

The Performance of Private Equity Funds

Ludovic Phalippou

University of Amsterdam

Oliver Gottschalg

HEC Paris

The performance of private equity funds as reported by industry associations and previous research is overstated. A large part of performance is driven by inflated accounting valuation of ongoing investments and we find a bias toward better performing funds in the data. We find an average net-of-fees fund performance of 3% per year below that of the S&P 500. Adjusting for risk brings the underperformance to 6% per year. We estimate fees to be 6% per year. We discuss several misleading aspects of performance reporting and some side benefits as a first step toward an explanation. (*JEL* G23, G24)

Capital committed to private equity funds increased from \$5 billion in 1980 to \$300 billion in 2004. According to the press, this spectacular growth is primarily due to a widely held belief of high past performance (Appendix). However, and despite the fact that private equity funds represent a major class of financial assets, we are still missing a comprehensive account of their historic performance.

The objective of this study is to estimate the performance of private equity funds both net-of-fees and gross-of-fees. We find that the average private equity fund underperforms the S&P 500 Index net-of-fees by about 3% per year and overperforms that index gross-of-fees by 3% per year.

The typical private equity fund partnership contract stipulates that funds have a life of 10 years, with a possible extension of 3 years. Throughout their life, funds call capital from their investors up to the level committed at inception and distribute dividends to their investors. On a quarterly basis, they self-report the net asset value (NAV) of the ongoing investments. To minimize the impact of these NAVs, we select only funds that are more than 10 years

We thank Gemma Postlethwaite, Jesse Reyes, and Thomson Venture Economics for making this project possible through generous access to their databases. We are especially grateful to Antoinette Schoar, who helped us to replicate her results. We also thank an anonymous referee, Bernard Dumas, Alexander Groh, Ulrich Hege, Raj Iyer, Tim Jenkinson (WFA discussant), Ron Kaniel, Robert Kosowski, Josh Lerner, Dima Leshchinskii, Alexander Ljungqvist, Florencio Lopez-de-Silanes, Pedro Matos, Andrew Metrick, Eric Nowak, Manju Puri, Anna Scherbina, Clemens Sialm (EFA discussant), Per Stromberg, Charles Trzcinka (FFC discussant), Annette Vissing-Jorgensen, Michael Weisbach (the editor), Maurizio Zollo, and audiences at INSEAD, Toulouse BS, UC Dublin, SIFR-Stockholm, EFMA meeting, CEPR meeting, EFA meeting, Frontiers of Finance Conference, Gutmann Center Symposium, Inquire-UK meeting, and WFA meeting for their comments. Financial support from the R&D Department at INSEAD, Fondation HEC, and the Wharton-INSEAD Alliance is gratefully acknowledged. Send correspondence to Ludovic Phalippou, University of Amsterdam, FEE, Finance Group, 11 Roeterstraat, 1018 Amsterdam, The Netherlands, telephone: (31) 20-525-4153. E-mail: l.phalippou@uva.nl.

old and funds with no recent sign of activities. This ensures that our sample consists of basically liquidated funds for which final NAV can be written off (see the text for further discussion and justification) and that performance can then be evaluated only on actual cash flows to/from investors. Compared to the standard approach of treating NAVs as accurate estimates of fund market value, writing them off decreases the average fund performance by 7% in terms of the profitability index (*PI*).¹ Therefore, our first finding is that funds that have reached their normal liquidation age still report large accounting valuations for ongoing investments, which biases performance estimates upward.

A second correction that we propose concerns performance weights. Standard practice is to weight each fund by the total capital committed by the investors at inception. However, the present value of money invested differs from capital precommitted as funds vary in both the speed at which they call capital and the fraction they actually call. If poorly performing funds invest *more* slowly, then capital-committed-weighted performance is *downward* biased (and vice versa). We thus weight fund performance by the present value of investments and find that this choice reduces the average *PI* by 2% compared to a standard capital-committed weighting. Therefore, our second finding is that standard aggregation choices bias performance estimates upward.

The sample we use is similar to that used for industry benchmarks and Kaplan and Schoar's (2005) study. These data are mainly provided by large private equity investors on a voluntary basis. The question of sample selection bias arises as these investors might not be representative. Lerner, Schoar, and Wong (2007) document large differences in the performance achieved by different types of investors in private equity.

In addition to the "base sample" used above, we use VentureXpert—a commercially available dataset—which contains information on investment exits (e.g., IPO, bankruptcy). Four hundred and seventy-six funds are present in Venture-Xpert, but not in our "base sample." For these "additional funds," performance is not available (by definition), but we observe their fraction of successfully exited investments (IPO or sale), which is a widely used proxy for performance (e.g., Hochberg, Ljungqvist, and Lu, 2007).

The unique aspect of our data is that we have a link between this commercial dataset and the "base" sample. By comparing success rates between additional funds and "base" funds, we learn about the extent of a sample selection bias in the dataset that is widely used for performance.

We find that the "additional" funds have a 45% success rate while "base" funds have a 50% success rate. The spread is statistically significant and a 5% spread has been deemed economically significant in the literature (e.g., Hochberg, Ljungqvist, and Lu, 2007). This spread translates into an average *PI* in the extended sample that is 4% below that of the base sample. Therefore, our

¹ The profitability index (*PI*) is the present value of the cash flows received by investors divided by the present value of the capital paid by investors. The discount rate is the realized S&P 500 rate of return; $PI > 1$ indicates outperformance.

third finding is that a commonly used dataset for private equity performance contains funds that perform better than average.

