Revisiting private equity performance computation for multi-asset investors

E. Nouvellon, H. Pirotte

Centre Emile Bernheim, Solvay Brussels School of Economics and Management
Université Libre de Bruxelles

May 2019

Abstract

Private equity has increasingly been used in portfolio for all types of investors as family offices or ultra-high net worth individuals. Financial Literature proposes different ways to compute private equity performances with results that can question the promised over-performance on public equities. The investment process in private equity funds with the system of committed capital and called capital can have a huge impact of the private equity performance in the whole portfolio and in multi-assets framework. This paper proposes an empirical study that integrates the J-curve effect on the private equity part of a portfolio and its scaling effect with the low rate environment.

KEYWORDS: private equity performance, opportunity cost of investment, general partner, limited partner, multi-asset investors, J-curve, capital calls, deposit rate, IRR, multiple on invested capital

JEL classification: G11

1 Introduction

The importance of the private equity industry has been increasing since 2000. Committed amounts have been growing from USD 10bn in 1991 to USD 180bn in 2000 (Jesse Reyes, 2002). After the exceptional turmoil in public equity markets in 2008, an impressive cycle of expansion started then for the private one. The fundraising private market compound annual growth rate (CAGR) reaches 15.1% between 2010 and 2015 according to Prequin, one of the most important private equity databases. The private asset's market size is nowadays close to USD 5.2 trillion (McKinsey, 2018).

Private equity investing (PE) is set up with a limited partnership structure involving a general partner (GP) who is the manager and limited partners (LPs) who provide capital to invest in different private companies through the general partner. Traditionally, LPs consist in several type of investors: endowment plans, family offices, foundations, public pension funds, sovereign wealth funds and private pension funds. More and more high net worth retail investors have been investing through private equity funds. LPs commit to provide a defined amount of capital to the GP through a closed-end fund with defined maturity. GP "calls" the amount in order to invest in the private firms he chooses to invest in, after a deep financial and legal due diligence. Usually, PE funds have a 10 or 12-year duration. The GP can call the capital during the investment period which is typically four or five years. The investment period follows terms and conditions; the GP cannot make any investments beyond this investment period. Then, 2 sub-periods can be drawn in a private equity funds investment cycle: the investment period during which the capital is called, and the period during which the capital is repaid to investors, with multiple committed capitals, depending on the success of the various investments. The investment period is not necessarily over when the GP starts to repay LPs. An investment made in year 1 can be exited in year 2 or 3. From the LP's point of view, GP capital calls are "contributions" whereas repaid amounts are "distribution", and constitute the cash flow stream. For non-exited investments, the estimated net asset value of the fund is taken as the last distribution.

Contrary to liquid assets like bonds or listed equities, private equity does not offer the possibility to have transaction-based prices. Measuring performance in this context is a challenging question. Traditional measures are used and are derived from contributions and distributions, such as the multiple on invested capital (MOIC) and the internal rate of return (IRR). The MOIC is the sum of distributions divided by the sum of contributions. It captures the return on invested capital without any timing considerations. The IRR measures adjusts the performance measure by considering the date of each contribution and distribution. But it does not include the (opportunity) cost of holding capital before the capital is called by the GP. The modified IRR (MIRR) corrects the IRR by capitalizing the contribution at the cost of funding. These traditional gross performance measures have been completed by market-adjusted return measures, aiming to the comparison between PE performance and benchmarks of listed equities, as in the seminal paper of Long and Nickels (1996). But contributing literature to this question is still being developed.

Both absolute return measures, such as IRR or MOIC, and market-adjusted return measures use LPs' cash contributions and distributions as inputs. As mentioned before, in the partnership, investors (LPs) commit capital ex-ante and fund managers (GP) call this capital at their own discretion during the investment period, i.e. the "contributions". In order to maximize IRR with an unchanged MOIC, GP calls capital only when an investment is closed, to minimize the "cash at work" period. On an entire investment period, the final total called capital can be less than the originally committed one. In reality, the called capital is close to 80% of the committed capital. That means that the IRR performances provided by GP to LPs neglect completely the period where the capital is committed but not called. The commitment states that LPs must provide an amount, potentially between 0 and 100% of the committed capital, within a fifteen-day (usual market practice) capital call notice. More and more private equity funds with strong creditworthiness use leverage. They contract banking bridge loans to close first investments before calling LP's capital. The leverage may range from 20% to 30% of commitments. In this case, the GP can distribute some cash before any contributions made by LPs. This is a way to reduce the number of capital calls and operational issues, but this is a clear way to improve the IRR as the bridge loan interest rate is very low compared to the return on investments made by funds.

Over the last years, IRR measures reported by GPs are often close to 10% or higher. The over-performance on listed equities is an open debate. Harris, Jenkinson and Kaplan (2015) found that the median private equity fund outperformed the S&P 500 by 1.75% per year in the 90's and 1.5% in the 2000's. The return necessary to compensate the added risks of LPs is widely viewed to represent 3%, mainly explained as a liquidity premium with respect to listed stocks. This reported 10% private equity IRR is "made" during the "cash at work" period, representing an average invested amount taking into account cash contributions and cash distributions. The literature on the Public Market Equivalent (PME) method and on realized private equity fund performances concentrates on this "cash at work" period performance. We advocate that the ultimate performance to the investor should be considered more globally, integrating the committed but uninvested cash portion as well.

One of the goals of the literature on the PME method is to compute the PE excess return on public equities to integrate the PE asset class in the context of the modern portfolio theory. Nevertheless, PE performances are shown as a standalone asset class and as if LPs were mono-asset class investors. The aim of this paper is to integrate the fact that LPs are multi-asset class investors and therefore highlight the need to correct PE performances for the bias introduced by the PE integration as a block in the whole asset allocation. This integration is firstly impacted by the shape of the contributions and the distributions commonly named J-curve, then by the risk-free deposit return. Within a multi-asset class allocation, a decision must be taken by the LPs on how the committed amounts are invested before they are called by the GP, since any reserved amounts would generate an opportunity cost. We could consider the allocation chosen by LPs on the cash distributed by the GP. The return generated on this cash, between the distribution and the end of the considered investment cycle would clearly improve the return on the mix of private equity and other asset classes. Nevertheless, the aim of our paper is not to measure the return of the entire cycle of private equity but to analyze the impact on the return of the commitment imposed to LPs at the sole discretion of the GP, who has the control on capital calls. This induces LPs to allocate their commitment to liquid and low risk assets. After the GP distributes the resulting cash, LPs have the possibility to invest it freely without any remaining commitment, which would be similar to any other investment case.

We can imagine different strategies to invest this cash before it is called. First, it could be invested in public holdings, whose portfolio is mainly allocated to PE funds, with a high correlation to the original PE invested part. Ang, Chen and Phalippou (2013) show the positive correlation of PE performance with the EV/EBITDA multiple and a negative correlation with high yield credit spreads. Using this finding and this correlation, we could find an appropriate liquid proxy portfolio to invest the committed cash before it is called by GPs. After defining the appropriate driving factors (if any) of PE and mapping them to a given public index, factor investing listed products can be a solution to replicate a listed portfolio that would "follow" PE performances except for the liquidity premium. Nevertheless, in order to manage market risk and liquidity risk, it seems appropriate for LPs to invest the non-called cash in treasury funds or deposits with a very short duration. In reality, the decision is not purely devoted to LPs as the main GPs (or the depositary bank) require the committed amount as collateral, forcing LPs to invest the committed amount in cash deposits or monetary funds (except the case where LPs accept to over-collateralize with riskier assets than cash). Then the realized return between the commitment date and the contribution dates will only be the short-term risk-free rate.

