PitchBook

PitchBook Data, Inc.

Nizar Tarhuni Executive Vice President of Research and Market Intelligence

Dan Cook, CFA Global Head of Quantitative Research & Market Intelligence

Institutional Research Group

Analysis

Hilary Wiek, CFA, CAIA
Senior Strategist
hilary.wiek@pitchbook.com

Zane Carmean, CFA, CAIA
Director, Quantitative Research
zane.carmean@pitchbook.com

pbinstitutionalresearch@pitchbook.com

Publishing

Designed by **Chloe Ladwig**

Published on January 30, 2025

Contents

Benchmarking and Returns: Why Are There So Many Numbers?

The ins and outs of investment evaluation metrics

PitchBook is a Morningstar company providing the most comprehensive, most accurate, and hard-to-find data for professionals doing business in the private markets.

Key takeaways

- Measuring investment performance is not as straightforward as it seems. Understanding which figures to use—and not to use—in which situations is important to appropriately measure the effectiveness of any investment.
- IRRs are most appropriate for investment pools with multiple cash flows over extended periods, which is why it is the generally accepted metric for private equity funds. They do not, however, measure the full experience of a capital commitment, as in a drawdown fund, the full allocation is not invested fully throughout the life of the fund.
- Interim IRRs and multiples are highly dependent upon the valuations of illiquid assets. For this reason, LPs should critically evaluate the valuation policies and their implementation during the life of a fund.
- There are advantages and drawbacks to every investment measurement discussed in this report. Gaining an understanding of these will aid readers in properly evaluating investments in a wide variety of scenarios.

Introduction

In our prior piece, Numbers Mean Nothing Without Benchmarking, we covered a lot of ground when it came to the rationale behind benchmarking, important characteristics of benchmarks, and how to think through which benchmark is appropriate to each situation in the investment experience.

This report continues the conversation, getting into more technical questions like what the different return metrics are, how to calculate them, when to use which ones, and why the different returns should be used only for specific scenarios. Some investors may be vaguely aware that time-weighted returns (TWRs) are not comparable to internal rates of return (IRRs) but may have no idea why. Others may not even have that vague awareness. Some are big fans of multiples, but they may not realize the drawbacks of these and other metrics. Others may be aware of public market equivalents (PMEs) but do not know what the numbers are, what they mean, or how to use them appropriately. This report provides insights into all of these questions.

Some investors may be vaguely aware that time-weighted returns (TWRs) are not comparable to internal rates of return (IRRs) but may have no idea why.

The metrics we use

In this section, we attempt to keep things as nontechnical as possible, but we do go into the equations as well as the rationales and drawbacks for using each in different situations. The first building block is the holding period return, which is how you would calculate the percentage change of anything from an egg price increase to a temperature change, but we state our equations in investment terms.

$$Holding\ period\ return\ (HPR) = \frac{End\ value - Initial\ value}{Initial\ value}$$

This equation most reflects reality if you make a one-time capital investment into a fund, or really any investment opportunity, and then measure what it was worth at any later date. If more money is added or some is extracted from this fund during the holding period, however, then the equation needs an adjustment so that the ending value is a reflection not only of the final valuation but also of any capital that flowed between the investor and the fund.

For a numerical explanation of this, we show three scenarios. In the first, an investment was made, and then after one year, the full investment was cashed out. In the second and third, the same amount of capital was withdrawn from the fund after one year, but more capital was added to the fund by the investor partway through the year.

In the first scenario, we have clearly made a 20% return. However, in the second and third scenarios, if an adjustment is not made for the additional capital injected after the initial investment, the HPR would not reflect reality. Clearly in the second scenario no money has been made, so an HPR of 20% would be incorrect. In the third scenario, less money has been made than in the first scenario, so the true return should be less than a simple HPR would provide.

Holding period return scenarios (Hypothetical data)

	Scenario 1	Scenario 2	Scenario 3
January 1, 2020	$100,000 investment	$100,000 investment	$100,000 investment
June 30, 2020		$20,000 added	$10,000 added
January 1, 2021	-$120,000 withdrawal	-$120,000 withdrawal	-$120,000 withdrawal
Cash profit	$20,000	$0	$10,000

Source: PitchBook

To adjust the returns for cash flows, there are two equations named after Peter Dietz called the Dietz method and the modified Dietz method. The first takes the holding period return, subtracts the net interim cash flows from the numerator (a contribution is a positive cash flow, while a distribution is a negative cash flow), and then adds the contribution to the denominator and divides it by two to account for the fact that the cash flows happened at a time other than the beginning of the period. This looks like:

$$Simple\ Dietz\ method[1]: \frac{End\ value - Initial\ value - Net\ inflow}{Initial\ value + \frac{Net\ inflow}{2}}$$

Using this equation, the second scenario from the corresponding table would provide a Dietz return of 0%, which feels undeniably correct given the $0 cash profit on this investment. For the third scenario, the Dietz method provides a return of 9.52%, which is curious at first glance. We know we put in $110,000 and got back $120,000 for a $10,000 profit. That feels like it should be roughly a 9.1% return. The adjustment to the denominator—dividing by two—assumes that the additional capital came in midway through the period—which is, in fact, accurate in this case, but in most real-life situations will not be. The equation thus assumes that the original $100,000 was held for six months, providing some return, and $110,000 was at work for six months, providing the rest of the return. If the return occurred in a straight line upward, then we should end up with a higher return than if the full, higher amount was invested for the full year, which aligns with the result this equation provided. Put another way, a $10,000 profit on a $100,000 investment should provide a better return result than a $10,000 profit on $110,000 invested.

