PitchBook

Does an Allocation to Private Equity Add Value?

Introducing a new way to measure PE performance

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

Credits & Contact

PitchBook Data, Inc.

John Gabbert Founder, CEO
Nizar Tarhuni Senior Director, Institutional Research & Editorial

Institutional Research Group

Analysis
Andrew Akers, CFA Senior Data Analyst
andrew.akers@pitchbook.com
pbinstitutionalresearch@pitchbook.com

Publishing

Designed by **Joey Schaffer**

Published on July 7, 2021

Key takeaways

- The addition of a 20.0% allocation to randomly selected PE buyout funds in a simple 60/40 portfolio added an average of 0.6% in annualized excess returns over 100 simulations from 1997 to 2020.
- Evaluating performance at the total portfolio level removes the need to use IRRs to measure performance because all the cash flows to and from the PE funds remain in the portfolio.
- Differences in fund selection created a wide dispersion of excess returns across the simulations—the annualized excess returns of the upper and lower quartiles were 0.9% and 0.4%, respectively.
- Driven in part by the lack of market valuations for PE funds during the global financial crisis (GFC), the relative value-add of allocating to PE was highly dependent on the timeframe considered.

Contents

Key takeaways	1
Introduction	2
Returns: More than just a number	2
Reviewing the public versus private equity performance debate	3
Methodology	4
Results	8
Conclusion	10

Introduction

When combined, the lack of transparency within the PE industry and the structure of private funds make it difficult to clearly understand performance in the space. This opacity has led to an intense debate within the investment and academic community about the benefits of investing in PE—especially relative to public equity. An investor should expect significant outperformance from PE to justify the asset class' substantially higher fees, illiquidity, and complexity. Unfortunately, calculating the performance of PE in a way that provides a meaningful comparison to public equity is not straightforward. In this note, we present a novel approach to analyzing the historical performance differences between private and public equity that focuses on simulating realistic investor outcomes. We use historical cash flow and NAV data for PE buyout funds to construct hypothetical PE portfolios in which funds are randomly selected as an allocation within a simple balanced portfolio. Our results indicate that the addition of PE to a balanced portfolio would have added value (net of fees) on average over the full analysis period from 1997 through 2020. However, the value added by PE has fluctuated over different periods, and excess returns have been less positive over the past decade ended 2020.

Returns: More than just a number

Differences in the implementation of typical public and private equity funds play a substantial role in obfuscating performance comparisons. These structural differences demand two distinct approaches to calculating investment returns to provide a meaningful measure of performance: time-weighted and money-weighted returns.

The time-weighted return is the most common way to calculate public market fund returns. It measures the compound growth of a portfolio over time, after removing the effect of any cash flows. Since the time-weighted return is not affected by the amount of money invested, it reflects the true performance of the underlying investments. PE funds, on the other hand, are typically structured as closed-end drawdown vehicles. From a return calculation perspective, the key difference compared with public market funds is that managers of PE funds control the timing and size of cash flows. Therefore, the time-weighted return, which removes the effect of cash flows, is an inappropriate measure of PE performance. Instead, a money-weighted return provides a more useful measure of PE performance—more commonly known as the IRR.

In addition to not being comparable to time-weighted returns, even for the same investment and cash flows, IRRs are harder to interpret because they do not reflect the actual value or benefit an investor receives. The reason for this is caused primarily by the IRR calculation's key assumption that the capital distributed to investors during the life of the fund is reinvested immediately and at the same rate as the IRR. This assumption is often unrealistic—particularly for high-performing funds—and typically results in inflated performance metrics. Furthermore, IRRs do not account for the opportunity costs of managing uncalled capital. Investors often need to keep a portion of the portfolio in low-returning, liquid assets to meet unpredictable capital calls for private funds, which can drag long-

term performance lower. To properly compare public and private equity performance, we must account for the inherent differences between time-weighted returns and IRRs—as well as the key assumptions in the IRR calculation.