The next section of this paper proposes a literature review on PE performances. The third section will analyze the J-curve's shape impact on the PE performance, as part of a building block of a larger asset allocation. The fourth section proposes an IRR correction to integrate the J-curve effects on LPs performance. Finally, the last section illustrates our corrected IRR measure for different deposit rates.

2 On private equity performance measures

The GP usually reports two performance measures to LPs: the multiple on invested capital (MOIC) and the internal rate of return (IRR). The MOIC is the ratio of the sum of distributions from GP to LPs divided by the sum of contributions by LPs. This measure doesn't consider any timing considerations, whilst the IRR is measured by all contributions and distributions at their respective dates.

In order to complete these absolute measures and improve the benchmarking of private investments, Long and Nickels (1996) initiate a debate on the appropriate relative performance measure for PE. They propose an "index comparison method" later recognized as the first Public Market Equivalent (PME) measure aiming at comparing PE performance to listed equities benchmark. The Long-Nickels PME creates a replication of cash contributions and distributions in a public market benchmark and compare the PE IRR and the public market replication IRR. A contribution is replicated by a public index purchase whereas a distribution is replicated by a public index sale, taking into account the evolution of the public index. This method can bring problems with high distributions being inducing a net short position in the public index. Rouvinez (2003) corrects this problem with PME+ by rescaling distributions.

Kaplan and Schoar (2005) introduce a relative KS-PME method that measures the multiple between PE and the public index performance. KS-PME exceeding 1 means the fund multiple at the end of fund's life is better than if the investor would have chosen listed equities in the first place. Since this paper, several PME methods has been proposed by different authors. Gredil, Griffiths and Stucke (2014) provide a comparison of all of them.

Before Kaplan and Schoar (2005), the literature on PE performance has been quite limited due to difficulties to find appropriate non-public data. In a seminal paper, Kaplan and Schoar (2005) conclude that average PE funds disclose similar returns than the S&P 500 index. They also analyse the persistence of returns, a subject enlarged then by other authors. Managers launch regularly new "series" to make sure investors in a given fund have had a similar lifetime in the fund. Nevertheless, investments are not dispatched through the various PE series of that manager but instead, each investment is made regarding it as an opportunity for a given fund at a time. Therefore, persistence is not a trivial issue. Kaplan and Schoar (2005) measure persistence with an AR(1) model for consecutive yearly funds pertaining to the same manager. They find a significant predictability from year to year and also establish a positive link between fund size and GP's experience.

Korteweg and Sorensen (2017) measure the link between persistence of private equity firms returns and their skills. Ang & al. (2013) draw a private equity time-series return based on LP cash flows for different types of funds. Then they propose a PE return decomposition into systematic and idiosyncratic components. They conclude on the PE beta being significantly greater than one with different exposure levels depending on the type of PE sub-asset class (buyout, ventures or real estate). The idiosyncratic portion of private equity returns is assimilated to the PE premium. They compare the PE performance to that of the S&P 500 index from 1993 to 2010, and also conclude that PE beats the listed index. This finding is consistent with Harris, Jenkinson and Kaplan (2014) and Robinson and Sensoy (2011). Finally, Barber and Yasuda (2017) analyze the link between interim performance of funds and their fundraising capacity. They show that interim performance affects the fund's raising capacities.

The practitioner's approach used by LPs is to benchmark the PE IRR measure on their private equity investments to stock market indices such as S&P 500 or Russel 300, plus 300 basis points. The 3% risk premium remunerates additional risks and illiquidity of PE investments compared to listed equities. Do PE funds still return this 3% risk premium is an open question. Appelbaum and Batt (2016) conclude the PE performance based on IRR is higher than their performances based on PME. The funds' performance heterogeneity brings a difference between median and average performances of buyout funds, i.e. some skewness. The median buyout fund outperformed S&P500 by about 1% per year whereas the average over-performance is closer to 2/2.5 % per year according to the financials (Appelbaum and Batt, 2014).

3 Asset allocation discontinuity with private equity

Main literature conclusions stipulate that PE performance is better than that of public equities, either measured in absolute terms by measures such as the IRR, or by market-related measures as the PME. This literature considers realized contributions and distributions to measure LPs' performance. This PE overperformance on listed equities does not consider the discontinuity in asset allocation due to the PE inclusion in wider portfolios. The LPs' commitment to provide cash to GP during a 4 or 5-year investment period at the GP discretion creates a discontinuity in asset allocation with a higher part in cash. Table 1 illustrates this discontinuity by showing an asset allocation of an LP, with 50% in equities and 5% cash (45% in other asset classes). In this example, we consider the PE portion as an extension of public equities, and we keep the weight of other assets unchanged, assuming the PE integration is originally allocated from the listed equities part of LP's portfolio. In order to manage the market and liquidity risk or simply to pledge the cash for GP at the commitment date, the cash increases from 5% to 25% for a committed amount of 20% of the portfolio. The asset allocation's discontinuity duration depends on the capital calls made by the GP.

Table 1
Timeline of cash allocation to equities (private and public)

This table shows the evolution of cash and its allocation through time, from the first commitment of cash to the integration of private equity in the portfolio. Steps 2 and 3 illustrate the delay between both timings.

Order	Allocation step	Public equities	Private equity	Cash weight
1	Before private equity commitment	50%	0%	5%
2	After private equity commitment and before any contribution	30%	0%	25%
3	After private equity contribution for 50% of the commitment	30%	10%	15%
4	After private equity contribution for 100% of the commitment	30%	20%	5%

The significant cash increase in asset allocation could be an opportunity cost for LPs, currently not measured in PE performance measures. This opportunity cost depends on the first part of the J-curve shape drawn by contributions and also on cash return. The next section proposes a corrected IRR measure to include the effects of this J-curve shape.

4 J-curve and corrected IRR proposal

The key concept we use in this paper for the private equity performance measurement is the J-curve. The J-curve represents a timeline of the contributions and distributions cash flows as a percentage of committed capital, that displays a convexity. The entire J-curve is known only *a posteriori* after the fund closes. It depends on the investments and exits realized by the GP. The opportunity cost relates to the first part of the J-curve during the investment period when contributions occur.

Table 2
Example of a J-curve in the allocation process

This table shows an illustrative example of a J-curve with yearly contributions and distributions (payouts) in % of committed cash at inception, over a 10-year horizon. Numbers are signed in the direction of the cash flow for LPs.

Years	Date	(1) Contributions	(2) Distributions	(3) Net cash flows	(4) Net "committed" cash flows
1	31/12/18	-20%		-20%	-100%
2	31/12/19	-25%		-25%	
3	31/12/20	-25%		-25%	
4	31/12/21	-20%	10%	-10%	10%
5	31/12/22	-10%	15%	5%	15%
6	31/12/23		25%	25%	25%
7	31/12/24		40%	40%	40%
8	31/12/25		40%	40%	40%
9	31/12/26		10%	10%	10%
10	31/12/27		10%	10%	10%
			IRR	10.01%	7.04%

Table 2 shows a typical J-curve for a PE fund. Effective LPs' contributions (column 1) range from year 1 to year 5. In year 4, the first investments to be exited start providing cash distributions (column 2). As from year 5 onwards, contributions and distributions provide a net positive cash LPs contribution. Please note that this J-curve shape does not include any use of leverage.

Figure 1
Illustration of the resulting J-curve
This figure illustrates the cumulative cash weight outcome through the timeline shown in Table 2.