The Dietz equation worked out in this simple scenario, but as time periods extend and more cash flows occur and at different times, the results become less intuitive and other methods are required to account for the amount of assets in the ground (money invested) when each return regime occurred. The modified Dietz method does handle additional cash flows better, but the formula is complicated and, as we will see, IRR is really what we are progressing toward as the best representative of a fund's return when there are cash flows on a variety of dates throughout the life of an investment.

1: "Simple Dietz Method," Wikipedia, n.d., accessed January 20, 2025.

To visualize the difficulty with different amounts invested at different time intervals, imagine that a $100,000 investment went up 15% in the first year, down 21.74% in the second year, and up 22.22% in the third year. At the end, the investment would be worth $110,000, which would be a 10% holding period return. But if the investor were to put in an additional $25,000 at the end of the first year because the investment was hot, then withdrew $25,000 at the end of the second year when things had soured, they would have had more money at work in the bad year, and the ending value would be only $103,354, or a 3.35% holding period return, despite the annual returns being identical. This example highlights the importance of money-weighting returns—cash flows and when they happen are important—something the IRR calculation takes into account, which will be covered later in this report.

Generally speaking, an investor should use these holding period returns for quick calculations to see how one investment has performed but should be cautious when using them to compare one investment to another or to a benchmark because time is not well captured in these simple equations.

Generally speaking, an investor should use these holding period returns for quick calculations to see how one investment has performed but should be cautious when using them to compare one investment to another or to a benchmark because time is not well captured in these simple equations. In the first set of examples, only one year had passed, but in the last discussion, the same $100,000 was invested over three years. Thus, the returns discussed are not properly comparable; making 10% over one year is a more celebratory accomplishment than taking three years to achieve the same result. For a proper comparison, returns compared to each other or to benchmarks should be normalized to adjust for the time element.

Annualizing returns

It is common practice to normalize returns for periods greater than one year by annualizing them. For periods shorter than one year, however, returns should not be annualized because this would be an extrapolation rather than a normalization. In other words, annualizing a shorter-than-one-year investment period is like saying that the investment achieved that return during the whole year, when in fact the capital for that investment was invested elsewhere, or possibly nowhere, for the portion of the year outside of the holding period.

Putting this in a simple and extreme example, if a person was to buy an apple for $1 today and sell it for $1.25 in one month, that would be a 25% return. If we annualized that one-month number (1,355%), this result assumes that the person was able to buy an apple and sell it for 25% more the next month every month of the year. If we also assume that on the day the first apple is sold that $1.25 is the going rate for apples that day, that would also mean that apple prices would be continually going up 25% every month, reaching $14.55 on the last day of the year. While hyperinflation could be an explanation in this scenario, we assume that was not the economic environment of the example, and common sense tells us that a 25% return every month is extremely unlikely for apple prices. Investors should be wary of any short-term numbers that have been annualized—they do show up in pitch decks from time to time. In fact, it is quite common, if misleading, to annualize monthly inflation figures rather than showing a rolling 12-month figure or just leaving the figure unannualized for shorter periods.

For periods longer than one year, however, returns can and should be annualized to make them more comparable to other investment opportunities. In order to compare the results of investments that were invested for different time periods, it is customary to normalize the figures to an annual rate, which may be done with this equation:

$$Annualized\ HPR = (1 + HPR)^{\frac{1}{N}} - 1$$

There are a few steps to annualizing a return. First, the return needs to be in decimal, rather than percent, form. This means that a 10% HPR should start off as 0.10 in the equation. Then 1 should be added to this raw number in order to avoid the issues that could come from the use of negative numbers in an equation. If the return was -10%, then 1 plus the return will become a positive 0.9 in the first part of the equation.

Next, we have the annualizing portion of the equation to consider, which is found in the exponent portion: 1 divided by the number of years of the holding period. The number of years does not need to be a whole number, so if an investment was purchased January 1, 2020, and sold June 30, 2023, then the number of years should be 3.5. The last part of the equation removes the 1 that we added to the initial return to get us back to a decimalized return.

For example, say an investment grew from $100 to $275 between January 1, 2020, and June 30, 2023. The annualization calculation on this 175% return would look like this:

{(1+1.75)^(1/3.5)} – 1, or 0.3351, which would be a 33.51% annual return in percentage terms.

Some fund managers would put in their materials that they earned a 175% return on this investment. While this is technically true, it is more appropriate for them to annualize that return for ease of comparison to other investments. Most every index return that reflects periods longer than one year will be an annualized figure, so it is important that when benchmarking the return of an investment against an index, they must be in the same annualized time basis for proper comparison.

Most every index return that reflects periods longer than one year will be an annualized figure, so it is important that when benchmarking the return of an investment against an index, they must be in the same annualized time basis for proper comparison.

Time-weighted return

Holding period returns work best when measuring a fairly stable portfolio of assets with few cash flows in or out. If cash flows occur in large amounts and at widely divergent times, it is common to build on the HPR concept by calculating HPRs for short periods like one quarter and then linking the returns together to get to a geometric mean, which in investment parlance is called a time-weighted return. The geometric mean formula takes into account the returns of each holding period and also incorporates the concept of compounding.

For our first example, we will assume no cash flows into or out of a fund, but when an investment goes from $100 to $110 in the first period, then the second period return will be on a $110 base.