A valid comparison of performance between two investments should also consider risk. An investment that generates higher returns because it assumes more risk does not necessarily add value. A simple way to incorporate risk into a performance comparison is to scale the excess return by volatility—commonly known as the Sharpe ratio. However, like returns, it is not straightforward to compare the volatility of public and private equity. We wrote about this in detail in a previous analyst note. The key takeaway is that the valuation practices of closed-end private funds artificially smooth reported returns, and, in turn, systematically lower volatility.

Reviewing the public versus private equity performance debate

A significant amount of research exists examining public versus private equity performance. Unsurprisingly, the debate remains unsettled; well-respected researchers have reached conflicting conclusions using different methodologies and datasets.

Most of this research focuses on market-adjusted measures of pooled or median/average fund performance, such as public market equivalents (PME).[1] Kaplan and Schoar created this metric for a paper originally published in 2003, one of the first research studies that used individual fund cash flow data.[2] The authors found that the average buyout fund with vintage years between 1983 and 1997 had a PME of 0.97, meaning that buyout funds had similar net of fee performance to the S&P 500 Index. Using the same dataset, Phalippou and Gottschalg concluded in 2009 that the average buyout fund had an excess net IRR of -6.0% versus the S&P 500 Index, which indicates material underperformance.[3] The authors made several performance adjustments that help explain the significantly different conclusion, including writing off positive NAV values after funds were more than 10 years old.

One of the critiques of these papers was that the authors used a dataset of self-reported manager returns, which could potentially lead to a biased sample. In 2003, Ljungqvist and Richardson used fund cash flow data from the investments of an anonymous LP as an alternative source and found that the median fund achieved an excess net IRR of +4.1%.[4] Using a broader LP-sourced dataset, Higson and Stucke concluded in 2012 that the median buyout fund had a PME of 1.03 against the S&P 600 Index, a small-cap benchmark.[5] On a pooled capital-weighted basis, the PME improved to 1.09, thus indicating that PE slightly outperformed the public markets.

1: The PME represents an adjusted cash multiple based on the performance of a public equity index, such as the S&P 500. A PME of greater than 1.0 implies that the PE fund outperformed an investment in the index with the same cash flow schedule.
2: "Private Equity Performance: Returns, Persistence and Capital," NBER Working Paper Series, Steven Kaplan & Antoinette Schoar, June 2003.
3: "The Performance of Private Equity Funds," Review of Financial Studies, Ludovic Phalippou & Oliver Gottschalg, January 2009.
4: "The Cash Flow, Return and Risk Characteristics of Private Equity," Stern School of Business New York University, Alexander Ljungqvist & Matthew Richardson, January 9, 2003.
5: "The Performance of Private Equity," Chris Higson & Rüdiger Stucke, March 2, 2012.

More than a decade after his first paper, Phalippou released a more comprehensive study on relative performance that included datasets from three different providers. He found that buyout funds with vintage years between 1996 and 2005 achieved PMEs above 1.2 across the three data providers and a variety of large-cap benchmarks, although relative performance was less impressive when compared with a small-cap index and active small-cap mutual funds. When looking only at buyout funds with vintage years from 2006 to 2015, Phalippou reported PMEs only slightly above 1.0, which suggests that PE performance has been on par with public equity performance in more-recent years.[6]

For more detailed information about each paper, please see the full studies linked in the footnotes.

Methodology

The previous section highlights the conflicting conclusions reached in previous research studies regarding the relative performance of PE versus public equity. The findings of each study depend most importantly on the dataset and period used in the analysis. Additionally, the focus of all these studies on pooled or average/median fund PMEs (or similar metrics) has several limitations, particularly for institutional investors trying to determine a reasonable estimate for how much value an allocation to PE could add to a portfolio. First, neither the pooled PE cash flows nor the average PE fund are investable. For industry participants to act on its conclusions, an analysis should account for the fact that investors must gain access to traditional PE through active fund management. Second, while PME and similar metrics are helpful when assessing relative PE manager performance, they do not reflect the realistic tradeoff investors face when deciding whether to allocate capital to PE or public equity. This is because PMEs normalize public equity performance by assuming an investment has the same cash flow schedule as PE, in which capital is invested over time. However, in a real-world setting, an investor that chooses to allocate to public equity rather than PE can soon be fully invested by purchasing an open-end fund. PMEs also do not make any assumptions about the ways distributions are reinvested, when, in practice, distributions are often reinvested back into new PE funds. More generally, previous studies have not addressed the challenges allocators face when investing in closed-end PE funds, most importantly cash flow management and commitment pacing.