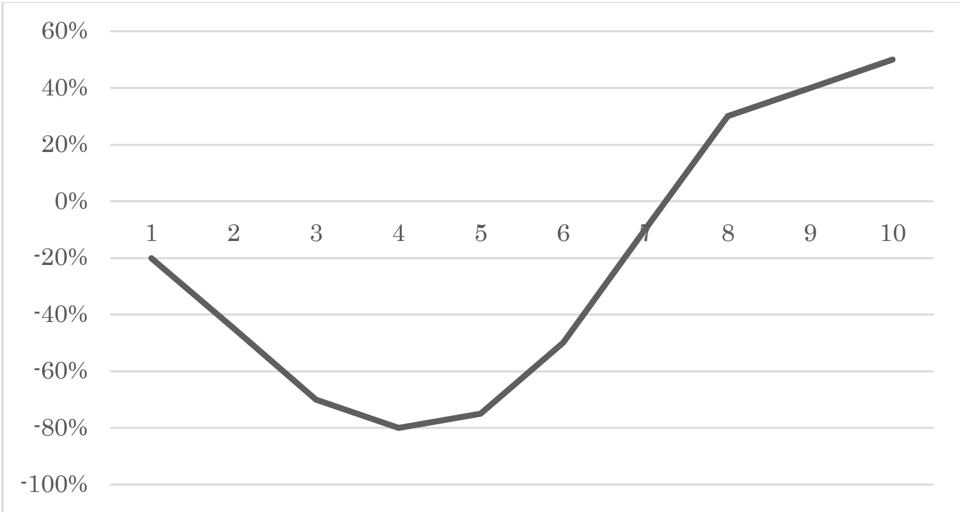

The example of J-curve in figure 1 provides a 150% MOIC and an IRR of 10.01% measured with contribution and distribution. Since the commonly accepted additional risk premium of PE to public stocks is 2.5-3%, with a historical average S&P 500 return of 7.5%, our illustrative IRR example is in line with real-life expectations.

The 10.01% IRR resulting from “net cash flows” (column 3) in table 2 only considers the “cash at work” period. In column 4 of table 2, a corrected IRR computation with “net committed cash flows” instead is proposed. This results in an IRR measure at 7.04%, assuming at this stage that there is no return on unutilized cash. We could argue that this IRR truncation is simply asymmetric, and it could integrate the return made on distributions from 2022 to 2027. But once the cash is distributed by the GP to LPs, LPs have full freedom to reinvest it.

The aim of our paper is therefore to focus on the impact of the cash commitment during the pre-contribution period. The investor does not know when the committed cash will be called by the GP. GP makes capital calls at his pure discretion for each investment’s closing. It is very difficult to have an idea of the path that will be followed by capital calls. It can depend on several variables:

- Market conditions. For PE investing, the level of EV/EBITDA multiple and credit spread are key variables to enter in a new company as a shareholder. The EV/EBITDA multiple is a component in the company's valuation by the GP. Credit spreads reflect the price of leverage that could be used by GPs to maximize their return on equity. In a low EV/EBITDA multiple and low credit spread environment, GPs can decide to take investment opportunities faster than in conditions where multiples are higher, and cost of leverage is expensive.

- GP strategy. Some GPs need to be active shareholders for several years in order to implement their industrial transformation, re-organization or scaling strategies. Other GPs have only a financial strategy that can be implemented in a short period of time.

- Type of funds**.** Secondary and co-investment funds usually call capital earlier than buy-out funds. Funds that invest in a higher proportion of debt instruments than in equity also call capital faster.

The corrected IRR should consider the return on deposit made with the committed but non-called capital. The last column of table 3 shows the impact of a correction with a 1% deposit interest rate on the IRR computation.

Table 3
Example of J-curve to compute corrected IRR with committed amount and a 1% deposit interest rate on the non-called capital
This table shows the example of J-curve and compute corrected IRR considering 100% committed cash as a contribution with a 1% interest rate.

Years	Date	Contributions	Remaining cash deposit	Deposit return	Distributions	Net "committed" cash flows with deposit
1	31/12/18	-20%	80%	0.80%		-99.20%
2	31/12/19	-25%	55%	0.55%		0.55%
3	31/12/20	-25%	30%	0.30%		0.30%
4	31/12/21	-20%	10%	0.10%	10%	10.10%
5	31/12/22	-10%	0%	0.00%	15%	15.00%
6	31/12/23				25%	25.00%
7	31/12/24				40%	40.00%
8	31/12/25				40%	40.00%
9	31/12/26				10%	10.00%
10	31/12/27				10%	10.00%
					IRR	7.35%

In a similar way, table 4 shows the corrected IRR considering different interest rate levels for the committed cash deposit's return for the J-curve example presented above with, as a result, a range equivalent to the commonly accepted PE equity premium.

Table 4
Corrected IRR as a function of the deposit rate
This table shows the proposed corrected IRR with different deposit rates, between 0% and 10%.

Deposit rate	Corrected IRR
0%	7.04%
1%	7.35%
2%	7.66%
3%	7.97%
4%	8.29%
5%	8.62%
6%	8.95%
7%	9.29%
8%	9.63%
9%	9.98%
10%	10.33%

As long as the opportunity cost is not compensated by an equivalent return on the unutilized cash, the IRR decrease results both from the J-curve shape and the lower level of the deposit rate.

5 Simulating different J-curve patterns on the corrected IRR

In this section, we generalize our previous finding on the combination of the J-curve with a deposit return for the unutilized cash portion of the commitment, by simulating many J-curves patterns. The model computes corrected IRRs associated to different J-curve shapes considering the same MOIC (multiple on invested capital), defined here at 150% for the sake of illustration.

Multiple J-curve patterns are drawn through Monte-Carlo simulations with a uniform distribution, for two different components of the cash flow stream, contributions and distributions. The first component relates to LPs' contributions during the first five years, considering a 5-year investment period. The procedure is run as follows. The total contribution is fixed to 100% of the total committed amount on the five years. The amount of the cash contribution of the first year, CO1, is drawn from a uniform distribution from [1,100]. The second-year cash contribution is then similarly drawn, but in the interval [1,100 – CO1]. The cash contributions of the three following years are computed accordingly, by reducing the size of the sample by the cash amount that has already been called.

The second component of the J-curve consists in the cash distributions by GPs to LPs. We apply the same routine as for the cash contributions except that the total amount is fixed to 150, to respect the same MOIC measure for all simulations, and can be allocated at any time during the 10 years of the time-horizon. Also, the model permits for cash contribution and cash distribution to occur simultaneously between years 1 and 5. Table 5 shows an example of a simulated J-curve.

Table 5
Example of a simulated J-curve

This table shows an example of J-curve drawing before the permutation routine.

Year	1	2	3	4	5	6	7	8	9	10
Cash outflows	-19%	0%	-5%	-75%	-1%					
Cash inflows	7%	0%	26%	0%	20%	22%	3%	18%	0%	54%

With this routine, the probability to have a concentration of high cash contributions and distributions in the first two years is high. In order to correct this bias, we propose 20 random permutations of the initial simulated J-curve. Table 6 shows the same example of a simulated J-curve with permutations.

Table 6
Example of simulated J-curves with permutations of contributions and distributions

This table shows 20 random permutations of the simulated pattern of the J-curve shown in Table 5, both for cash contributions and cash distributions. This is performed in order to diversify the cases and avoid having too much-repeated concentration in the first years, due to the algorithm used.