A “standard” approach to compute fund performance (i.e., using base-sample funds only, treating NAV as correct and weighting by capital committed) leads to an average *PI* of 1.01, which means a slight outperformance of the S&P 500. After our three adjustments mentioned above (i.e., considering additional funds, writing off NAV, and weighting by the present value of invested capital), we find an average *PI* of 0.88, which means that private equity funds have lost 12% of the value invested in present value terms. To obtain a more tangible measure of performance, we compute an alpha and find that it is -3% per year on average (with beta set to 1).

A fourth issue is that of risk adjustment. Given the high leverage used by buyout (BO) funds and the nature of venture capital (VC) investments (resemble small growth companies, which typically exhibit high systematic risk), the assumption of a beta as 1 is likely to overstate relative performance. Using “industry/size-matched cost-of-capital,” we find a risk-adjusted *PI* of 0.75 for BO funds and 0.77 for VC funds. These correspond to an alpha of about -6% per year. Then, our fourth finding is that risk adjustments decrease performance substantially.

To gauge the value added for private equity fund managers, we calculate the performance gross-of-fees by applying a typical fee structure on the observed time series of cash flows. Our fifth finding is that the fee is more than 25% of the value invested (6% per year). In addition, we observe that two-thirds of the fee comes from management fees while a third comes from incentive fees.

Our sixth finding is the quantification of the value of different fee arrangements. Previous literature (Gompers and Lerner, 1999; and Metrick and Yasuda, 2007) relies on simulated cash flows to gauge the total value of fees and their sensitivities to different arrangements. As assumptions about cash flow timing and amount are intimately related to the value of the fees, having *actual* cash flows represents a clear advantage. Using this feature of our data, we show that moving from 2% to 2.5% management fees translates into a 1.3% decrease in alpha, that reducing the amount of fees in the postinvestment period (from 2% of capital committed to 1%) is only worth 0.6% in terms of alpha, etc.

Our analysis is about the “average” fund, and large discrepancies exist across funds. Our seventh finding is that top quartile funds outperform the S&P 500 and that there is evidence of performance persistence that is robust to using different samples, control variables, and controlling for sample selection bias. We find that prior fund performance actually subsumes all other fund characteristics in explaining fund performance.

Determining average past performance is important *per se* for direct applications, such as portfolio benchmarking and optimal asset allocation, but it also fits in a large literature on the performance of professional asset managers. Indeed, financial economists have long been interested in such an issue to assess both market efficiency and investor sophistication. In the rational model of financial

intermediation by Berk and Green (2004), financial intermediaries provide zero “alpha” to their investors and capture a rent that is commensurate with their abilities. Empirically, the mutual fund literature has reached a quasi consensus showing that aggregate abnormal performance is negative after fees and positive before fees (e.g., Wermers, 2000). For hedge funds, Fung et al. (2006) show that average net “alphas” are zero. Our results for private equity funds indicate that the rent of the average private equity fund may be excessive. In addition, our results complement the finding of Moskowitz and Vissing-Jorgensen (2002) of low returns on entrepreneurial investments and bring additional evidence for the existence of what they call the “private equity premium puzzle.”

The rest of the paper is structured as follows: Section 1 presents the base sample and Section 2 covers the methodological corrections we operate. Section 3 is dedicated to correcting the sample selection bias. Section 4 presents final estimates of net-of-fees and gross-of-fees performance. Section 5 contains further tests and analyses. Section 6 discusses potential explanations for the observed underperformance and Section 7 briefly concludes.

1. Sample and Data

This section provides background information and reports the content and source of our data and the selection scheme for our “base” sample. We also provide descriptive statistics, describe how performance is typically computed in the literature, and replicate standard performance estimates.

1.1 Background

Investors in private equity funds are principally institutional investors, such as endowments and pension funds. These investors, called limited partners (LPs), commit a certain amount of capital to private equity funds at fund inception. Fund managers—called general partners (GPs)—search out investments and “calls” for money when needed, up to the amount committed by LPs. When a divestment occurs, the GP distributes the proceeds to its LPs (minus fees). The timing of these cash flows is typically unknown *ex ante*. A fund typically has a life of 10 years, which can be extended to 13. Stakes in private equity funds are basically nontradable and funds self-report quarterly an NAV that reflects the value of ongoing investments.

1.2 Data source and sample selection: Base sample

Our base sample is derived from a cash flow dataset maintained by Thomson Venture Economics (TVE). It contains the amount and date of all cash flows (to/from investors), as well as the quarterly NAVs from 1980 to 2003. Cash flows are net of fees as they include all fee payments to GPs, as well as carried interest. The dataset also contains information on fund characteristics (size, investment focus, regional focus). Earlier versions of this dataset have been used by Kaplan and Schoar (2005) and Jones and Rhodes-Kropf (2003).

Table 1
TVE cash flow dataset coverage

Year	Universe						Coverage TVE cash flow dataset (%)					
	VC + BO		VC		BO		VC + BO		VC		BO	
	N	Size	N	Size	N	Size	% N	% Size	% N	% Size	% N	% Size
1980–1984	746	29,087	662	15,529	84	8,752	0.31	0.64	0.32	0.68	0.30	0.65
1985–1989	1,043	125,486	737	41,646	306	78,809	0.46	0.76	0.46	0.72	0.46	0.79
1990–1993	866	22,207	517	21,457	349	48,342	0.33	0.49	0.32	0.57	0.35	0.57
1994–2002	7,393	864,735	5,339	366,852	2,054	677,710	0.19	0.55	0.16	0.42	0.27	0.60
1980–1993	2,655	176,780	1,916	78,632	739	135,903	0.38	0.70	0.37	0.67	0.39	0.70

This table reports the overall number of private equity funds and the sum of committed capital in millions of 2003 U.S. dollars (size). The universe is obtained by taking the maximum figures from two sources: (1) “TVE investment” dataset we received from TVE in 2003 and (2) VentureXpert “fund-raising statistics” accessible via VentureXpert.com. This table also reports the coverage of the TVE cash flow dataset relative to the universe. We distinguish between buyout (BO) and venture capital (VC) funds.