The methodology developed for this note seeks to address the issues discussed above and contribute to the body of work on public versus private equity performance by estimating the marginal impact of adding a PE allocation to a simple 60% equity and 40% bond portfolio. To achieve this, we leverage simulation to create hypothetical PE portfolios from PitchBook's fund cash flow data that incorporate the nuances of managing investments in closed-end funds. This methodology differs from those found in previous research studies because it focuses on investor outcomes rather than a pure performance comparison.

6: "An Inconvenient Fact: Private Equity Returns & The Billionaire Factory," University of Oxford, Said Business School, Ludovic Phalippou, June 10, 2020.

The data for this analysis consists of cash flows and NAVs from 1,044 US buyout funds with vintage years from 1997 to 2020, which are sourced from both fund managers and LPs. The accompanying table summarizes the characteristics of included funds by vintage year.

Summary of included funds*

Vintage year	Count	Median fund size ($B)	Average fund size ($B)	Median PME
1997	15	$510.0	$1,201.2	1.08
1998	32	$686.5	$1,061.4	1.10
1999	25	$770.0	$1,069.2	1.42
2000	36	$1,200.0	$1,775.2	1.26
2001	18	$700.0	$1,217.1	1.31
2002	21	$950.0	$1,775.7	1.33
2003	14	$415.7	$824.6	1.25
2004	33	$650.0	$1,158.6	1.34
2005	45	$918.0	$1,659.5	1.05
2006	63	$750.0	$2,511.0	1.01
2007	60	$791.9	$2,323.8	0.96
2008	52	$958.3	$2,284.0	1.00
2009	21	$435.0	$883.0	1.20
2010	22	$520.8	$623.4	0.95
2011	37	$717.5	$1,830.9	1.08
2012	58	$725.0	$1,788.7	1.06
2013	52	$552.8	$1,388.1	1.06
2014	56	$704.5	$1,664.7	1.12
2015	64	$1,025.0	$1,767.5	1.03
2016	55	$765.0	$2,053.2	1.01
2017	73	$860.0	$2,344.9	0.95
2018	71	$1,084.0	$2,335.8	0.89
2019	76	$1,407.3	$2,598.3	N/A
2020	25	$1,500.0	$4,390.7	N/A

Source: PitchBook | **Geography:** US
*As of September 30, 2020

The simulation begins in 1997 with two identical portfolios that consist of 60.0% public equities and 40.0% core fixed income. We proxy these allocations with the Russell 3000 Index and the Bloomberg Barclays US Aggregate Bond Index, respectively. One of the portfolios, which we will refer to as the simulated portfolio, makes a commitment to three randomly selected PE buyout funds from the 1997 vintage year cohort with a target allocation of 20.0%.[7] The other portfolio maintains its original 60/40 allocation and becomes what we will refer to as the benchmark. The main conclusions from the simulation will focus on the difference in performance

7: To lessen the time it takes to reach the allocation target, the initial commitment size is 24.0% of the total portfolio.

between the simulated portfolio and the benchmark because it isolates the impact of adding a PE allocation to a more traditional portfolio. Both the simulated portfolio and the benchmark are rebalanced quarterly back to the 60/40 target allocations. In the simulated portfolio, the total equity including PE allocation is rebalanced to 60.0%. Additionally, to reflect the need for liquidity to meet net capital calls, the simulated portfolio includes a 1.5% cash allocation that comes from fixed income. The 1.5% reflects a conservative estimate that the net capital calls are unlikely to exceed 7.5% of the committed capital to PE in a single quarter. In the rare case that this does occur, we assume that the capital can be sourced from elsewhere in the portfolio.