Permutation number / Years	Cash contributions (%)					Cash distributions (%)									
	1	2	3	4	5	1	2	3	4	5	6	7	8	9	10
1	-19	0	-5	-75	-1	7	0	26	0	20	22	3	18	0	54
2	-1	-19	-5	-75	0	22	18	3	20	0	26	0	7	54	0
3	0	-19	-75	-5	-1	0	22	0	54	18	20	0	7	26	3
4	-1	-75	-19	0	-5	0	0	20	18	0	7	54	26	3	22
5	0	-5	-19	-75	-1	54	18	22	7	0	0	0	26	3	20
6	-5	-75	-1	-19	0	20	7	26	0	54	22	3	0	18	0
7	-1	0	-75	-19	-5	18	0	22	54	0	0	20	3	26	7
8	-5	-1	-75	-19	0	22	0	0	0	7	3	26	54	20	18
9	-1	-5	-75	-19	0	0	22	26	54	0	20	7	18	0	3
10	-75	-19	0	-1	-5	0	0	22	20	7	3	0	18	26	54
11	-75	-19	-5	0	-1	0	0	20	3	54	18	26	0	22	7
12	-19	0	-5	-75	-1	26	3	0	20	54	7	0	18	0	22
13	-75	-19	-5	0	-1	26	22	3	18	0	20	0	0	7	54
14	0	-5	-1	-19	-75	0	20	0	54	0	7	18	26	22	3
15	-1	0	-5	-19	-75	0	18	22	20	0	0	54	7	26	3
16	-19	-5	-1	0	-75	0	3	26	22	7	54	0	20	18	0
17	-75	-5	-19	0	-1	22	7	0	54	26	18	3	0	20	0
18	-1	-5	-75	-19	0	18	20	0	0	26	54	22	7	3	0
19	-5	-75	0	-19	-1	0	22	54	20	26	0	18	3	7	0
20	-19	-5	-75	-1	0	7	0	0	26	18	20	0	54	22	3

Table 6 shows examples where this routine does not constraint the shape of the J-curve, as well as the possibility to have a distribution higher than the contribution of the same year, in year 1. This reflects the typical situation where the fund would be using leverage, for example. This leverage can be partial or total. The fifth permutation is an example where the fund uses 100% leverage to close an investment in year 1 and exit this investment in the same year with a positive multiple. In this special case, it is not impossible to have no contribution but still a positive distribution providing an infinite IRR for year 1. We do not consider yearly IRRs in our analysis, but only IRRs on the entire cycle. Permutations #2, #7, #9, #15 and #18 show cases where exits are realized with investments in year 1 funded by partial leverage.

5.1 First results

We propose here below the results stemming from a simulation of 10'000 J-curve patterns iterations, and their permutations, with a null return on the unutilized cash portion of the commitment. Table 7 shows the average cash contribution and distribution per year.

Table 7
Average of simulated cash contributions and cash distributions
This table shows the yearly average amount for the 10'000 J-curve cash contributions and cash distributions.

Year	Contribution	Distribution
1	-20.62	15.37
2	-19.87	14.57
3	-19.71	14.98
4	-19.89	15.12
5	-19.91	14.56
6		15.06
7		14.38
8		15.11
9		15.26
10		15.59
Total	-100.00	150.00

Figure 2 shows the distribution of the 10'000 corrected IRR measures issued from the simulated J-curves. The average corrected IRR is 6.24%. More than 93% of the simulated IRR are below 10% which is the initial J-curve's IRR. The corrected IRR minimum is 3.04% whereas the maximum is 33.25%.

In order to provide more transparency on the materiality of the results, and the impact of extremes, we will now focus on special cases where simulations represent leveraged situations. These leverage cases are realistic despite the fact that the IRR computed solely on the cash flows of 2018 appears to be infinite. We firstly consider cases where permutations include no contribution and distribution on year 1. These situations cover 7% of the 10,000 simulations. The average IRR on theses case is 7.68% with minimum and maximum values of 3.36% and 33.25%, respectively. In these cases, only 2.04% on the total show what could be considered as extreme case. We define extreme cases where the distribution is higher than 30, i.e. where an investment of 20 is financed by leverage and returns a MOIC of 1.5. An amount of 20 is deemed an arbitrary reasonable leverage with respect to the overall commitment of 100. The average corrected IRR on these extreme cases is 10.68%, which interestingly shows that producing a corrected IRR close to the "promised" one (10.01%, as assumed in the original J-curve) could require such a special investment case.

Secondly, we consider cases where permutations include no contribution neither in 2018 nor in 2019, with distributions starting as from 2018. These cases represent 1.96% of the total number of simulations with an average corrected IRR of 8.13%, with minimum and maximum values respectively of 3.75% and 33%.

We finally consider cases where permutations include no contribution neither in 2018, 2019 nor in 2020, with distributions starting as from 2018. These cases cover 0.7% of the total number of simulations with an average corrected IRR of 8.9%, with minimum and maximum values respectively of 4.32% and 33%.

Figure 2
Computed corrected IRR from simulated J-curves
This figure plots the sorted corrected IRR measures for the 10'000 simulated J curves, obtained with a null return on the unutilized cash portion of the commitment. The corrected IRR lies in the range [3.04%, 26.96%] with the distribution per quantile of the simulation shown here below.

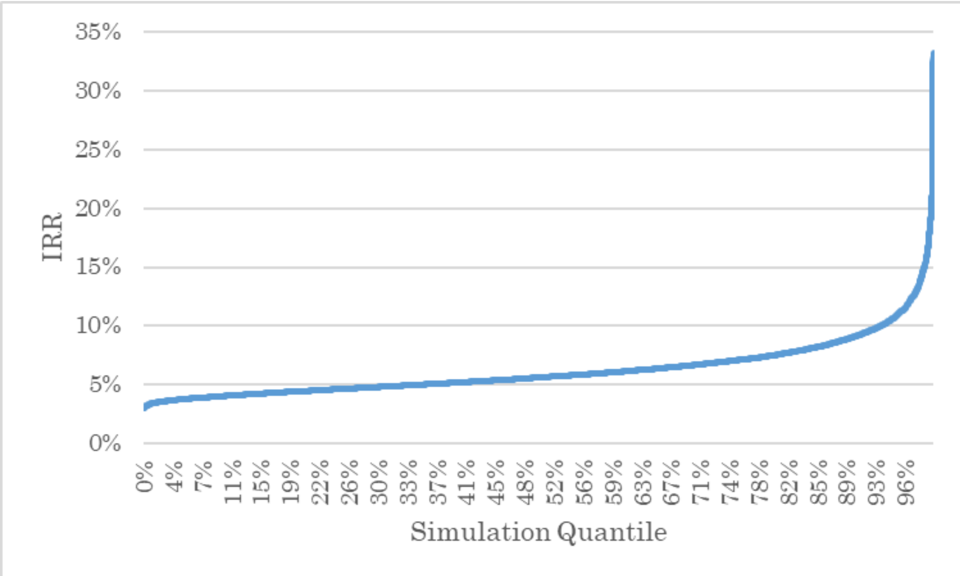

Table 8
Corrected IRR distribution metrics
This table shows descriptive statistics on corrected simulated IRR.

IRR bucket	% of total simulations	Average IRR	Median IRR
[3.04%;4%[	9.07%	3.71%	3.75%
[4%;6%[	49.05%	4.98%	4.97%
[6%;8%[	25.60%	6.84%	6.78%
[8%;10%[	9.5%	8.85%	8.8%
[10%;12%[	3.6%	10.83%	10.78%
[12%;33.25]	3.15%	15.19%	14.03%

Table 8 shows that 83.72% of simulated IRRs are below 8%. The most represented class is 4%-6%, which covers almost 50% of the simulations. The IRR distribution shows the IRR sensitivity to the J-curve shape.

5.2 J-curve shape's impact on IRR with no deposit return

We analyze in this section the link between the previously simulated J-curves and the corresponding corrected IRR, considering first there is no return from the deposit, or that the deposit rate is null.

