions to appropriate a "net proxy" for such data cuts to better understand the returns from an LP's perspective.

Figure 1: An example of since-inception and backward-looking returns
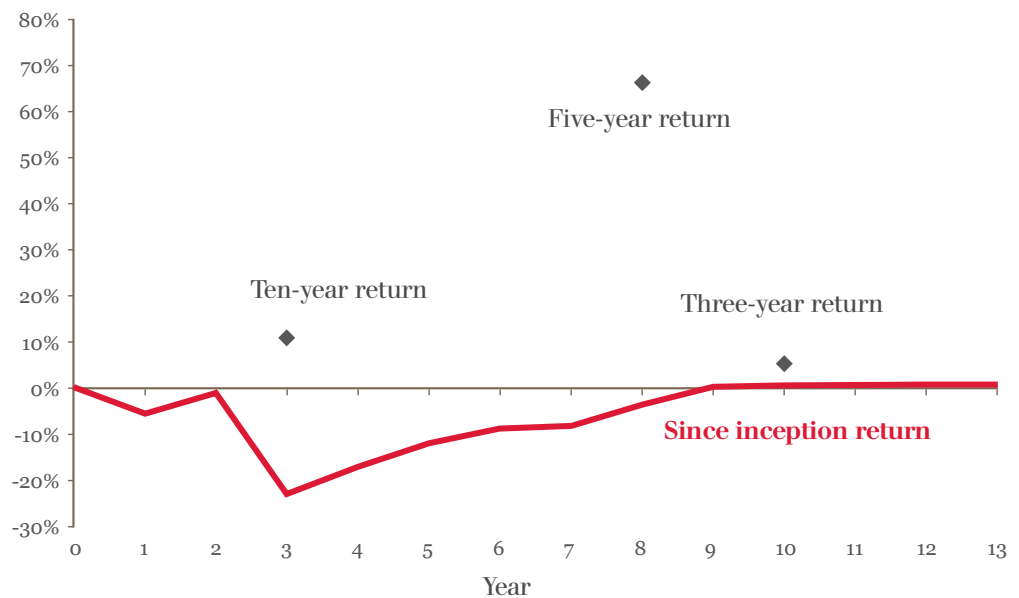
Modified Internal Rate of Return (MIRR)

It is worth noting in the context of the IRR that a metric known as the Modified Internal Rate of Return (MIRR) has been developed to tackle some of the problems faced by the IRR - namely that it no longer assumes that cash flows can be reinvested at the same rate. Rather, the MIRR assumes a reinvestment rate, separate from the rate of return of the project, at which cash flows can be reinvested. It then calculates the rate of return by looking at all of the cash flows of the project, and accounting for the time value of money. As such, the MIRR is designed to better take into account what is done with cash flows once they are received and better reflects the profitability of an investment.

The MIRR is a theoretically strong method for calculating the overall returns from a private equity investment, though it is not widely used in practice. As a result, we have not covered it in depth in this paper, although further information on its uses and limitations have been included in the Appendix.

[1] This example is based on a real-life fund drawn from the BVCA's Performance Measurement data set.

Multiples of Invested Capital

Money multiples are another metric that measure returns from an investment, providing a cash-on-cash measure of how much investors are receiving. They are calculated by dividing the value of the returns by the amount of money invested. Two multiples that are typically reported by funds are distribution to paid-in capital (DPI) and total value to paid-in capital (TVPI), which differ in terms of whether or not they include residual values.

Multiples are often used in the PE industry as they offer an easy way to show the scale of the returns an investment has given. While the IRR and MIRR provide useful ways of showing returns from an investment, they cannot provide a scale of returns without knowing the length of the fund. Multiples, however, offer a fast and easy way to show this, and when used in conjunction with the IRR or MIRR, can paint a quick and clear picture of a fund's performance.

It is convention in the industry to show TVPI and DPI on a net basis, but often a GP's return profile will have a gross multiple i.e. money in – money out as a headline figure in any performance table. As with IRR and other performance metrics, investors should exercise care when understanding whether a presented metric is gross or net.