The simulation tracks the quarterly contributions, distributions, and NAV of the PE portfolio as well as the performance of the public market allocations, rebalancing accordingly along the way. Then, at the end of the year, an additional commitment is made to randomly selected buyout funds in the next vintage year cohort—that is, at the end of 1997, funds with a 1998 vintage year are selected.[8] Optimizing the total commitment to new funds based on the expected cash flow schedule of the existing PE portfolio allows us to maintain the target allocation to PE within the total equity portfolio, which is 33.0%. We calculate the expected cash flow schedule using PitchBook's existing single fund cash flow model, and we determine the number of new funds to which investors commit by setting the maximum allowable commitment size to any one fund at 5.0% of the overall portfolio value. For example, one could consider a simulated portfolio at year-end with an allocation to PE of 27.0% of the total equity portfolio, representing an approximately 6.0% underweight relative to its target of 33.0%. The accompanying chart displays the expected relative allocation to PE given the expected cash flow schedule and different commitment sizes.

Relative PE allocations with different commitment sizes*

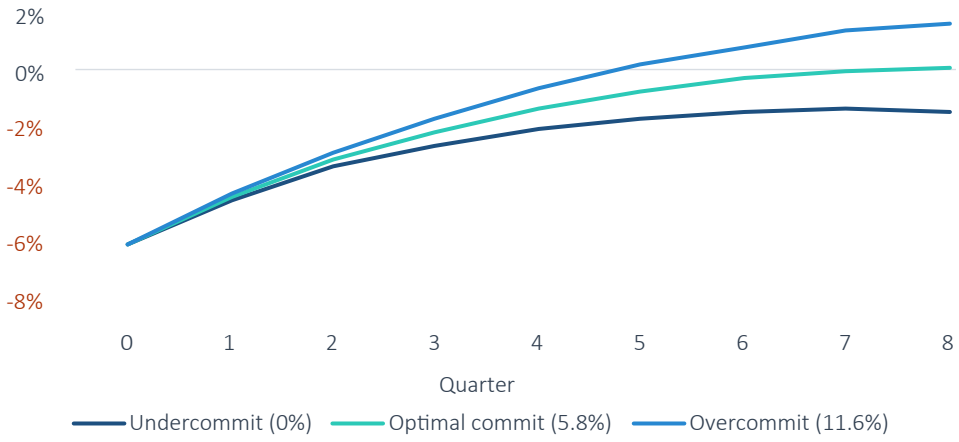

Source: PitchBook
Note: For illustrative purposes only.

8: Each fund in the vintage cohort has an equal probability of being selected, and we assume that the hypothetical investor can invest in all funds in the sample.

The optimal commitment to new funds achieves the target allocation without overshooting. When this process is repeated at the end of each year, the portfolio will eventually maintain the target allocation to PE, assuming that the expected cash flows are equal to the actual cash flows. This is, of course, an unrealistic assumption, and differences between expected and actual cash flows will cause the allocation to PE to fluctuate around its target over time. Additionally, because it is more difficult to correct for overcommitting than undercommitting, we scale the optimal commitment size by 70.0%. The commitment optimization process also assumes that when PE is above its target, the best action is to do nothing, which produces undesirable swings in the relative allocation and skips investing in some vintage years. To account for this, we implement a minimum annual commitment of 2.0% of the overall portfolio.

We repeat the process of committing additional capital to random buyout funds at the end of each year until the simulation's end in 2020. We can then calculate the absolute and relative performance of the simulated portfolio. Evaluating performance at the total portfolio level removes the need to use an IRR because the cash flows to and from the PE funds remain within the portfolio. Therefore, it is possible to calculate a time-weighted return that we can compare directly with the benchmark portfolio. The excess return of the simulated portfolio over the benchmark reflects the value an investor would have ultimately received by allocating to PE.

Approximately 20.0% of the ending simulated portfolio value is based on manager self-reported NAV, which may not be an accurate indication of the true value an investor could receive in the secondary market. PE transactions in the secondary market are typically completed at a discount to NAV. Therefore, to make the simulated portfolio performance more comparable to the benchmark, which is based entirely on market returns, we use historical data on secondary buyout sales to reasonably adjust the ending NAV of the PE portfolio. In 2017, Nadauld et al. studied proprietary data on secondary transactions for private market funds between 2007 and 2014 and found that, on average, buyout funds were priced at a 15.0% discount.[9] However, the magnitude of the average discount was heavily affected by market conditions, with discounts of close to 50.0% reported in 2009 compared with only 7.0% in 2014.[10] More recently, anecdotal evidence suggests that some funds are trading at a premium. However, to provide a conservative performance estimate, we calculate the impact of incorporating a 15.0% NAV discount. When examining performance over a long period, the results illustrate that this has a modest effect on relative performance.