This dataset is also the basis of the industry standard performance benchmark published by TVE and various industry associations.

An accurate performance estimate is only possible for sufficiently mature funds because there is no market value for ongoing investments. We thus select funds that have reached their “normal” liquidation date (i.e., funds that are older than 10 years (raised between 1980 and 1993). We eliminate funds smaller than \$5 million (1990 U.S. dollars) and funds with cash-flow activities over the last six quarters of our observation period (as in Kaplan and Schoar, 2005). This “filter” should exclude both evergreen funds (i.e., funds with infinite life) and funds that have been granted an extension (hence, cannot be considered liquidated after 10 years). This leaves us with 852 funds that we consider “effectively” liquidated (half of them are also “officially” liquidated according to TVE). These funds are referred to as our base sample.

The cash flow dataset of TVE is considered to be the most comprehensive source of performance information on private equity funds. In Table 1, we show an estimate of the universe of funds and the coverage of the TVE cash flow dataset.² We find that the average coverage is 66% (in terms of size) for both VC funds and BO funds and that larger funds are overrepresented in our sample. We also note that the total amount raised between 1980 and 1993 represents only 18% of the amount raised between 1980 and 2002. It is therefore important to keep in mind that our performance estimate will evolve over the subsequent years (see also Section 5.5). Nonetheless, our approach and results are expected to remain relevant.

² As private equity firms are not required to disclose fund-raising information, a full census of their activity is unavailable. To define the “universe,” we take the highest numbers out of two major sources of information. The first source is TVE fund-raising statistics, which are based on annual surveys. It is available at the Analytics-Commitments section of VentureXpert.com. Both the United States and the EU Private Equity Association base their fund-raising statistics on these data. The second source is the TVE investment dataset, which contains information on private equity funds and their underlying investments. We use the 2003 version we received from TVE. Summary statistics from the investment database are available in the Profiles-Funds section of VentureXpert.com.




Figure 1

Multiple distribution by fund type

This graph shows multiple quartiles per fund type (venture capital, VC; buyout, BO; US; and European, EU). Two multiples are used: (1) total value over paid-in capital (TVPI), which is the sum of distributions plus latest NAV, divided by the sum of all takedowns, and (2) distributed over paid-in capital (DPI), which is the sum of distributions divided by the sum of all takedowns. The fourth quartile is capped at the 95th percentile.

1.3 Standard performance estimate for base-sample funds

1.3.1 Multiples. The widespread belief of good past performance of private equity mentioned in the Appendix is often based on so-called multiples. The first of these is the “total value over paid-in capital” (TVPI); it is defined as the sum of all cash distributions plus the latest NAV, divided by the sum of all takedowns. The second of these is the “distributed over paid-in capital” (DPI); it is defined as the sum of all cash distributions divided by the sum of all takedowns.

Figure 1 reports descriptive statistics. Median TVPI is 1.73 for BO funds with a U.S. investment target (U.S.-focused), 1.57 for U.S.-focused VC funds, 1.56 for EU-focused BO funds, and 1.46 for EU-focused VC funds. This confirms that U.S. funds typically perform better than EU funds (see also Section 5.4). We also observe that the interquartile range is very large. The 75th percentile fund returns are about double those of the 25th percentile fund for each subgroup (U, EU, VC, and BO).

When NAVs are written off, we find that the median multiple (DPI) is 1.57 for U.S.-focused BO funds, 1.42 for U.S.-focused VC funds, 1.21 for EU-focused BO funds, and 1.12 for EU-focused VC funds. The relative size of the final NAV is thus larger for EU-focused funds. This might be due to lower accounting standards. The interquartile range increases a lot when writing off NAVs. For U.S.-focused VC funds, the 75th percentile fund returns almost three times as much as the 25th percentile fund. For EU-focused VC funds, the

75th percentile fund returns more than four times as much as the 25th percentile fund. This is consistent with poorly performing funds being more aggressive with their NAVs (see also Section 2.2.2).

1.3.2 Kaplan and Schoar sample. Kaplan and Schoar (KS, 2005) also report fund performance and thus provide a comparison for our study. When we restrict our sample to the subsample described in KS, we obtain similar IRR (19% for us and 18% for KS) and similar number of observations (851 for us and 905 for KS). After applying the “six-quarter of inactivity” filter described above (same as used in KS), however, we find an average *PI* of 1.01 (1.05 for KS), an average IRR of 15% (19% for KS), and 599 observations (746 for KS). Given the nature of the data, it is difficult to match their sample with precision and some discrepancies are present. Nonetheless, our study focuses on the impact on performance of sample selection bias, accounting values, and fees. Hence, the starting point for performance is not of primary importance.

In addition, and more importantly, whether we restrict our sample to the subset that corresponds to the KS time period and selection criteria or whether we use our “base” sample, the average performance is the same, so discrepancies are not due to selection schemes. Indeed, our base sample differs from that of KS along four dimensions. First, their performance is “cut” as of 2001 (versus a 2003 “cut” in our study). Second, selected vintage years are 1980–1995 (versus 1980–1993 in our study). Third, selected funds are those that are inactive from June 2000 to December 2001 (June 2002 to December 2003 in our study). Fourth, selected funds are U.S.-related (i.e., either have a U.S. investment target or are headquartered in the United States versus worldwide funds in our study). In nontabulated results, we find that the 1980–1995 batch has an average performance of 1.01 both as of 2001 and as of 2003. The same holds true for the 1980–1993 batch. Hence, the first two dimensions mentioned above have no impact on performance. The third and fourth dimensions, however, change performance. The 2003 cut increases performance by 0.06 relative to the 2001 cut while adding non-U.S.-related funds reduces performance by as much. Hence, whether we work on a subsample corresponding to the KS dataset or on our base sample, our starting point for performance is the same.