Compounding returns (Hypothetical data)

Date	Value	Holding period return
January 1, 2020	$100.00	
January 1, 2021	$110.00	10%
January 1, 2022	$99.00	-10%
January 1, 2023	$128.70	30%
	Cumulative total return:	**28.7%**
	Annualized return:	**8.77%**

Source: PitchBook

Not that we suggest that anyone do this, but using a simple average of the returns gets a 10% result: (10 + -10 + 30)/3 = 10. But because more money was at work in the down year, we got a poorer result than the $133.10 we would have finished with if we had earned 10% every year. ($100 + 10% = $110. $110 + 10% = $121. $121 + 10% = $133.10). But we only got to $128.70, so the actual return must be lower than what the simple average provides.

The TWR formula, with the annualizing feature built in, looks like this:

$$TWR = [(1 + HPR_1) * (1 + HPR_2) * \dots (1 + HPR_n)]^{\frac{1}{N}} - 1$$

The end result of this formula is the mean return of this portfolio on an annual basis. HPR1 is the holding period return for the first period, HPRn is the holding period return for the last of n periods, with each other HPR return treated the same way. For the same reason as in the HPR equation, we add 1 to account for negative returns, then reverse out the 1 when most of the math has taken place. With the above example, the calculation looks like this: [(1.1)*(0.9)*(1.3)]^(1/3) = 8.77%, which is indeed lower than the return that would come from a simple average.[2] If we did an annualized holding period return of the growth of $100 to $128.70 in three years, we would get the same result, so the geometric mean—or time-weighted return—gets us to the answer we expect.

A more complex example of TWR:

Time-weighted returns (Hypothetical data)

Date	Cash flow	Value	Annual HPR
January 1, 2020	$100,000 investment	$100,000	
January 1, 2021		$115,000	15.00%
January 1, 2022		$90,000	-21.74%
January 1, 2023	-$110,000 redemption	$0	22.22%

Source: PitchBook

2: People do still use simple averages to come up with an approximation of returns. As a quick and dirty calculation, it gets one in the ballpark, but these returns should never be reported or used as a figure in benchmarking exercises.

This is still a fairly simple example for which a few methods would work to calculate the return. But to show the math, the TWR would be [{(1.15)*(0.7826)*(1.2222)}^(1/3)] – 1 = 3.23%. Because there were no interim cash flows, an annualized holding period return would have come back with the same result of 3.23%. Before annualizing either figure, the return would be 10%, which fits with what we observed with an investment growing from $100,000 to $110,000. While we have not discussed the calculation yet, an IRR calculation would also result in a 3.23% return in this simple scenario.

If we add cash flows into the scenario, we begin to see the weakness of the TWR:

Time-weighted returns with cash flows (Hypothetical data)

Date	Cash flow	Value	Annual HPR
January 1, 2020	$100,000	$100,000	
December 31, 2020		$115,000	15.00%
January 1, 2021	$25,000	$140,000	
December 31, 2021		$109,564	-21.74%
January 1, 2022	-$19,562	$90,002	
January 1, 2023	-$110,000	$0	22.22%
		Annualized TWR:	**3.23%**
		Dietz method:	**1.46%**
		IRR:	**1.36%**

Source: PitchBook

In this modified scenario, the TWR comes back with the same 3.23% answer because each annual return was the same in this scenario versus the previous. But a sense check finds a problem: We actually invested $125,000 in total and redeemed $129,562—we made $5,438 less than in the first example. The fact that TWRs do not take into account the amount of money invested in each time period creates a problem in accurately reflecting reality when the size of the capital pool varies over time.

If we use the Dietz method to adjust for the interim cash flows in this scenario, we get an annualized 1.46%. This feels closer to the truth, as we know we made less money in this scenario. This is because we added capital when the investment had just finished an up 15% year, but the next year, when more capital was at work, the fund declined 21.74%. We then took some money out at the low, and that capital missed out on the best year of the three. An IRR calculation, explained later in this report, comes up with a 1.36% return, close to the Dietz method.

Private market funds

These examples have assumed that when the investor decided to invest in a fund, the allocated capital was provided to the fund on day one, but this is not what happens with the typical drawdown fund found in the private markets. In fact, it is

not uncommon for the full commitment of an LP to never be fully at work during the lengthy life of the fund, as capital is sometimes returned before the full commitment is called, particularly when capital is reserved for follow-on investments in existing portfolio companies beyond the investment period.

Simple PE example

PE returns (Hypothetical data)

Date	Cash flow	Remaining value	Year-end HPR
January 1, 2020	Close date	$0	
June 30, 2020, call	$50,000 (portco 1)	$50,000 portfolio company (portco) 1	
January 1, 2021		$55,000 portco 1	10.00%
June 30, 2021, call	$50,000 (portco 2)	$55,000 portco 1 $50,000 portco 2 $105,000	
January 1, 2022		$60,000 portco 1 $50,000 portco 2 $110,000	4.76%
June 30, 2022, distribution	-$90,000 (portco 1)	$45,000 portco 2	
January 1, 2023		$45,000 portco 2	22.73%
January 1, 2024		$35,000 portco 2	-22.22%
June 30, 2024, distribution	-$20,000 (portco 2)	$0	-42.86%
		HPR:	**N/A**
		Dietz method:	**7.46%**
		TWR:	**-10.96%**
		IRR:	**5.27%**

Source: PitchBook
Note: Remaining value is also called net asset value (NAV).

Putting some numbers behind the PE case, the simple holding period return is useless in this situation, as both the beginning and ending values are $0. When we calculate a return incorporating cash flows using the Dietz method, we get an annualized 7.46% return. This feels high, as we put in $100,000 and got back $110,000 over time and the fund was active for four years. This equation is now being asked to handle four cash flows that took place at various times, but the assumption of the Dietz method is that they all took place at the midpoint. It is not recommended that Dietz be used to report on the full experience of any investment with multiple cash flows occurring at multiple times over the life of an investment.