9: "The Liquidity Cost of Private Equity Investments: Evidence from Secondary Market Transactions," Journal of Financial Economics, Taylor D. Nadauld, et al., November 22, 2018.
10: Ibid.

Results

The accompanying table shows the average results of 100 simulations across different periods. Each simulation invests in a unique set of PE funds.

Annualized simulation results (net of fees)*

	Portfolio return	Portfolio return adjusted**	Benchmark return	Excess return	Excess return adjusted**	Portfolio volatility	Desmoothed portfolio volatility***	Benchmark volatility
1997-2020	8.2%	8.0%	7.5%	0.6%	0.5%	8.0%	9.4%	10.3%
2000-2020	6.8%	6.8%	6.1%	0.7%	0.7%	7.3%	8.7%	10.1%
2000-2009	3.7%	3.6%	2.4%	1.4%	1.2%	7.9%	8.7%	10.7%
2010-2020	9.4%	9.4%	9.2%	0.2%	0.2%	6.5%	8.5%	9.4%

Source: PitchBook | **Geography:** US
*As of September 30, 2020
***15.0% discount applied to beginning and ending PE NAV.*
****Portfolio returns were desmoothed using the technique described in this analyst note.*

Over the full simulation, the inclusion of PE added an average of 0.6% to annualized returns relative to the benchmark. If a NAV discount of 15.0% is deducted from the final PE portfolio value, the annualized excess return falls modestly to 0.5%. It is important to note that because the NAV discount is annualized over a 23-year period, it has only a small impact on the excess return. Realized returns drive almost all the relative value-add throughout the simulation. This point highlights the fact that annualized excess returns in this context are more robust over longer timeframes because they are less dependent on NAV.

The average excess return masks a fairly wide range of outcomes in the individual simulations, which are caused by manager selection. For additional context, the distribution of the excess returns over the full timeframe is illustrated in the accompanying chart. The upper-quartile excess return was 0.9%—more than double the lower-quartile excess return of 0.4%—thus confirming a common viewpoint that manager selection is

Distribution of excess returns (1997-2020)*

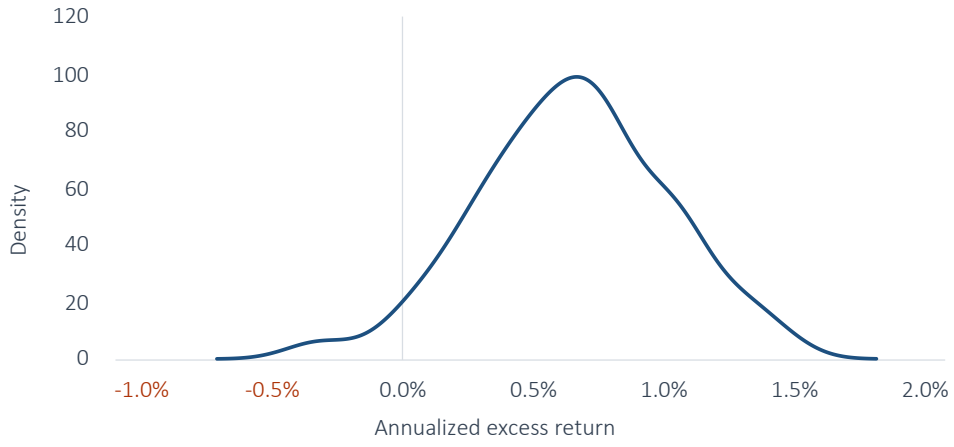

Source: PitchBook | **Geography:** US
*As of September 30, 2020

an important consideration when allocating to PE. Underperformance of the simulated portfolio was rare—only four of the 100 simulations had a negative excess return relative to the benchmark.