2. Methodology

In this section, we operate two corrections to the standard performance estimate. We first change the weights of individual funds in our aggregate estimate. Second, we write off NAVs.

2.1 Value-invested weighting

It is a standard practice to weight each fund by its capital committed (e.g., Kaplan and Schoar, 2005). However, the value invested differs from capital committed as funds do not invest all capital upfront and vary in the speed at

which they call capital. If poorly performing funds invest *more* slowly, then capital-committed-weighted performance is *downward* biased (and vice versa). We thus weight from here onward each fund's *PI* by the present value of the stream of investments ("value invested").

The consequence of such a change can be seen in panel A of Table 2. In terms of overall average performance, it decreases the average *PI* from 1.01 (capital-committed-weighted) to 0.99. In certain vintage years, the impact is dramatic. For instance, in 1989, the capital-committed-weighted *PI* is 1.12 while the value-invested-weighted *PI* is 0.96.

2.2 NAVs

2.2.1 Relevance of NAVs. As mentioned above, funds self-report accounting valuations for ongoing investments (NAVs). Their magnitude decreases as funds approach liquidation date, but we often observe positive NAVs after year 10 (usual liquidation time), after year 13 (maximum liquidation time), and even after funds are officially liquidated. We face a tradeoff between sample size and selecting funds that are old enough for us to say something about their latest NAV. Also, selecting only funds that are officially liquidated and report zero final NAV solves the NAV problem, but it (1) reduces sample size by a third and (2) introduces a clear sample bias as better funds both liquidate faster (see Section 5.2) and report lower final NAVs (see Section 2.2.2).³

As mentioned above, we select the batch of funds that has reached its usual liquidation time. This way, we keep the sample bias to a minimum, have a reasonably large sample size, and can more easily judge NAVs.

In the literature, two different assumptions have been made concerning the treatment of final NAVs. The first and most frequent one treats the final NAV as a cash inflow of the same amount at the end of the sample time period. That is, NAVs are assumed to be an unbiased assessment of the market value of a fund (e.g., Kaplan and Schoar, 2005, and industry benchmarks). The second one writes them off (e.g., Ljungqvist and Richardson, 2003). We follow this second option as it is both transparent and reasonable for a sample of old and inactive funds like ours.

We show in Table 2 that the treatment of NAVs has a significant effect on performance. Writing them off instead of treating them as correct decreases the average *PI* by 7%. This is due to the fact that as many as half of our funds report positive final NAVs despite the fact that they have reached their usual liquidation age and have not shown recent signs of activity (e.g., cash distribution, investment, or fee collection for more than six quarters).

2.3 NAV dynamics

In this subsection, we investigate the evolution of NAVs over time in order to gain further insights into the appropriateness of writing them off.

³ The performance of these funds is about that of the S&P 500 (*PI* = 1).

Table 2
Fund performance

Year	N	Total size	PV Inv.	PI-VW	PI-PVW	IRR-VW	IRR-PVW	IRR-TPVW	Aggregated IRR	TVPI
Panel A: Performance of funds per vintage year with final NAVs, base sample										
1980	22	4,049	1,049	1.22	1.21	22.50	21.69	20.68	21.19	2.83
1981	28	1,789	746	0.75	0.74	10.74	10.59	9.68	11.77	2.08
1982	32	2,249	783	0.48	0.45	4.85	3.96	3.97	5.30	1.53
1983	64	6,390	2,421	0.85	0.82	14.32	13.23	10.73	12.29	1.99
1984	80	7,762	3,379	0.96	0.95	14.07	13.97	11.35	15.05	2.49
1985	73	5,631	2,509	1.35	1.32	22.86	21.83	13.88	23.61	2.49
1986	64	7,071	3,019	0.89	0.87	10.32	10.05	10.66	12.32	2.38
1987	107	15,376	7,326	1.02	0.98	16.30	14.68	12.02	14.01	1.97
1988	86	18,451	7,569	0.93	0.94	12.39	12.54	11.99	13.46	1.87
1989	99	11,013	6,309	1.12	0.96	15.69	12.47	9.49	13.78	1.93
1990	56	9,939	5,309	1.05	1.05	18.18	18.43	13.95	17.78	2.04
1991	49	7,137	3,491	0.96	1.00	13.43	14.83	11.77	16.54	2.00
1992	40	6,615	3,753	1.09	1.07	17.93	17.54	16.86	20.99	2.24
1993	52	6,410	3,722	1.10	1.09	15.44	15.50	12.80	18.89	2.27
Total	852	109,882	51,385							
Mean VW				1.01		15.20		12.26	15.51	2.11
Mean PVW					0.99		14.64	12.22	15.66	2.10

Table 2
(Continued)