Distributions to Paid in Capital (DPI)

■   DPI measures the ratio of money distributed by a fund against the total amount of money paid into the fund.

■   At the start of the investment, this ratio will be zero, and will begin to increase as distributions are paid out. When the DPI is equal to one, the fund has broken even, as money paid in is equal to money distributed, and any number above this indicates that the fund has paid out more than has been paid in.

Total Value to Paid-in Capital (TVPI)

■   TVPI measures the overall performance of a PE fund with a ratio of the fund's cumulative distributions and residual value to the paid-in capital. It calculates what multiple of the investment would be returned to investors if the unrealised assets were sold at current valuations and added to distributions that had already been received.

■   Unlike DPI, TVPI includes the residual value in its calculation. While this can suggest a more complete picture on the returns from the fund, it is significantly impacted by the valuation that is placed on the unrealised investments remaining in the fund (although the impact should reduce as the fund matures and investments are realised).

■   While DPI provides a clear metric of the actual multiple of cash invested which has been received by an investor, TVPI provides a metric that accounts for potential returns that are the result of increased valuations of portfolio companies as they approach exit. Given this difference, many LPs rely on TVPI earlier in the life of a fund and DPI towards the end.

Uses and strengths

■   Multiples offer a quick and readily digestible means of indicating the performance of a fund. Any multiple above 1 shows that that the fund has returned more than its initial investment.

■   They also offer a way for projects of different scales to be compared. As returns are expressed as a proportion rather than as an absolute figure, investments of different sizes can be easily compared.

Limitations

■ Multiples ignore the time value of money. For example, a 2x multiple tells investors that for every one GBP invested into a fund, they received back twice as much in return. However, the relative attractiveness of this investment would be markedly different if it had taken 10 years to produce that return than if it had taken two.

■ The metric is also less useful when comparing investments in their early stages. When draw downs are still being made, Multiples can be highly variable, reducing their usefulness. As such, they are most relevant as a fund matures.

■ While multiples are good at showing returns as a proportion, they do not provide an indication of the scale of the project nor of the size of the absolute returns. This means their use is limited when comparing mutually exclusive projects.

■ It is also possible that when looking at multiples on a fund basis, one large deal can have a disproportionately large positive or negative impact on a blended basis.

Public Market Equivalent (PME)

Because it is not practical to create time weighted returns based on private equity investments, comparisons with time weighted returns of the public market have the potential to be misleading – an IRR and a time weighted return are not directly comparable. To attempt to address this and provide an accurate comparison with the market, the Public Market Equivalent was developed.

The Public Market Equivalent (PME) is a returns measure which allows investors to compare an IRR to the performance the public market would have generated over the same period using the same investment timings. The PME is generated by creating a hypothetical investment vehicle which purchases and sells shares in the public market index in such a way that mimics a PE fund's irregular cash flows. As discussed earlier, net cash flows should be used to provide the most transparent return profile for the underlying fund(s). On a related point, industry convention generally dictates the use of indices with reinvested dividend income to provide the most "apples to apples" comparison.

Essentially, the metric looks at all the cash flows for a given private equity fund. It then buys an equivalent value on the market when a draw down from an investor is made (for instance, if a draw down of £50m was made by the private equity fund, then £50m would be hypothetically invested into the market. Similarly, if £50m was distributed back to investors from the private equity fund, then £50m would be sold from the hypothetical market vehicle). The two investments (private equity and market) will therefore have the same cash flows over time, but different end values, giving two directly comparable IRRs.

It is used to give a like-for-like comparison between private equity and the markets, and is a means by which private equity's irregular cash flows can be accounted for.

A challenge with using the PME (and also PME+ and other derivatives of index-based performance measurement) is deciding which index is most relevant to the private equity strategy. This is a subjective area and beyond the scope of this paper but it's important to flag as a relative comparison between private equity cash flows and an inappropriate index (i.e. developed market returns vs an emerging market index or equity returns vs a credit index) can provide misleading results.

Uses and strengths

■ Both the IRR and Multiples are unsuitable for comparing private equity to typical public market investments, which tend to give long-term and regular returns. PME, on the other hand, works by creating an investment vehicle that takes into account a PE fund's irregular cash flows.