It is firstly interesting to see the link between the J-curve shape and the corrected IRR. In order to analyze which J-curve type leads to a given level of corrected IRR, six IRR "classes" are arbitrarily defined, each representing a determined IRR range, and the average J-curve is then drawn for each class. Figure 3 shows the average contribution and distribution for the 5-year and 10-year time-horizon respectively, for each of these IRR classes. The choice of the average as a statistical metric is only for a visual point of view. The high dispersion (in terms of standard deviation) of the J-curve shapes for the whole set of 10,000 simulations and within each corrected IRR class are not in line with a single modal distribution. The range of curves is not clustered tightly around the average due to the close relationship between the corrected IRR and the J-curve shape as it will be explained below.

Figure 3
Average cumulated contributions and distributions from simulations

This figure shows the average cumulated contributions (over the first 5 years) and separately the distributions (over the 10-year duration) for each corrected IRR class (six IRR ranges are proposed). For the purpose of the visualization, both components (the stream of contributions and the stream of distributions) are shown separately.

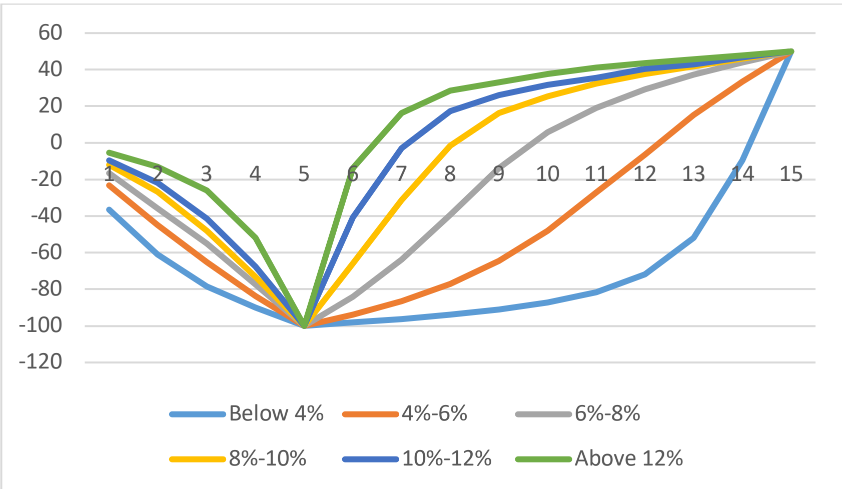

Figure 4 shows the average J-curve for each IRR class, having netted the contributions and distributions from the previous illustration. The classical J-curve shapes for the classes "Below 4%" and "4%-6%" present a time delay between contributions and distributions in line with traditional PE fund management: investments are made by GP during the whole investment period and exits are realized within a reasonable time delay.

The class "6%-8%" shows a quasi-flat J-curve shape meaning the distributions by GP are quite contemporaneous to contributions. This can represent the case of a PE manager that realized a 1.5 MOIC exit with short-duration investments.

The remaining three highest IRR classes offer a curve without the J shape meaning the time distribution of cash contributions and distributions have a "non-traditional" form.

Figure 4
(Net) J-curve for each simulated IRR class
This graph shows the J-curve by netting average cumulated contributions (5 years) and distributions (10 years) for each IRR class.

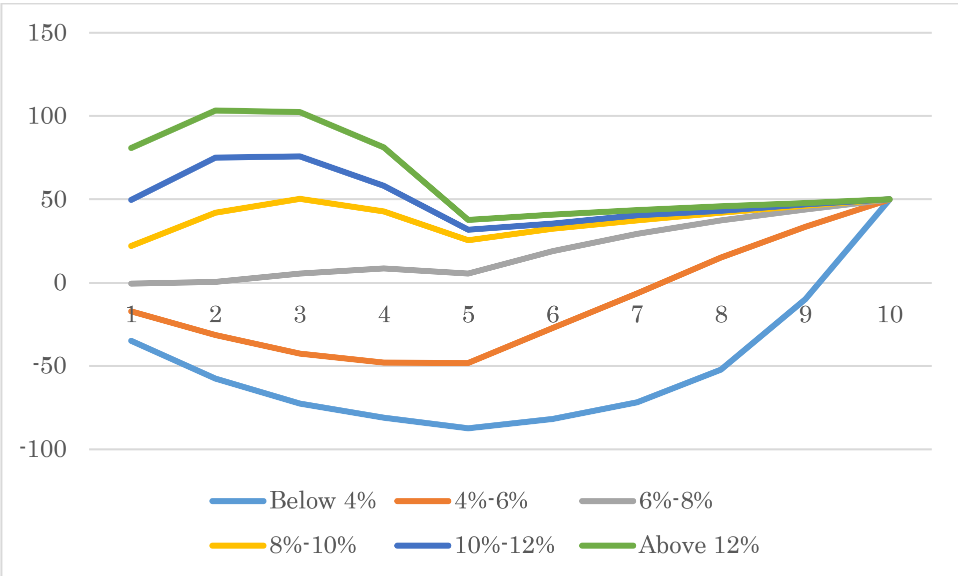

In order to show analytically the link between the J-curve shape and the corresponding corrected IRR, we propose a metric consisting in the difference between the weighted average delay of contributions and the weighted average delay of distributions, namely:

$$Weigthed\ delay\ difference\ (WDD) = \frac{\sum_{t=1}^{n} t\ Contribution_t}{\sum_{t=1}^{n}\ Contribution_t} - \frac{\sum_{t=1}^{m} t\ Distribution_t}{\sum_{t=1}^{m}\ Distribution_t}$$
$$= \sum_{t=1}^{n} w_{C_t}\, t - \sum_{t=1}^{m} w_{D_t}\, t$$

where t represents the year of occurrence of a given payment, $Contribution_t$ and $Distribution_t$ represent the payment flow, respectively for contributions and distributions, in year t. The total number of years being different for contributions and distributions, we represent them with variables n and m respectively. w_{C_t} and w_{D_t} are the resulting weights for contributions and distributions.

In the illustrative case presented above, the average delay of contributions weights the yearly cash contribution by its year of occurrence from 1 to 5, and then the sum is divided by the total contribution of 100. The average delay of distribution weights the yearly cash distribution by its year of occurrence from 1 to 10, the sum is divided by the total distributions of 150. The metrics of interest is the difference between theses two delays. It provides a cash at work duration for LP's. The bigger is the time gap between contributions and distributions, higher would be the cash at work duration. For the initial J-curve presented is figure 1, the delays difference is 4.28 years with an average contribution delays of 2.75 years and a average distribution delays of 7.03 years.

Figure 5
Link between delays difference and corrected IRR
This graph shows the 10,000 simulated corrected IRR and their associated delays difference.