Year	Base sample					Extended sample					Universe	
	<i>N</i>	Size	PV Inv.	PI	Alpha	<i>N</i>	Size	PV Inv.	PI	Alpha	<i>N</i>	Size
Panel B: Performance of funds per vintage year without final NAVs												
1980	22	4,049	1049	1.21	3.84	50	5,247	1,349	1.13	2.62	78	3,617
1981	28	1,789	746	0.73	−5.40	82	5,020	2,484	0.68	−7.40	125	5,259
1982	32	2,249	783	0.45	−11.64	88	4,543	2,055	0.43	−12.78	133	6,532
1983	64	6,390	2,421	0.81	−4.32	122	9,360	4,113	0.64	−7.94	196	10,550
1984	80	7,762	3,379	0.94	−1.32	129	9,704	4,721	0.80	−5.57	214	12,624
1985	73	5,631	2,509	1.31	7.68	116	8,848	4,246	1.10	2.00	188	12,408
1986	64	7,071	3,019	0.85	−2.88	94	9,493	4,196	0.84	−3.44	174	14,528
1987	107	15,376	7,326	0.96	−1.20	125	16,479	7,932	0.94	−1.59	207	33,173
1988	86	18,451	7,569	0.91	−2.40	114	20,770	8,981	0.88	−2.80	217	29,302
1989	99	11,013	6,309	0.91	−2.52	123	13,265	7,679	0.90	−2.33	257	36,075
1990	56	9,939	5,309	0.98	−0.72	78	11,953	6,322	0.98	−0.70	247	20,200
1991	49	7,137	3,491	0.82	−6.36	57	8,578	4,121	0.88	−3.82	138	16,636
1992	40	6,615	3,753	0.96	−1.56	67	12,936	7,311	0.99	−0.31	205	23,030
1993	52	6,410	3,722	0.79	−7.56	83	10,002	5,961	0.86	−4.33	276	30,660
Total, 1980—1993	852	109,882	51,385			1,328	146,196	71,470			2,655	254,594
Mean				0.92	−2.26				0.88	−2.94		
Venture Cap.	616	49,261	22,088	0.88	−2.48	1,014	72,008	37,966	0.82	−4.22	1,916	95,930
Buyout	236	60,621	29,297	0.96	−1.82	314	74,187	33,504	0.95	−1.65	739	158,664

This table shows the number of funds, fund size in real terms (million of 2003 U.S. dollars), present value of all the investments (PV Inv.), average PU (either weighted by size, VW, or present value of the investments, PVW), IRR, total value over paid-in capital (TVPI), and alpha. TVPI is the sum of distributions plus final NAV, divided by the sum of takedowns. Alpha is defined as a constant to be added to the S&P 500 return to set NPV to zero and is computed on the aggregated cash flows for each vintage year. Three IRRs are shown: one value weights each fund, one weights each fund by the product of their duration and the present value of capital invested (TPVW), and one is computed on aggregated cash flows for each vintage year. At the foot of the panels, we report the average performance across vintage years. In panel A, final NAVs are assumed to equal the market value of the fund while in panel B, final NAVs are written off. In panel B, averages are weighted by the present value of the investments.

The 462 base-sample funds that report a positive NAV represent 50% of the capital invested. We first classify these 462 funds into different categories as a function of (1) the net change in NAV between the end of 2000 and the end of 1998 (upmarket period) and (2) the net change in NAV between the end of 2000 and the end of 2003 (downmarket period).

To better capture the effect of write-ups or write-downs of NAVs, we adjust the changes in NAVs by intermediate cash flows. Hence, net change in NAV between T and $T + 1$ is defined as

$$RV_{T+1} - RV_T = R\tilde{V}_{T+1} - R\tilde{V}_T + I_T - D_T, \quad (1)$$

where (1) $R\tilde{V}_{T+1}$ denotes reported NAV at date $T + 1$, and I_T and D_T denote the investments and distributions made between dates T and $T + 1$.

Table 3 reports the proportion of funds in each category in terms of numbers, value invested, and NAVs. The largest category (cat. 1; 46% of the funds, 40% of total NAV) contains funds that have had no change in NAV in either period. We also see that they do not have any cash flow activity. This means that they have been completely inactive over at least the last 5 years. Even with NAVs treated as correct, this large group of funds exhibits the lowest performance of all categories ($PI = 0.71$ with NAVs treated as correct, $PI = 0.58$ otherwise).

The next two most important categories of funds are those that changed their NAVs between 1998 and 2000, but not since then (cat. 2 and cat. 3). In addition, they show no cash flow activities since January 2001. They represent as much as 31% of total NAVs and exhibit low performance (0.87 for cat. 2 and 0.89 for cat. 3). Given their inactivity over at least the last 3 years, the NAVs reported by these funds appear highly likely worthless and reflect so-called living dead investments.

Altogether, funds in categories 1, 2, and 3 represent 71% of the total reported NAVs and 58% of the value invested by the nonfully liquidated funds. If we write off only these most obvious living deads, and treat all the other NAVs as accurate, then we can deduce from Table 3 that the average PI would decrease by 0.05 (instead of 0.07 when all NAVs are written off).

Determining which investment is a living dead and thus to what extent accounting valuation is exaggerated (or conservative) is bound to be a subjective exercise. We argue that it is most reasonable and transparent to write off all NAVs, but it is important to bear in mind that erasing only the most obvious and treating as correct the remaining ones would affect average performance only marginally. Most of the final NAVs are rather obvious living deads.

3. Correction of Sample Selection Bias

In this section, we evaluate data on additional private equity funds to determine whether the base sample is biased toward successful funds and to correct accordingly our performance estimate.

Table 3
Changes in NAVs

	Categories		Funds in each category (out of all funds with NAV ₀₃ > 0)			% with no CF from 2000 to 2003 (in each category)			PI w/t NAV	PI with NAV
	NAV ₀₀ -NAV ₉₈	NAV ₀₃ -NAV ₀₀	% N	% PV invested	% NAV	% N	% PV invested	% NAV		
Cat. 1	=0	=0	0.46	0.33	0.40	100%	100%	100%	0.58	0.71
Cat. 2	<0	=0	0.11	0.13	0.15	100%	100%	100%	0.69	0.87
Cat. 3	>0	=0	0.06	0.12	0.16	100%	100%	100%	0.70	0.89
Cat. 4	<0	<0	0.23	0.30	0.12	43%	40%	23%	1.19	1.24
Cat. 5	>0	<0	0.11	0.07	0.06	85%	69%	85%	1.30	1.40
Cat. 6	<0	>0	0.02	0.02	0.02	80%	70%	24%	0.81	1.02
Cat. 7	>0	>0	0.01	0.03	0.09	50%	82%	92%	0.35	0.82

This table shows the evolution of the NAVs for base-sample funds that report a positive NAV at the end of our sample period. We first classify these funds into seven categories as a function of (1) the net change in NAVs between end of 2000 and end of 1998, and (2) the net change in NAVs between end of 2003 and end of 2000 (see the text for details). The fraction of funds in each category is expressed in terms of number, present value of the investments, and NAV. For each category, the fraction of funds in this category that have cash flow activities between (end of) December 2000 and December 2003 is displayed. Finally, average PIs (value weighted) are shown with NAVs either treated as correct or written off.