■ PME therefore allows public market investments to have their performance compared to private investments that could have taken place over the same time period.

■ The method is particularly accurate when it is used to look at more mature investments where the NAV is a smaller fraction of the distributions, as any errors in calculating the NAV will have a smaller impact on the final result.

Limitations

■ It is possible that the public market NAV could become negative in cases where the PE portfolio greatly outperforms the benchmark. This can potentially result in a largely nonsensical comparison of the performance of a long-only PE portfolio being compared against a short position in the public market.

■ PME also faces the same issues as IRR, as the calculations are effectively identical. As such, this method does not account for the scale of the project and can encounter issues with abnormal cash flows.

■ There is also an argument that the PME is not a true comparison, as it is reliant on the timings of the PE fund. This means that the PME vehicle may be forced to buy or sell the market at times that would not ordinarily be chosen, adversely affecting the performance of this vehicle and overstating the performance of PE relative to this vehicle.

■ PME (and indeed PME+ discussed below) does not attempt to specifically adjust for risk differentials between the portfolio cash flows and the public index (i.e. higher leverage ratios on private equity companies) or the tax impact of returns.

Public Market Equivalent + (PME+)

PME+ is a methodology developed to address the problem of short exposure that can be present in a PME calculation[2]. It does this by fixing the closing NAV of the benchmarking vehicle to be the same as the closing NAV of the private equity fund. This is achieved by scaling distributions until the final valuation of the public market index is equivalent to the private equity NAV. By selling a fixed proportion of the corresponding private equity cash flows instead of an equal amount, PME+ ensures a positive end balance whilst preserving the overall cash flow pattern. However, it does not guarantee that the exposure stays positive over the entire investment period.

A scaling factor that is smaller than one shows that the public market has outperformed the private market.

[2] If the private equity investment significantly outperforms the market, it is possible under a standard PME metric that the hypothetical market vehicle will have to sell more shares than it has available to match a private equity distribution to investors. As a result, the market vehicle goes short.

Uses and strengths

■ PME+ shares some of the same advantages as PME, namely, that it allows public market investments to have their performance compared to the private market.

■ The method further builds on PME by overcoming the problem of short exposure, as outlined above. By ensuring a positive final NAV, the results have greater validity when comparing performance.

Disadvantages

■ Similarly, PME+ shares a number of its disadvantages with PME and IRR. The method relies on the benchmark index for calculating the NAV of the hypothetical fund. For this purpose only total return indexes should be used, as measures that ignore dividends render the index invalid and can result in distortions. Likewise, PME+ does not account for the scale of the project and can encounter issues with abnormal cash flows.

■ Due to the introduction of scaling, PME+ does not exactly replicate PE cash flows in the public market. Since cash flows are being moved, this can have a consequent effect on the comparative results.

Summary

It is important to note that none of the methods described above can in isolation give a complete picture of PE performance, nor is there a single 'right' way to measure it. One of the challenges associated with this is the fact that managers can claim to be "top quartile" by creatively choosing benchmarks that position themselves favourably. An LP needs to fully understand the constituents of any benchmark that a manager is using in order to fully assess the relevance to the manager's strategy. In addition, as there is no industry standard to benchmarking, there are potentially significant biases in the available benchmark data.

There are also many questions around the detail behind the performance metrics. For example, whether cash-on-cash multiples are more meaningful than IRRs, how to gauge impressive performance as a result of unrealised versus realised results, the impact of the "J-curve" and the gross to net performance spread.

The issues described above are potentially significant and are the subject of further commentary outside the scope of this paper. Suffice it to say that while the performance measurement techniques described above are a useful tool for assessing PE fund performance, individual investors need to decide which metrics are most appropriate for each situation. In addition, investors generally go well beyond simply comparing performance statistics and often employ detailed analyses such as value creation attribution, capital efficiency including the gross-to-net ratio, consistency over multiple time periods and individual analyses focusing on sectors, geographies, sizes, types and key professionals.