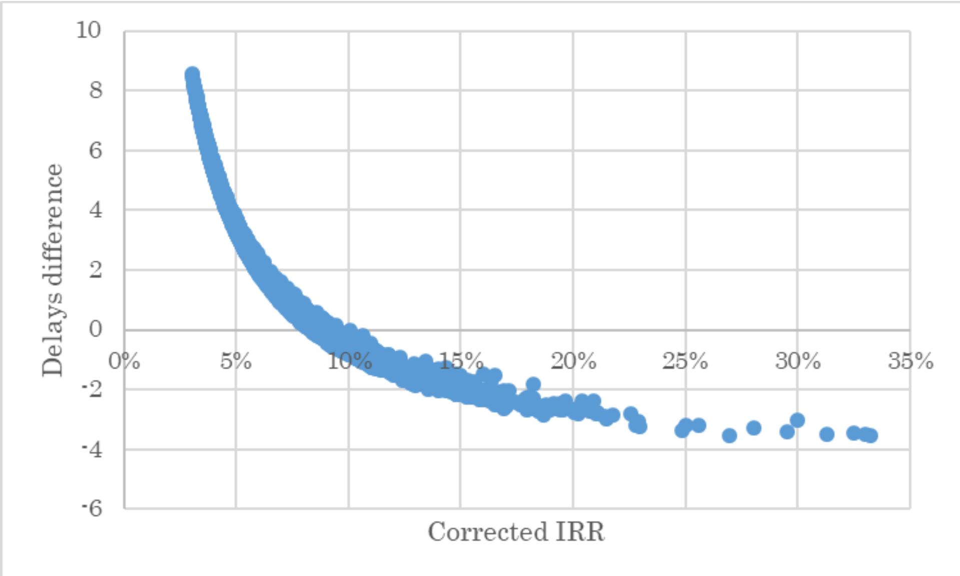

Figure 5 shows a clear visual negative relationship between the corrected IRR and the delays difference (the WDD metric presented above). The higher delays difference brings a lower corrected IRR. The correlation between both metrics is -0.88. The lowest IRR comes with a delays difference of 8.56 years. Every delays difference are positive for IRR higher than 8.35%. Higher IRR than 8.35% occur with a negative delays difference. Figure 6 to 11 propose the relationship for each IRR bucket. Figure 7 is interesting as it represents the "mode" situation that covers close to 50% of simulation cases. For this bucket, the delays difference ranges between 1.92 and 5.43 years which can be considered as reasonable levels. Figure 10 shows the bucket from 10% to 12%, including the initial J-curve's IRR. Most of IRR measures close to 10% and higher show a negative delay difference. This means that the 10% target IRR is achieved only when the time-weighted average of contributions on the 5-year period is shorter than the time-weighted average of distributions on the 10-year period. This means that the GP would realize a return on investment on a shorter period than the time he has waited until calling the capital from LPs. There is no link between the time the GP will take to find the appropriate investment and the time he will create value and exits by realizing a multiple. If the average delay of contributions would be assumed to be 2.75 years as in the initial J-curve, the GP has to exit investments in a delay below or equal to 2.75 years and with an average MOIC of 1.5. This situation implies a yearly 15.89% return on equity average on all investments, which seems relatively optimistic.

Figure 6
Link between delays difference and corrected IRR
This graph shows the simulated corrected IRR and their associated delays difference for the bucket [3.04%;4%[

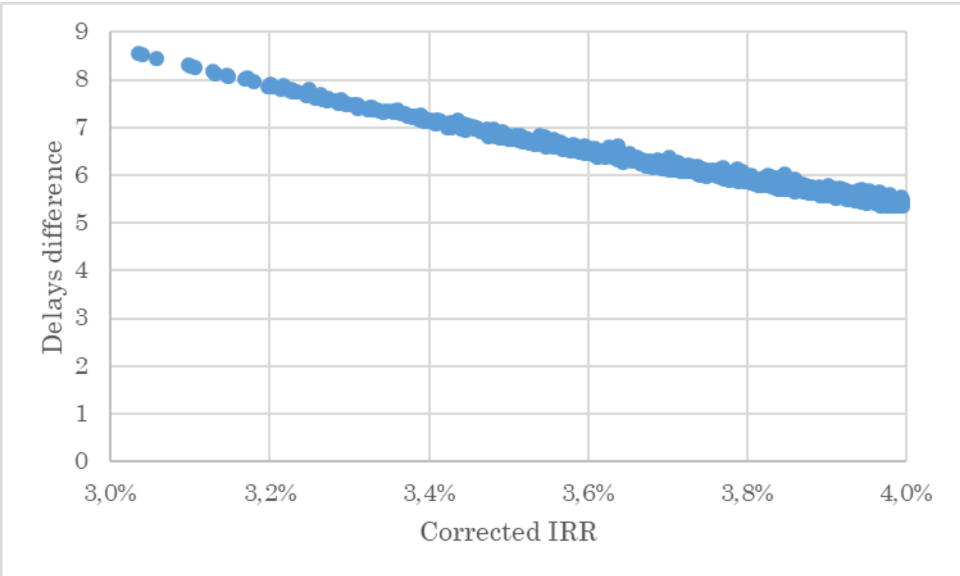

Figure 7
Link between delays difference and corrected IRR

This graph shows the simulated corrected IRR and their associated delays difference for the bucket [4%;6%[

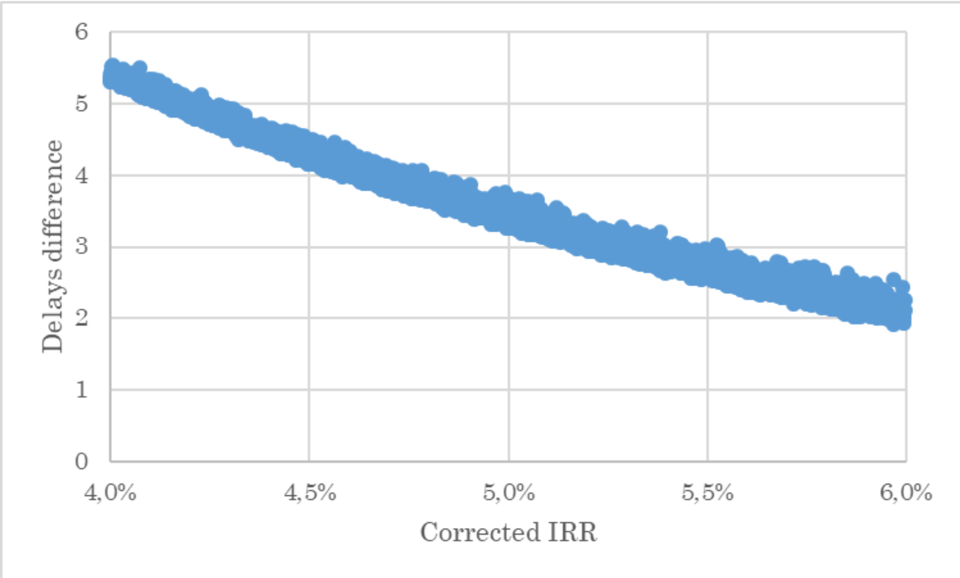

Figure 8
Link between delays difference and corrected IRR

This graph shows the simulated corrected IRR and their associated delays difference for the bucket [6%;8%[

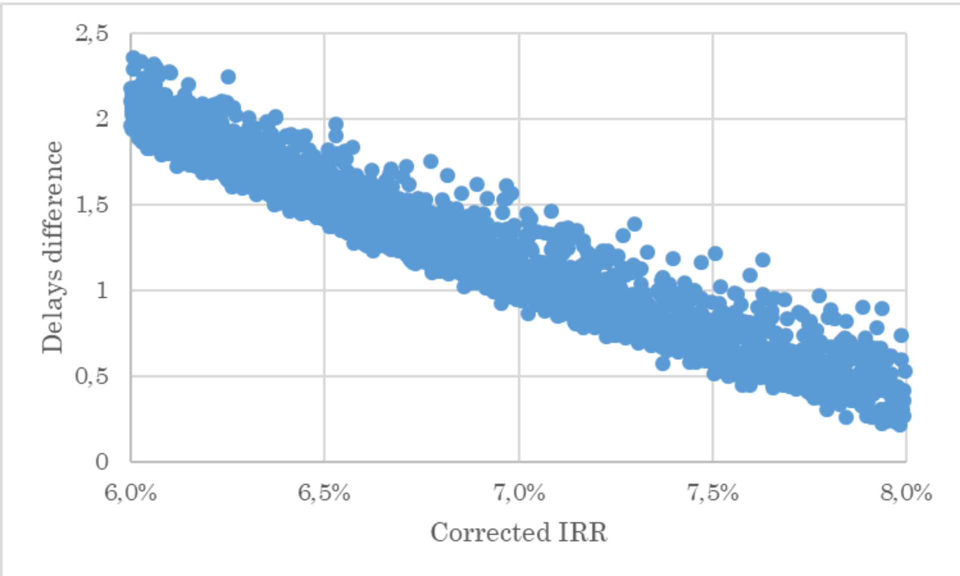

Figure 9
Link between delays difference and corrected IRR
This graph shows the simulated corrected IRR and their associated delays difference for the bucket [8%;10%[