3.1 VentureXpert data

TVE also collects information on underlying investments of private equity funds in a separate and widely available “investment dataset” known as VentureXpert. Data include information about 29,739 companies (location, industry description, and age), their investment characteristics (time of investment, stage, equity invested, exit date, and mode), and funds (fund size, investment focus, vintage year, and headquarters).⁴

3.2 Extension of our sample

The unique feature of our dataset is that we have a link between the “investment” dataset and the cash flow dataset. In the investment dataset, we can observe the characteristics of funds that are not included in the cash flow dataset and thus of funds that are not included in either widely used industry benchmarks or the previous fund-level studies like Kaplan and Schoar (2005). We count 476 funds raised between 1980 and 1993 with more than five investments in the investment dataset, size above \$5 million (1990 U.S. dollars), which are *not* part of the base-sample group of funds; they constitute the “additional sample.” The increase in coverage due to the addition of these funds can be seen in panel B of Table 2. The merged dataset (base + additional) is called the “extended sample,” contains 1,328 funds, and represents 57% of the universe in terms of size. In comparison, the base sample contains 852 funds and represents 43% of the universe.

3.3 Differences between base-sample funds and extended-sample funds

TVE obtains data mostly from fund investors as most fund managers refrain from giving out information. A priori, fund investors do not have an incentive to report only good or bad performance. Nonetheless, fund investors that report to TVE might differ from the representative investor. They might be large private equity investors with potentially privileged access to larger and more established funds or may avoid certain funds (e.g., first-time funds) by preference. Whether the funds held by reporting investors have different performance and characteristics than the average fund is an empirical issue.

Table 4 reports the characteristics of both the additional funds and base-sample funds.⁵ We observe that additional funds are smaller than base-sample funds and are less experienced (average sequence of 2.39 versus 3.01 in the base sample). Importantly, additional funds have fewer investments exited via an IPO or an M&A (45% versus 50%; *t*-stat is 3.26 for the spread). Finally,

⁴ We operate a minor correction to the dataset as we extrapolate missing fund size (11% of the cases) by running an OLS regression with the rest of the funds using as explanatory variables: year-fixed effect, a dummy for VC, a dummy for Europe-focus, fund sequence, and the sum of all investments (in log). The *R*-square is 50%. This extrapolation enables us to increase our sample size but leave results unaffected. We also verify that different treatments of missing investments and missing investment size in the investment dataset have no impact on our results.

⁵ VCs have a higher fraction of successful exits than BOs because they have less inactive investments at the end of our time period and not because of a lower bankruptcy rate.

Table 4
Descriptive statistics—extended and base sample

	Extended sample	Additional sample			Base sample			Spread	<i>t</i> -stat
	VC + BO	VC	BO	VC + BO	VC	BO	VC + BO		
Size	110	57	174	76	80	257	129	53	4.11
Sequence	2.78	2.40	2.35	2.39	3.08	2.83	3.01	0.62	4.37
IPO+M&A	0.48	0.48	0.31	0.45	0.53	0.36	0.50	0.05	3.26
IPO	0.19	0.19	0.11	0.18	0.22	0.13	0.20	0.02	2.71
U.S. investments	0.89	0.88	0.87	0.88	0.94	0.72	0.91	NA	NA
U.S. headquarters	0.72	0.87	0.87	0.87	0.70	0.50	0.64	NA	NA
U.K. headquarters	0.10	0.05	0.04	0.05	0.10	0.21	0.13	NA	NA
FRA headquarters	0.02	0.00	0.00	0.01	0.02	0.05	0.03	NA	NA
NL headquarters	0.01	0.00	0.00	0.00	0.00	0.04	0.01	NA	NA
Other headquarters	0.05	0.08	0.09	0.08	0.02	0.04	0.03	NA	NA
U.S. focus	0.47	0.00	0.00	0.00	0.81	0.55	0.74	NA	NA
EU focus	0.17	0.00	0.00	0.00	0.19	0.45	0.26	NA	NA
No. of observations	1,328	398	78	476	614	238	852	NA	NA

This table gives descriptive statistics for the extended, additional, and base samples. The “additional sample” consists of funds added to the “base sample” to obtain the “extended sample.” We report the average of (1) capital committed in millions of 2003 U.S. dollars (size); (2) fund sequence number; (3) the proportion of investments exited via either an M&A or an IPO; (4) investments exited via an IPO; (5) investments made in the United States; (6) funds headquartered in the United States, UK, France, the Netherlands, and other countries; and (7) funds targeting the United States and the EU. Base-sample funds without sufficient data about their investments are not included in averages (except for size and sequence). *t*-stat for testing that the base sample and additional sample means are equal is reported whenever appropriate. Information about headquarters and targeted region is sometimes missing (hence fractions do not add up to one).

we note that 81% of the funds in our extended sample are headquartered in the United States and 89% of the investments are made in the United States.