Appendix

IRR and Multiples: Worked example

The example set out below illustrates how the IRR generated by a fund is influenced by both the timing and scale of distributions, even when the total returns from the different funds are the same. The three funds below have different cash flows, but the same total returns and distributions. All three funds have the same initial outlay of £100 in the first period, and have the same total distributions of £200 over the lifetime of the fund, and the same final valuation of £100. Fund A distributes no returns to investors until 2008, when they receive £50. The fund then distributes £50 the year after and £100 the following year. This gives an IRR of 11.4%. Fund B, however, only has to wait until 2003 to distribute its first returns, and after 2006 makes no more distributions during the period looked at here. The result is an IRR of 17.4%. Fund C however distributes all of its returns by 2004. This gives the greatest IRR of 26.4%.

The bottom section of the table shows the multiples calculated for the same cash flows and valuations. It demonstrates how multiples do not account for the time value of money, with all three funds showing the same result on DPI and TVPI basis . This is in contrast to the IRR and MIRR results, which take into account the time value of money and therefore result in varying returns depending on the timings of the cash flows.

Year	Fund A	Fund B	Fund C
01/01/2000	-£100	-£100	-£100
02/10/2001	£0	£0	£30
27/02/2002	£0	£0	£30
05/06/2003	£0	£20	£40
17/04/2004	£0	£30	£100
25/03/2005	£0	£50	£0
20/08/2006	£0	£100	£0
10/10/2007	£0	£0	£0
23/11/2008	£50	£0	£0
08/03/2009	£50	£0	£0
15/09/2010	£100	£0	£0
31/12/2010*	£100	£100	£100

*Values on this date are Net Asset Valuations

IRR	11.4%	17.4%	26.4%
DPI	2	2	2
TVPI	3	3	3

Modified IRR

Below is some further information on the Modified IRR, looking at its uses and strengths, as well as its limitations. A worked example is also included.

Uses and strengths

■ The MIRR shares many of the same advantages of the IRR. It provides a method to universally compare projects of different sizes, allows for irregular cash flows and accounts for the time value of money.

■ The MIRR offers improvement on one of the IRR method's major shortcomings. Using the IRR can give an overly optimistic picture of returns where the rate of return on any reinvestments does not match that offered by the initial investment. The MIRR tackles this reinvestment assumption and does not assume that the returns from the investment can be reinvested at the same rate of return. Rather, the methodology uses an assumed reinvestment rate which is usually lower than the return of the investment – as may be the case in real life situations.

■ The MIRR does not suffer the same issues as the IRR in terms of alternating positive and negative cash flows. While the IRR algorithm can fail to find a solution in some situations, the MIRR approach will always give one answer.

Limitations

■ Using the MIRR requires a reinvestment rate to be estimated. Because the reinvestment rate has a significant impact on the overall return under this method, questions arise over whether it accurately reflects returns or whether it is too influenced by the reinvestment assumption. Furthermore, the MIRR assumes that the fund continues long after the final cash flow distribution. This means that the final return estimate is biased towards the investment rate which is assumed in the calculation. Relying on this type of estimate can adversely impact the accuracy of the final result and result in errors.

■ Like the IRR, the MIRR can make a smaller project seem more attractive and ignore the fact that the larger project will generate significantly higher cash flows and a greater total return. Similarly, the MIRR can rank a short term project with high returns over a long term project with lower returns, ignoring the fact that the larger project can generate significantly higher cash flows.

Worked example

The worked example below compares the same investment on an IRR and MIRR basis. For the MIRR, a reinvestment rate of 10% has been assumed. Because this is lower than the implied rate under the IRR, the MIRR produces a lower return in this situation. It would equally be possible to assume a higher reinvestment rate and produce a commensurately higher MIRR.

Date	Cashflow	Investment Returns	Cumulative Returns
01/01/2010	-50	0	0
01/01/2011	0	0	0
01/01/2012	10	0	10
01/01/2013	20	1	31
01/01/2014	30	3.1	64.1
01/01/2015	40	6.41	110.51
01/01/2016	50	11.051	171.561
IRR=	27.9%		
MIRR=	22.8%		