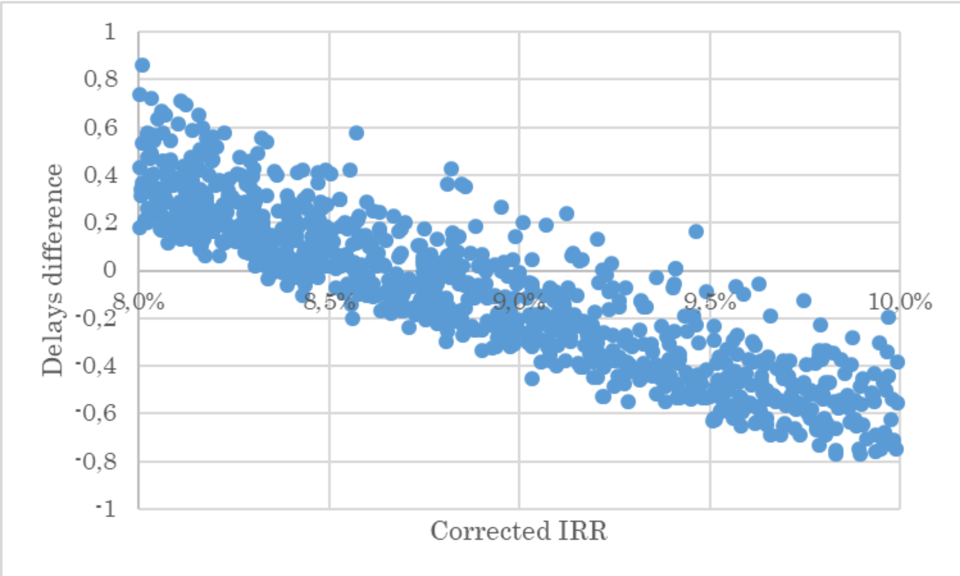

Figure 10
Link between delays difference and corrected IRR
This graph shows the simulated corrected IRR and their associated delays difference for the bucket [10%;12%[

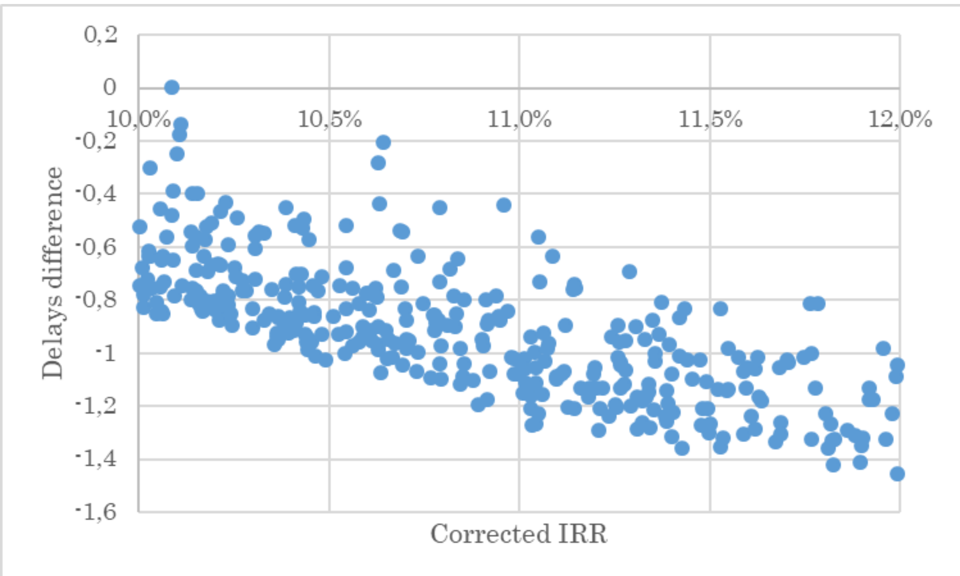

Figure 11
Link between delays difference and corrected IRR
This graph shows the simulated corrected IRR and their associated delays difference for the bucket [12%;33.25%[

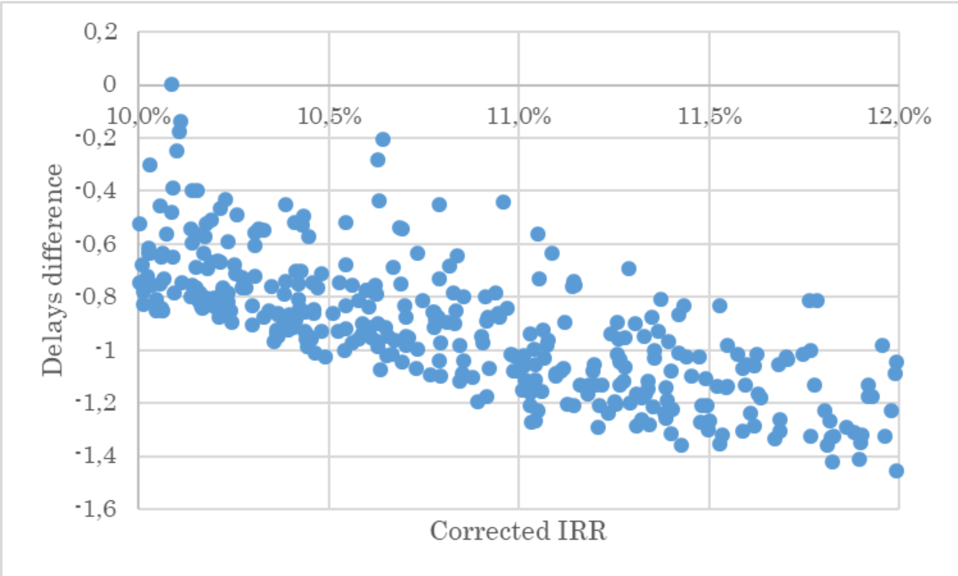

5.3 Results with the integration of the cash deposit's revenues

The previous section shows the impact of the J-curve shape in the computation of the IRR for LPs, as we assume that these LPs made 0% return on the period between the commitment and the effective capital call by the GP, to focus on the primer effect. In this section, we allow for a higher rate of return on the unutilized cash fraction. We assume that the capital call amount is known at the beginning of the year and the non-called amount is invested in a 1-year deposit. The return on the cash deposit is integrated in the cash flow stream. The cash contribution appears to be decreased by the deposit gain.

Table 8
Example of a simulated contribution with a positive unutilized cash return
This table shows a contribution example with deposit rates ranging from 0% to 4%, for a given simulation between years 1 and 5.

Year	1	2	3	4	5	Total
Cash outflows without deposit	-19.00%	0.00%	-5.00%	-75.00%	-1.00%	-100.00%
Cash outflows with 1% rate deposit	-18.19%	0.81%	-4.24%	-74.99%	-1.00%	-97.61%
Cash outflows with 2% rate deposit	-17.38%	1.62%	-3.48%	-74.98%	-1.00%	-95.22%
Cash outflows with 3% rate deposit	-16.57%	2.43%	-2.72%	-74.97%	-1.00%	-92.83%
Cash outflows with 4% rate deposit	-15.76%	3.24%	-1.96%	-74.96%	-1.00%	-90.44%

Table 8 details the previous example of cash contributions including the deposit for different interest rates, ranging from 1% to 4%.