3.4 Correction for sample selection bias

In the literature, exit success is frequently used as a proxy for fund performance. Certain researchers use the fraction of investments exited via an IPO as a proxy for performance (e.g., Sorensen, 2007) while other researchers use the fraction of investments exited via either IPO or M&A as a proxy for performance (e.g., Hochberg, Ljungqvist, and Lu, 2007). The widely held belief is that exit success and performance are highly related. As exit success is available for both additional-sample funds and some base-sample funds; this variable is used to approximate the performance of extended-sample funds. We proceed in two steps. First, we verify that exit success is significantly related to true performance. Second, we extrapolate performance of the “additional” funds.

3.4.1 Exit success as a proxy for true performance. We assume that the true performance (*PI*) equals a performance proxy (exit success) plus noise. We further allow this noise to be dependent on fund characteristics such as size, size-squared, fund sequence, fraction of non-U.S. investments, and fraction invested in VC. The model writes

$$\ln(1 + PI) = a + b * Exit + c * [size \ size^2 \ sequence \ nonUS \ VC]' + u. \quad (2)$$

In Table 5, we report the results from the estimation of Equation (2). We find that exit success is positively and significantly related to true performance across various specifications.⁶ Of interest, we find that if we redefine exit success as only the fraction of investments that exit via an IPO, as is sometimes the case in the literature, the relation between exit success and performance is weaker. Consequently, this result shows that it is likely better to use both M&A exits and IPO exits as successful exits, as done in Hochberg, Ljungqvist, and Lu (2007). We also find that the relation between cash-flow-based performance and exit success is lower for VC funds, but not statistically significantly so (specification 3). Finally, results are similar when using IRR and multiple-to-measure performance.

3.4.2 Performance extrapolation. As mentioned above, Equation (2) can be used to extrapolate the performance of additional-sample funds. Results are reported at the bottom of Table 5. We find that additional-sample funds have an expected *PI* that is 0.13 lower than that of base-sample funds, which is statistically significant (specification 1). This translates into a decrease in the overall average *PI* of 0.04. We conclude that performance estimates based on funds for which cash flow information is available have to be considered as significantly overstated. Our results are consistent with Kaplan and Schoar's (2005, p. 1796) observation that

One potential bias in our returns sample, therefore, is toward larger funds. We also over-sample first-time funds for buyout funds [*opposite for venture capital funds*]. As we show later, larger funds tend to outperform smaller ones, potentially inducing an upward bias on the performance of funds for which we have returns. Also, first-time funds do not perform as well as higher sequence number funds. Therefore, our results for average returns should be interpreted with these potential biases in mind.

Finally, our sample bias correction is conservative as there are still some funds that are not included in our computations. These funds not only do not have cash flow data, but they do not have data about their investments either. Such funds are likely not to have had many successful exits (otherwise TVE would have identified them). Consequently, their performance is expected to be even

⁶ Equation (1) is estimated with a subsample of funds, which may create a bias for the coefficients of interest. This bias is also called sample selection bias but it is distinct from the sample selection bias we focus on in the paper (the fact that performance may be different in the base sample compared to the universe). To control for this potential bias, we use a Heckit regression (nontabulated; see Greene, 2003, for details). We first note that in Equation (1), only exit success and size are significant. We thus consider this reduced equation as the equation of interest. Second, the so-called lambda is derived from a Probit regression with a dummy variable for being base sample as a dependent variable, and as independent variables, both the variables of interest (exit success and size) and instruments (e.g., fund sequence, VC dummy). Third, we run a regression of performance on exit success, fund size, and lambda. We find that the latter is not statistically significant. This indicates that the appropriateness of using exit success as a proxy for performance is not sample dependent. This result is intuitively appealing and allows us to extrapolate performance for extended-sample funds, hence correcting for the sample bias that we focus on.

Table 5
Performance and exit success

Dependent variable	PI Spec 1	PI Spec 2	PI Spec 3	PI Spec 4	PI Spec 5	IRR Spec 6	DPI Spec 7
(log) IPO + M&A (%)	0.29*** (2.61)		0.27** (1.94)	0.27* (1.79)	0.23*** (3.39)	0.11*** (3.52)	0.62*** (3.41)
(log) IPO + M&A (%) * VC dummy			0.02 (0.20)				
(log) IPO (%)		0.22*** (2.41)					
Control variables	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Net of fees	Yes	Yes	Yes	No	Yes	Yes	Yes
Time F.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Written-off NAVs	Yes	Yes	Yes	Yes	No	No	Yes
R-squared	12%	11%	14%	13%	15%	12%	8%
	Extrapolated performance (additional sample)						
Average	0.79	0.81	0.81	1.04	0.87	0.13	1.51
Difference with base performance (<i>t</i> -statistic)	0.13*** (3.36)	0.11*** (2.76)	0.11*** (2.80)	0.12** (2.38)	0.12*** (2.72)	0.02*** (2.57)	0.47*** (3.34)

This table reports OLS regression results. The dependent variable is either $\log(1 + \text{PI})$, $\log(1 + \text{IRR})$, or $\log(1 + \text{DPI})$. Explanatory variables include: (1) $\log(1 + \text{fraction of investments that exit via either an IPO or M\&A})$; (2) same as (1) times a venture capital dummy variable; and (3) $\log(1 + \text{fraction of investments that exit via an IPO})$, and a set of control variables ($\log$ of fund size and its square, dummy variable for VC focus, dummy variable for EU investment focus, $\log$ of fund sequence). Time-fixed effects are included in each specification. Inference is based on heteroskedasticity-consistent standard errors. *t*-statistics are reported in parentheses. To be included in the regression, funds should have at least five exits reported in the TVE "investment" dataset ($N = 618$). Results on extrapolated fund performance for each specification are at the bottom. Averages are weighted by the present value of investments.