In this same PE example, the TWR also falls apart as an accurate rendering of the experience. If an investor were to geometrically average the five returns (two

half years and three full years) and then annualize over four years, the result is a -10.96% return, yet we can see that the investor is $10,000 better off. In this case it is because the return of portfolio company 1, the better-performance asset, was diluted by overlapping in the portfolio with portfolio company 2, but when portfolio company 2 was sliding in value, portfolio company 1 was already gone and the asset base was much lower. By giving equal weight to each of the returns despite different amounts of money being invested at different times, a positive money-weighted return turned into a negative TWR.

Holding period returns can be a decent approximation of how an investment performed when adjusted for cash flows—though the approximation degrades dramatically when more time is involved and cash flows are widely dispersed throughout.

In sum, holding period returns can be a decent approximation of how an investment performed when adjusted for cash flows—though the approximation degrades dramatically when more time is involved and cash flows are widely dispersed throughout. The more cash flows that are included in one holding period—and if those cash flows are spread across the holding period—the less exact the HPR will be, as different amounts will be invested at different points of that period.

TWRs are a good approximation of how an investment performed when the pool of assets remains relatively constant with few cash flows in or out of the fund. Indexes and funds of liquid assets reasonably report returns that make this assumption, though the experience of the underlying investors is often vastly different if they trade in and out.

TWRs are a good approximation of how an investment performed when the pool of assets remains relatively constant with few cash flows in or out of the fund. Indexes and funds of liquid assets reasonably report returns that make this assumption, though the experience of the underlying investors is often vastly different if they trade in and out.[3] Depending on the liquidity of the underlying assets, indexes, and mutual funds, some evergreen funds calculate daily, monthly, or quarterly holding period returns, then geometrically link the results, such as calculating TWRs, to report longer time horizons.[4]

Based on this discussion, an allocator that has many moving parts within a complex portfolio may think that TWRs will be of no use to them, but at the total portfolio level—the collection of all of the holdings of public equity, private equity, and fixed income, among others—a monthly or quarterly holding period return based on the total portfolio value from one point in time to another may be reasonably accurate if few assets are being contributed or withdrawn from the total portfolio. Then those holding period returns may be linked to arrive at a TWR reflecting longer periods able to be benchmarked against the TWRs being reported by indexes.

Internal rate of return

As depicted in the previous example, private market funds have the complication of not only cash flows going in and out at various moments over the life of the fund but also a long holding period over which this is happening, so a return metric that accounts for both dynamics is required. Given that a PE fund starts with $0 invested and ends with $0 invested, a holding period return calculation is problematic, even with an adjustment for cash flows. We also saw above that returns calculated over shorter holding periods linked into a TWR may also do a poor job of reflecting reality. This brings us to the IRR calculation.

3: "Why Investors Missed Out on 15% of Total Fund Returns," Morningstar, Jeffrey Ptak, CFA, August 15, 2024.
4: For more on evergreen funds, please read our piece called The Evergreen Evolution.

An IRR is sometimes called a dollar- or money-weighted return. In essence, this equation accounts for how much capital flows into or out of a fund and the dates on which these flows occurred. The first point is important in the context of a PE fund with multiple holdings. If the best-returning company in a $100,000 fund grew 100% but only $5,000 in capital was invested in it, then the impact of that one investment on the total fund will be muted. Time is important because how much is invested at the times the fund is doing its best or worst will meaningfully impact the money made and the overall return reported by the fund.

The IRR equation:

$$\text{Net present value (NPV)} = \sum_{t=0}^{T} \frac{Net\ cash\ flows_t}{(1+IRR)^t} = 0$$

In technical lingo, the IRR represents the rate at which a historical series of cash flows are discounted so that the NPV of the cash flows equals zero. The formula assumes that the NPV is 0 because it is supposed to represent the full life cycle of the fund after the investments have been liquidated.

While sitting for a Chartered Financial Analyst (CFA) exam, a candidate may be asked to build out the formula and hand calculate an IRR with their trusty HP or TI calculator,[5] but most people use a tool such as the XIRR function in Excel,[6] as the equation and its explanation are complex for real-world examples. This formula requires all of the cash flows of an investment, the current NAV if the fund has not fully liquidated, and the dates on which they occurred. The resulting return is the number that solves for what return would make all of those numbers work out for the amounts and dates provided.

Unlike holding period returns, the IRR, when utilizing the XIRR function in Excel, comes out already annualized because the timing of each cash flow, represented by the t, is incorporated into the equation. In industry parlance, the IRR on a PE fund will be called "the return" just like other investment returns on other types of assets. But it is important to be aware of which calculation is involved when looking at any of the returns discussed in this report, as it is often not appropriate to compare the results of two different return calculations to one another.

In industry parlance, the IRR on a PE fund will be called "the return" just like other investment returns on other types of assets. But it is important to be aware of which calculation is involved when looking at any of the returns discussed in this report, as it is often not appropriate to compare the results of two different return calculations to one another.

Simple PE example redux

Looking at the simple PE example again, the IRR comes out to 5.27%—an annualized return taking in the full lifespan of the fund. This figure is higher than the negative TWR but lower than the Dietz method, so which makes the most intuitive sense? We know that the fund made $10,000 on $100,000 invested, so the negative TWR immediately feels wrong. We know the fund cash flows were scattered across the full life of the fund, while the Dietz method assumes the June 2021 capital call for portfolio company 2 and the June 2022 distribution for the sale of portfolio company

5: "CFA Exam Calculator Policy," CFA Institute, October 14, 2024.