Table 9
Corrected IRR distribution for different deposit rates
This table shows some descriptive statistics on corrected IRR with different deposit rate

Deposit rate	**0%**	**1%**	**2%**	**3%**	**4%**
Min	3.04%	3.05%	3.07%	3.08%	3.10%
Max	26.96%	30.95%	35.74%	41.95%	58.62%
Average	6.20%	6.59%	7.02%	7.47%	7.98%
% below 10%	93.27%	90.67%	87.88%	84.58%	80.69%

Table 9 shows some descriptive statistics on the corrected IRR measure for different deposit rates of the unutilized cash. The impact of deposit rates on average corrected IRRs is high, with results spread between 6.2% and 7.98%. Moreover, this table shows that the PE IRR does not even reach the expected 7.5% historical average return for public equities mentioned above, below a 3% deposit rate. Even with a 4% rate, more than 80% corrected IRRs are below 10%, the normal IRR with the assumed J-curve shape presented originally in table 1.

Figure 5
Computed corrected IRR from simulated J-curves and different deposit rates
This graph plots the sorted IRR for the 10'000 simulated J-curves, together with deposit rates ranging from 0% to 4%. The x-axis represents a quantile of the 10'000 simulations.

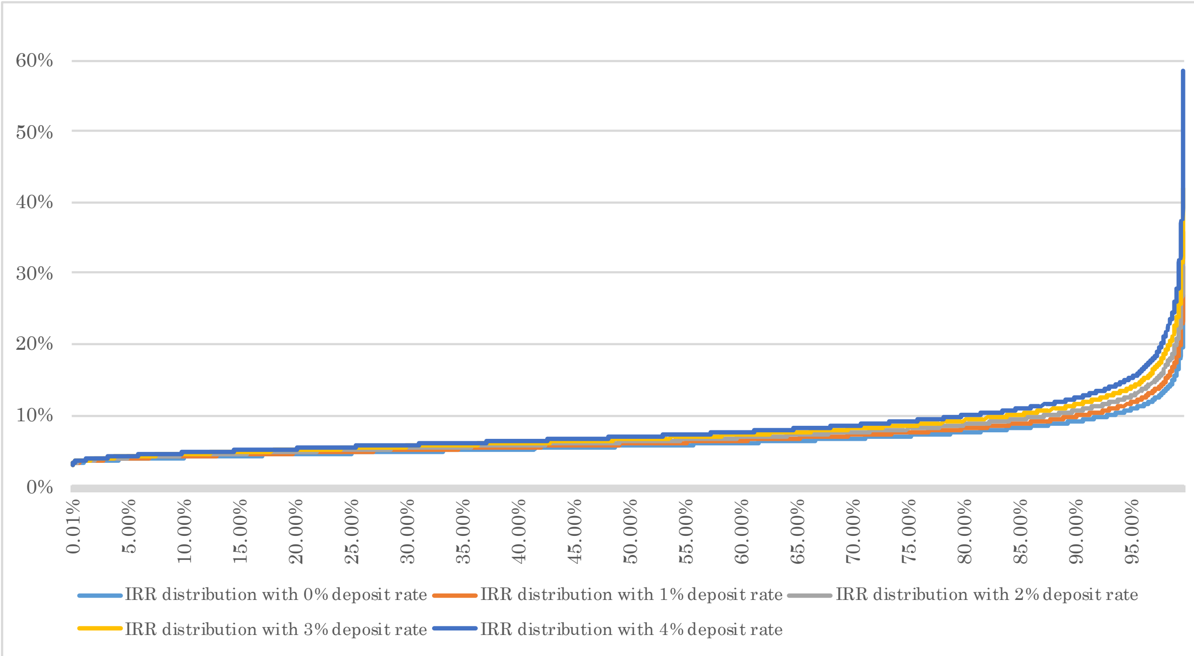

6 Conclusion

This paper revisits the private equity performance considering the asset allocation discontinuity created by the way private equity funds work, focusing on the real time distribution of the cash transfer between GPs and LPs. The main existing literature relies on private equity performances as if they were a "stand alone" investment and not in a multi-asset allocation framework. Our paper highlights the importance of the J-curve shape in the private equity performance as it is considered as a block of the whole asset allocation. During the investment period, the non-called committed cash is not allocated by investors to the risky asset class and is "locked", invested at low deposit rates. Our paper analyses the impacts of J-curve shapes and low rates environment on private equity performance. We conclude that the promised private equity over-performance on public stocks should be deeply challenged considering these effects on the asset allocation and the high fraction of cash as a consequence of the private equity commitment. The 300 basis points promised risk premium above that of listed equities, which should bring private equity performance close to 10% seems very optimistic to capture for LPs, knowing that they must integrate this block in a multi asset-class framework.

The private equity commitment practice gives the possibility to the GP to balance from cash to private equity whereas they report their performances only during the period on which the committed amount is really invested. This bring a problem at several levels. Firstly, it overestimates the private equity performance when considering multi-asset class investors as in the modern portfolio theory. Secondly the alignment of interest between GPs and LPs should be reviewed under this consideration. The IRR reported by the GP is the source of computation for the carried interest. From the GP's point of view, the IRR is not dependent on the timing of its investment pattern. Thus the carried interest paid by the LPs does not depend on the opportunity cost "paid" by LPs. We could imagine a solution where the GP mandate is larger than only private equity with the possibility to invest in public equities during the investment period before the capital is called. Then the GP would have to manage the equity market risk during the period before he proceeds to invest.

Finally, in the case of portfolios of high net worth individuals (HNWI) that display a multi-generational layering of investments in PE, the effect of the pre-investment committed capital might be "amortised" through the various commitments. Given this, this type of investor in that context might be less sensitive to the first year return and the liquidity urgency on a single PE case. This paper provides a starting point to open the debate for further research.

References

Ang. A.. Chen. B. Goetzmann N and Phalippou L (2013). Estimating Private Equity Returns from Limited Partner Cash Flows. Working paper Columbia University and NBER.

Appelbaum. E. & Batt. R. (2014). Private Equity at Work: When Wall Street manages Main Street. New York: Russel Sage Foundation Press.

Appelbaum. E.. & Batt. R. (2016). Are Lower Private Equity Returns the New Normal? (No. 2016-10). Center for Economic and Policy Research (CEPR).

Barber. B. M.. & Yasuda. A. (2017). Interim fund performance and fundraising in private equity. Journal of Financial Economics. 124(1). 172-194.

Gredil. O. B. Griffiths. and R Stuck. 2014. Benchmarking Private Equity The Direct Alpha Method. Unpublished working paper.

Harris. R. S. Jenkinson. T.. & Kaplan. S. N. (2014). Private equity performance: What do we know?. The Journal of Finance. 69(5). 1851-1882.

Henly .S. (2013). PME Benchmarking Methods. Pitch Book.

Jegadeesh. N.. Kräussl. R.. & Pollet. J. (2009). Risk and expected returns of private equity investments: evidence based on market prices (No. w15335). National Bureau of Economic Research.

Kaplan. S. N.. & Schoar. A. (2005). Private equity performance: Returns. persistence. and capital flows. The Journal of Finance. 60(4). 1791-1823.

Korteweg. A.. & Sorensen. M. (2017). Skill and luck in private equity performance. Journal of Financial Economics. 124(3). 535-562.

Long. A. M.. & Nickels. C. J. (1996). A private investment benchmark. AIMR conference on Venture Capital Investing. Working paper.

Mustad. H. J.. & Pedersen. O. O. (2018). Risk-adjusted index for private equity evaluation based of predictive firm characteristics: An empirical study of the public firms taken private by financials buyers from 1997 to 2017 in the US (Master's thesis).

Public Market Equivalent (PME) Benchmarking (2015). Prequin Special Report. xx pages.