The value-add from PE was also dependent on the time horizon of the calculation, which is likely due in part to overly optimistic private fund NAVs in periods of market stress. For example, the average excess return in the decade ended 2009 during the GFC was 1.4%, as the benchmark returned only 2.4%. The strong performance of public markets in the following decade, which included part of the recovery from the GFC, helped reverse some of this outperformance. The average excess return of the simulated portfolios over the past 10 years was only 0.2%, and 35.0% of the simulations underperformed the benchmark. To illustrate this time dependence in more detail, the accompanying chart displays the rolling five-year excess return of each simulation and the average excess return.

Rolling five-year annualized excess returns*

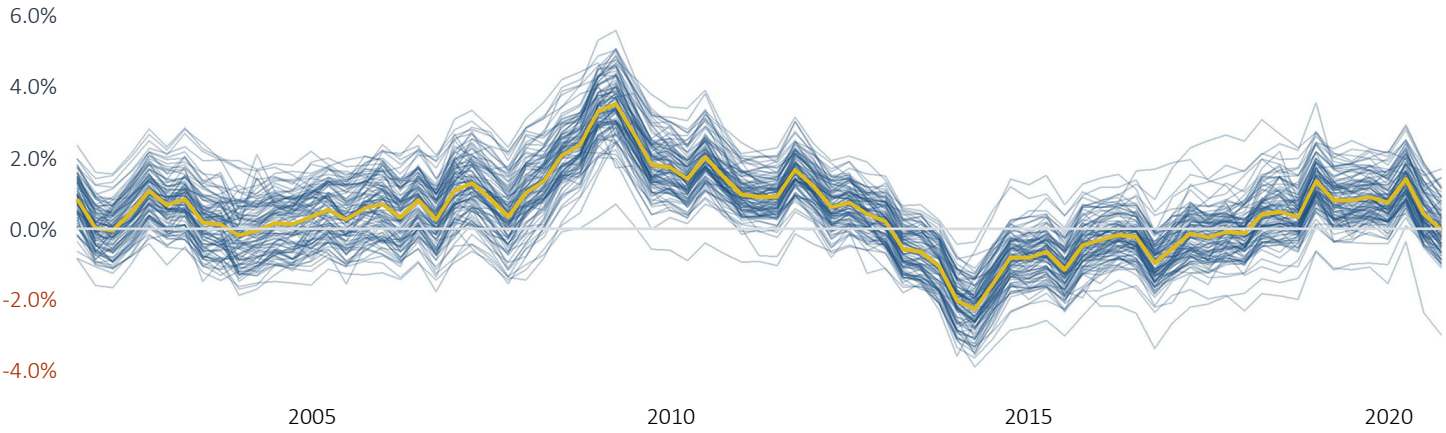

Source: PitchBook | **Geography:** US
*As of September 30, 2020
Note: Gold line represents the average.

The risks associated with investing in PE are also an important consideration. For this analysis, we define risk as the standard deviation of quarterly returns—typically referred to as volatility. Over the full period, the average annualized volatility of the simulated portfolios was 8.0%, versus volatility of 10.3% for the benchmark. While it appears that allocating to PE materially lowers overall portfolio risk, a significant caveat exists. As previously mentioned, PE returns are artificially smoothed, which underestimates volatility. After desmoothing the PE returns, the average volatility of the simulated portfolios increases to 9.4%, which is more in-line with the volatility of the benchmark. Using the adjusted volatility results in an average portfolio Sharpe ratio of 0.67, versus 0.54 for the benchmark. In addition to investment risks such as volatility, investors also need to consider the liquidity risk added from an allocation to PE—especially given that committed capital is typically locked up for seven to 10 years.

Conclusion

Despite mixed conclusions from prior academic research on the relative performance of PE, the results from this analysis illustrate that an allocation to PE would have typically increased the performance of a balanced portfolio over the past 20+ years. These results provide institutional investors with a useful framework to estimate the historical long-term outcomes of allocating to PE within a broader portfolio. This framework incorporates many of the complexities of investing in PE—including cash flow management—and can be tweaked to better reflect an individual investor's situation. Additionally, we can leverage the methodology developed for this analysis as a foundation for further research. In an upcoming note, we will dive more deeply into the impact of PE manager selection on portfolio performance.