6: There are three IRR functions in Excel, but when working with cash flows that happen at irregular intervals, the XIRR is the most appropriate one to use.

1 both occurred at the midpoint of the fund's life, which would have been June 2022. The IRR accounts for the actual timing and for the longer holding period for the poor asset, which put downward pressure on the IRR in relation to the Dietz figure.

Interim valuations and IRRs

IRRs may be calculated before a fund has reached its natural end, with the current NAV of the fund used as a "final" cash flow. The substitution of the current NAV for the final distribution amount in the interim IRR equation assumes that the current NAV is correct, which could be a poor assumption given that GP-determined valuations are estimations that are often a step removed from a market-based determination. Illiquid assets must use proxies for valuations rather than a traded price on that particular company. Because of the use of NAV in interim IRRs, biases may be built into these calculations that favor the GP's view of the portfolio, something LPs should stress test from time to time.

Hypothetical PE returns with interim IRRs

Date and investment action	Cash flow for June 2022 interim IRR	Cash flow for June 2023 interim IRR	Cash flow for June 2024 final IRR
January 1, 2020, fund commitment	$0	$0	$0
June 30, 2020, call portco 1	$50,000	$50,000	$50,000
June 30, 2021, call portco 2	$50,000	$50,000	$50,000
June 30, 2022, portco 1 distribution + portco 2 valuation	$135,000	-$90,000	-$90,000
June 30, 2023, NAV		$35,000	
June 30, 2024, portco 2 distribution			-$20,000
IRR	**21.76% interim**	**13.33% interim**	**5.27% final**

Source: PitchBook

Simple PE example with interim IRRs

Using the above example but calculating the IRR at various points in the fund's life, the return at two years is stunning at 21.76%. This is a function of portfolio company 1 having a successful exit, the fact that only two years have passed, and portfolio company 2 having not lost all of the value that it eventually will. At the three-year mark, the successful portfolio company 1 is in the past, time is passing, and portfolio company 2 is losing value, bringing the IRR down to 13.33%. The final sale amount of portfolio company 2 in 2024 is only $20,000, a 60% loss on the original $50,000 investment. This fund, which looked great in 2022, had a very different IRR two years later. While these moves may not be representative of a real-life fund, it is not uncommon for the IRR on a portfolio to change dramatically from interim valuations to the final IRR when all the portfolio companies have been exited and results are locked in.

The value-add that an investment manager can bring to a private market investment comes in three major phases: buying at a good price, managing the portfolio

investments in a way that increases their value, and selling at a good price. If an LP is planning to invest for the entire life of a drawdown fund, an argument could be made that measuring the success of a fund before the fund manager has had a chance to implement all three phases of its approach is an incomplete picture of manager skill and potentially unfair to the fund manager. Knowing that they are being evaluated, however, there is evidence that fund managers may determine their portfolio company valuations in such a way as to make the manager look good, particularly when they are in the process of raising the next fund.[7]

In our previous research on performance persistence that analyzed interim IRRs, we found little relation between how well a fund was doing at the time it came back to market for a follow-on fund and how that fund actually performed by the time it had fully liquidated. In other words, if Fund II was top quartile when the fund manager went to market with Fund III, there was no predictive information here as to which quartile Fund II or Fund III would eventually land in. This indicates that while people may feel it is their responsibility to see and evaluate interim IRRs, these returns may not be as informative as they think. Furthermore, in 2022, we reported that PE and VC valuation markups on unrealized investments were at historically high rates—a situation that caused many in the LP community to raise questions when the public equity markets turned but their private portfolio valuations did not. There are signs that LPs remain skeptical about NAVs and are probing the valuation policies of their GPs as a consequence.[8]

A huge caveat exists when it comes to IRRs that is the root cause for the admonition not to directly compare IRRs to TWRs: An LP's total commitment is fully invested for what may be only a small portion of the life of a PE fund.

A huge caveat exists when it comes to IRRs that is the root cause for the admonition not to directly compare IRRs to TWRs: An LP's total commitment is fully invested for what may be only a small portion of the life of a PE fund. One reason for this is that funds may start returning capital before they have called the entirety of the original commitment. The IRR thus only accounts for the portion of committed capital that is invested in the fund at any given time. Other returns assume the capital invested was in place throughout the life of the investment. While IRRs and TWRs are both reported as a percentage and investors call them both returns and they know returns should be compared to one another for context, one is measuring a stable pool of assets while the other is measuring a wildly varying amount of assets over the life of the investment.

Finally, while evergreen funds in the private market space report TWRs, investors should be aware that estimated valuations on illiquid assets influence the reported returns from these funds for any given period. In a drawdown fund, such as the typical PE fund, periodic valuations are interesting to observers, but until capital is returned, valuations are not typically actionable. In an evergreen fund, however, when some level of liquidity is made available on predetermined dates, people are buying and selling interests in the fund at valuations that may bear little relation to the true worth of the portfolio because these assets do not trade regularly like a publicly traded stock. In fact, the evergreen asset manager often has cash incentives tied to the estimated valuations on these funds—something that would not happen

7: "Interim Fund Performance and Fundraising in Private Equity," ScienceDirect, Journal of Financial Economics, Brad M. Barber and Ayako Yasuda, April 2017.
8: "LPs Harbor Valuation Concerns, Survey Shows," PitchBook, Jessica Hamlin, August 21, 2024.

in the typical PE fund where carried interest incentives are not generally paid until the illiquid asset is sold. Investors should pay attention to the returns coming out of evergreen funds as a result. In this case, it is not a problem with which return is used, but rather with the valuation numbers being used to do the calculation.

A total portfolio view

How would an investor then approach calculating the full return on the assets committed to a private market drawdown fund?

Imagine an investor commits $1 million to a traditional drawdown private equity fund. Assuming that investor actually has $1 million ready to invest, over the life of that fund, greater or lesser portions of that $1 million will be invested in the fund, but the remainder will likely be in other investments. Assuming the ability to perfectly swap between public equities and funding capital commitments—an admittedly unrealistic assumption—here is a simple five-year example of the accounting for this investment capital:

A $1 million commitment to a fund, with the capital invested in stocks when not in the fund (Hypothetical data)

Date	Private equity fund	Stock portfolio	Total portfolio value
January 1, 2010	$0 invested	$1,000,000 in stocks	$1,000,000
January 1, 2011	$500,000 called to purchase portco 1 NAV = $500,000	$1,200,000 in value Sell $500,000 $700,000 in stocks	$1,200,000
January 1, 2012	$500,000 called to purchase portco 2 Portco 1 held at $750,000 NAV = $1,250,000	$600,000 in stocks Sell $500,000 $100,000 in stocks	$1,350,000
January 1, 2013	Portco 1 sold for $900,000 Portco 2 valued at $700,000 NAV = $700,000	$150,000 in stocks Buy $900,000 in stocks $1,050,000 in stocks	$1,750,000
January 1, 2014	Portco 2 valued $600,000 NAV = $600,000	$750,000 in stocks	$1,350,000
January 1, 2015	Portco 2 sold for $400,000 NAV = $0	$800,000 in stocks Buy $400,000 in stocks $1,200,000 in stocks	$1,200,000

Source: PitchBook

In this case, we started with $1,000,000 and ended with $1,200,000, which is a 20% return in five years, 3.71% annualized. But how did the PE fund do? It called $1,000,000 and returned $1,300,000, so that looks like a 30% return that would annualize to 6.78%. But it only had the money for a portion of that time, with the remainder parked in stocks when not in use by the fund. We can examine the fund and stock returns separately as an initial exercise.

Why IRR and not TWR for PE portfolio? (Hypothetical data)

Date	PE cash flows	Remaining value	Cash flow-adjusted annual return
January 1, 2010, commitment date, $1,000,000			
January 1, 2011	$500,000 call	$500,000	
January 1, 2012	$500,000 call	$1,250,000	25.00%
January 1, 2013	-$900,000 distribution	$700,000	28.00%
January 1, 2014		$600,000	-14.29%
January 1, 2015	-$400,000 distribution	$0	-33.33%
IRR	**13.43%**	**TWR**	**-2.22%**

Source: PitchBook

By running the XIRR function in Excel on the calls and distributions of the fund for the dates they happened, we determine the IRR on the fund comes out to 13.43%, well above 6.78%—in large part because the amount invested in the fund was much lower than $1 million at most points in time, but also because the first capital call happened one year after the closing date on the fund, so the holding period was shorter than the total investment period. If we were to use the annual marks on the portfolio to get to a TWR, however, a very different picture results. In the example, portfolio company 1 performs very well in a fairly short period. Portfolio company 2 had a period of initial optimism and then a lingering decline, leading to multiple years of poor performance for the fund. The TWR comes out at -2.22%, which does not feel right when we know the fund provided a $300,000 profit.

But what about the rest of the assets not in the hands of the PE manager for the full duration? The full experience of the LP's $1 million allocation must account for where the assets were parked when the PE manager did not have them. Because capital is flowing in and out of the stock portfolio, an IRR, or money-weighted return, also gives a good picture of the performance of those assets.

Why IRR and not TWR for a remainder portfolio with significant cash flows going in and out? (Hypothetical data)

Date	Stock cash flows	Value	Cash flow-adjusted annual return
January 1, 2010		$1,000,000	
January 1, 2011	-$500,000 to PE	$700,000	20.00%
January 1, 2012	-$500,000 to PE	$100,000	-14.29%
January 1, 2013	$900,000 from PE	$1,050,000	50.00%
January 1, 2014		$750,000	-28.57%
January 1, 2015	$400,000 from PE	$1,200,000	6.67%
IRR	**-3.22%**	**TWR**	**3.29%**

Source: PitchBook

The PE portfolio provided $300,000 more than it invested, so if the final portfolio value was $1.2 million, the stock portfolio must have lost $100,000. Which return best reflects this?

Once again, we know directionally what the return should look like—the remainder portfolio lost $100,000, so the return should be negative. But a geometric average of the annual returns, when three years out of five were positive, would have us believe that the stock portfolio had an annualized return of 3.29%. Because the amount of money invested in stocks each year was different—more was invested in stocks in the down years—a better reflection of how this portion of the portfolio did is the IRR calculation, which came up with -3.22%.

So then how best to approach the total portfolio return? The total portfolio in this scenario was a closed-loop system—no cash flows entered or exited the system during the life of the investment. Because of this, both returns will provide the same answer. The "cash flows" for calculating the IRR are the starting and ending values five years apart ($1 million and $1.2 million), while the path to get there, which is tracked by annual TWRs, does not matter. Both will provide an annualized return of 3.71%, which is nearly 1,000 basis points below the return of just the PE fund. An LP excited about its 13.43% PE return might decide that the stock index return of 3.29% (assuming the remainder portfolio was in a stock index, the reported index return would be the annualized TWR) was a poor showing and PE is far superior to public equities. But the numbers are not comparable. A different experience in the remainder portfolio would provide a different answer, but if one does not look at the whole portfolio, the PE returns may be extremely overstated in terms of investor experience on the total assets committed over the full five-year period.

If only a portion of an investor's assets targeting PE are actually invested in PE, then the return that investor gets on PE does not tell the whole picture. This highlights why cash flow modeling and commitment pacing is so important to LPs—if an investor wants their PE allocation to be earning PE returns, they need to have as much of that money as possible in the hands of PE fund managers, not waiting on the sidelines.

As we have seen, if only a portion of an investor's assets targeting PE are actually invested in PE, then the return that investor gets on PE does not tell the whole picture. This highlights why cash flow modeling and commitment pacing is so important to LPs—if an investor wants their PE allocation to be earning PE returns, they need to have as much of that money as possible in the hands of PE fund managers, not waiting on the sidelines.

By having the assets for less time, PE firms have a time advantage in the calculation of returns, as the shorter the hold period, the less the annualization of the fund return brings the IRR down. Put another way, an investor doubling their money in two years is much more interesting than doubling it over 10 years. Even shortening the time a fund manager has an investor's money by a partial quarter can influence returns, which is one reason GPs have become fond of subscription lines.[9] This is also why the Institutional Limited Partners Association has released templates for fund and performance reporting that will allow investors to strip out the effect of fund financing for a better assessment of a GP's skill.[10]

Fees are another interesting question for this discussion on returns. The returns that most data providers, including PitchBook, report are all net of fees because we

9: "Inflating Returns With Subscription Lines of Credit," MSCI, Patrick Warren and Luis O'Shea, January 9, 2024.
10: "ILPA Releases Updated Reporting Template and New Performance Template for Industry Adoption," Institutional Limited Partners Association, January 22, 2025.

are not provided cash flows in enough granularity to provide gross-of-fee returns. LPs should have that granularity and can calculate their own net- and gross-of-fees figures based on their own cash flow data between themselves and their fund managers. Making a return net of fees (meaning that the fees paid to the fund manager have been accounted for in the return) requires treating a fee like a capital call whose value immediately goes to zero. From the LP's perspective, fees are just negative returns. Some GPs may feel that for a true apples-to-apples comparison, it is fair to use gross returns to show their investment skill versus another fund manager. But if the fees are different, that could also be seen as a choice by the fund manager. As it impacts the ultimate return to the LP, it is important not to gloss over that part of the return experience. LPs should be aware of which returns are being reported as they evaluate their GPs and ensure that the benchmarking solutions they select use the same treatment of fees.

Public market equivalents (PMEs)

To create a more apples-to-apples comparison of returns in private and public markets, several variations of PMEs exist. PMEs are a method of evaluating whether an investor was better off having invested in a private market fund or a public market alternative. An assumption is made that as capital is called down or distributed by a PE fund, cash flows to and from a hypothetical investment in an index like the S&P 500 are created for the same amounts and time periods. There are a few variations of the metric, but we often use the Kaplan-Schoar PME (KS-PME) and direct alpha calculations for our performance reporting.

The KS-PME calculates a multiple of invested capital in a private fund relative to what it would have been in a public market index. A multiple above 1x implies that the private fund outperformed the index, while a reading below 1x suggests the private fund underperformed the index.

$$KS\text{-}PME_T = \frac{\frac{NAV_T}{Index_T} + \sum_{t=0}^{T} \frac{Distribution_t}{Index_t}}{\sum_{t=0}^{T} \frac{Contribution_t}{Index_t}}$$

Because the KS-PME is a multiple and does not provide a time element, an annualized return can also be created with the adjusted cash flows and NAV values using the XIRR formula. This provides a metric called "direct alpha" and is a measure of the underperformance/outperformance of a private fund in annual percentage terms.

The choice of the index used for comparison should be considered carefully. The S&P 500 is common, but it is questionable as a fair comparison with middle-market buyout and venture funds, given the types of companies, their size, and the stage of business that is typically held in one of these portfolios. It would be even less appropriate to use the S&P 500 for an evaluation of real estate and infrastructure funds. For buyout, many use the Russell 2000 for comparison, but even that has significant differences in leverage and sector exposures. We have done some work on a buyout replication portfolio using a modified small-cap index that could provide a better comparison for evaluation.

KS-PME and direct alpha (Hypothetical data)

Date	PE cash flows	Index	Cash flow/index
January 1, 2011	-$500,000	100.00	-$5,000.00
January 1, 2012	-$500,000	85.71	-$5,833.33
January 1, 2013	$900,000	128.57	$7,000.00
January 1, 2014		91.84	
January 1, 2015	$400,000	97.96	$4083.33
		PE IRR	**13.43%**
		KS-PME	**1.02**
		Direct alpha	**1.00%**

Source: PitchBook

Taking the example from the prior section and using KS-PME and direct alpha methodologies, our example private fund outperformed the public index, but only by about 1% annually. This highlights how misleading the 13.43% return is. The index would have been up 3.29% using annual returns in a TWR, leading some to believe that PE outperformed by over 1,000 basis points. The PME and direct alpha calculations make the adjustment that show the effect of timing on the PE fund relative to the public markets. The large distribution of $900,000 occurred after public markets experienced their best period for returns, so on a relative basis, the outcome for the private fund was not as strong as it first appeared when looking at just the PE fund and stock portfolio IRRs.

As mentioned, there are variations of the PME, including the Long-Nickels, PME+, and modified PME. For a more technical explanation of the KS-PME and other variants, check out our white paper.

Multiples

There are three multiples typically used in assessing fund performance: distributions to paid-in capital (DPI), remaining value to paid-in capital (RVPI), and total value to paid-in capital (TVPI), which is a simple combination of the first two. The math is pretty simple in its most basic form, ignoring real-world accounting issues such as recycling provisions, recallable capital, and fund expenses, among other things.

DPI multiple: A measurement of the capital that has been distributed back to LPs as a proportion of the total paid-in, or contributed, capital. DPI is also known as the cash-on-cash multiple or the realization multiple.

$$DPI_T = \frac{\sum_{t=0}^{T} Distribution_t}{\sum_{t=0}^{T} Capital\ call_t}$$

The resulting figure, which will grow throughout the life of the fund, will largely be informative only in the context of expectations or the reported figures from other funds. Early in a fund's life, DPI will be minimal, as portfolio companies are still being acquired and the expectation is that the fund manager will hold those investments for multiple years before putting the companies back up for sale. Thus, in the first five or so years of the fund, investors should expect a low DPI and maybe even be suspicious of a high DPI if expectations were for long-term hold periods on each fund investment.

RVPI multiple: A measurement of the unrealized value of a fund as a proportion of the total paid-in, or contributed, capital.

$$RVPI_T = \frac{NAV_T}{\sum_{t=0}^{T} Capital\ call_t}$$

The term "unrealized" in PE lingo means that a portfolio company is still held in the fund and has not yet been sold to a new buyer. After a sale from a fund, the investment is considered realized. The unrealized, or remaining, value of each investment left in a fund is added together to become the NAV of this equation. As paid-in capital is a hard figure representing all of the capital calls contributed to a fund and the NAV is an estimation, the RVPI should be recognized as a multiple highly dependent upon these valuations, which can change based on GP expectations for each portfolio company. These expectations should be based on how the companies are doing operationally and how similar companies are being priced in the current market environment, but interim valuations may bear little relation to what capital these companies ultimately provide to investors when they are sold.

TVPI multiple: A measurement of both the realized and unrealized value of a fund as a proportion of the total paid-in, or contributed, capital. Also known as the investment multiple, TVPI is calculated by adding together the DPI and RVPI of a fund.

$$TVPI_T = \frac{NAV_T + \sum_{t=0}^{T} Distribution_t}{\sum_{t=0}^{T} Capital\ call_t} = RVPI_T + DPI_T$$

Because the RVPI is one of the components of TVPI, TVPI will also be dependent on GP valuations and should be relied upon as a signal of the GP's feelings about the portfolio rather than as an indication of the fund's final outcome—at least until a good portion of the fund has been realized, when more of the TVPI will be based on the cash multiple of DPI rather than the unrealized valuation of RVPI. As more of a fund's portfolio companies have been exited, the DPI will converge toward the TVPI and the RVPI will approach zero.

Fund multiples (Hypothetical data)

Date	PE cash flows	Remaining value	Paid-in capital	RVPI	DPI	TVPI
January 1, 2010, commitment date, $1,000,000						
January 1, 2011	$500,000 Call	$500,000	$500,000	1.0x	0.0x	1.0x
January 1, 2012	$500,000 call	$1,250,000	$1,000,000	1.25x	0.0x	1.25x
January 1, 2013	-$900,000 distribution	$700,000	$1,000,000	0.7x	0.9x	1.60x
January 1, 2014		$600,000	$1,000,000	0.6x	0.9x	1.50x
January 1, 2015	-$400,000 distribution	$0	$1,000,000	0.0x	1.30x	1.30x

Source: PitchBook

Using the earlier example, we can see how the multiples on the PE fund would look over time.

DPI has been a particular focus of late, as LPs feel constrained in committing to new funds while distributions have been meager. LPs fear that capital calls on new commitments will have to be funded from other parts of the portfolio rather than cash returned from older funds. Another reason commitments have been muted: If, as we witnessed going into 2025, many fund managers are holding on to portfolio companies longer, then LPs may have more money invested in funds than their cash flow models would have predicted when they made their commitments. This leaves their asset allocation fully funded to private markets, making it difficult to justify further commitments until capital is returned and allocations thus decline. Put another way, falling allocations will necessitate further commitments in order to maintain the target allocations.

Importantly, while multiples are simple to calculate and generally representative of reality, they have no time component. The 1.3 TVPI in the above example may seem equal to a 1.3 multiple on another fund, but if one fund took five years to reach that figure and another took 10, the experience would have been very different for the same cash multiple.

Conclusion

Some investors are mainly interested in multiples, while others are more focused on IRR. Both attitudes are flawed as sole points of evaluation for reasons that have been hinted at already. IRR is not a true reflection of the return of total committed capital, and interim IRRs can be manipulated by subscription lines and unrealized valuations determined by GPs. Multiples, on the other hand, ignore the time component, and, until a fund is fully liquidated, are also at the mercy of GP valuations. Using several valid metrics is a great way to ensure all angles have been covered when evaluating how an investment has performed.

As explained in the prior report, benchmarking should be situational, tied to decision-making, and appropriate to the time period and assets being evaluated. Certain measurements handle these marching orders better than others in various settings. We advise LPs to examine which returns are being presented, ensure they are valid for the situation, and probe the weaknesses that may be reflected in the numbers. Investment returns can feel unquestionable because they are numbers and numbers look scientific, but the many cases outlined in this report should help investors take a smarter approach to investment performance measurement.