*, **, *** statistical significance at the 10%, 5%, and 1% levels, respectively.




Figure 2
Impact of corrections on multiples

This figure shows average multiples per vintage year. We report (1) the total value over paid-in capital (TVPI), defined as the sum of the capital distributed plus final NAV divided by the sum of capital invested and (2) the distributed over paid-in capital (DPI), defined as the sum of the capital distributed divided by the sum of capital invested. DPI is shown for both the base sample and the extended sample.

lower than funds included in the extended sample. Such funds are 47% of the universe during our sample period.

4. Final Performance Estimate

4.1 Net-of-fees performance

We have reported above that “uncorrected” average *PI* of base-sample funds is 1.01, changing the weighting scheme from capital committed to value invested reduces average *PI* by 0.02, writing off NAVs reduces *PI* by 0.07, and including the projected *PI* for additional funds decreases the average *PI* by 0.04. These three corrections, therefore, reduce the average *PI* to 0.88, which is statistically significantly below 1 (*t*-stat is 2.25; nontabulated). Moreover, underperformance is present for both BO funds and VC funds, but more so for VC funds.

Because profitability indices might be difficult to interpret, we convert them into a more intuitive figure by computing corresponding “alphas” (sometimes called excess-IRR). If we assume that fund returns are given by the CAPM with a beta of 1 and a constant alpha, then alpha is the constant to be added to the realized S&P 500 returns to make *PI* equal to 1. We find an average yearly alpha of -2.94% after the three corrections (panel B of Table 2).⁷

Figure 2 shows the impact of our corrections on multiples. The average (size-weighted) multiple is 2.10. After writing off NAVs, it is 1.88. Finally, the average multiple for the extended sample is 1.77.

⁷ The correspondence between alpha and *PI* is linear when *PI* is close to 1 but is irregular when *PI* moves away from 1. As *PI* tends toward 1, it can be shown that $PI = (\text{modified Macaulay's duration}) * \alpha + 1$ with duration defined as the (present value of cash flows weighted) average time at which distributions are made minus weighted average time at which investments are made.

Note that our net-of-fees performance is not net of *all* fees and the above estimate is thus biased upward. First, about 20% of LPs (see Lerner, Hardyman, and Leamon, 2004) hire gatekeepers. These intermediaries typically charge 1% of the fund's size in addition to a 5–10% carried interest. Also, LPs without these gatekeepers spend considerable resources on screening funds. Such expenses can be substantial compared to the cost of investing in an S&P 500 Index fund. Second, if LPs need to liquidate their position before the closure of the fund, a penalty is charged. Third, distributions are often made with shares rather than cash. These shares have a lockup period and LPs typically incur a severe price impact when selling these shares (see Gompers and Lerner, 1999). As none of these costs are reported in our dataset, we cannot compute their exact impact.

The above performance estimate is, therefore, an upper bound for true investor performance.

4.2 Gross-fees performance

It is interesting to assess the gross-of-fees performance of funds in order to relate to both the literature on gross-of-fees performance of private equity investments (e.g., Cochrane, 2005) and the literature on the added value of professional asset managers.

We begin by assessing the magnitude of management fees. Given that we cannot reliably determine which cash flow in our data corresponds to a fee payment, we use an approximation based on the fee contract descriptions provided by Gompers and Lerner (1999) and Metrick and Yasuda (2007). Using a sample of both VC and BO funds raised over recent years, Metrick and Yasuda indicate that most management fees are set to 2% of capital committed during the investment period (year 1 to year 5). Then, during the postinvestment period (year 6 to year 10), the two most common arrangements are to charge 2% of net invested capital (42% of the funds) and charge less than 2% of net invested capital (33% of the funds).

Gompers and Lerner (1999) analyze a sample of VC funds raised between 1978 and 1992, and thus corresponding to our time period, and find that a 2.5% management fee for the entire fund life is most common. Consequently, we consider five management fee structures ranked from the least to the most costly: (1) 2% of committed capital during the investment phase (year 1 to year 5) and 1% of committed capital during the postinvestment phase (year 6 to year 10); (2) 2% of committed capital during the investment phase and 2% of NAV (proxy for net invested capital) during the postinvestment period; (3) 2% of committed capital throughout; (4) 2.5% of committed capital during the investment phase and 2.5% of NAV during the postinvestment phase; and (5) 2.5% of committed capital throughout.

Results in Gompers and Lerner (1999) and Metrick and Yasuda (2007) and our conversations with investors indicate that for our time period, arrangement 5 was the most common. To be conservative, however, we will focus on

Table 6
Gross-of-fees performance

		(1)	(2)	(3)	(4)	(5)	(6)	
Management fees		0%	2% Kcom	2% Kcom	2% Kcom	2.5% Kcom	2.5% Kcom	<i>Delta</i>
Investment phase		0%	1% Kcom	2% NAVs	2% Kcom	2.5% NAVs	2.5% Kcom	
Postinvestment phase		0%	1% Kcom	2% NAVs	2% Kcom	2.5% NAVs	2.5% Kcom	
Alpha								
Carry	Hurdle							
0%	NA	-2.26	1.11	1.33	1.68	2.22	2.64	-2.12
20%	8%	0.14	3.30	3.45	3.80	4.33	4.71	0.00
20%	0%	0.52	3.58	3.82	4.12	4.64	5.04	0.32
	<i>Delta</i>	-3.66	-0.50	-0.35	0.00	0.53	0.91	
Profitability index								
Carry	Hurdle							
0%	NA	0.92	1.04	1.05	1.07	1.09	1.10	-0.09
20%	8%	1.01	1.13	1.14	1.16	1.18	1.19	0.00
20%	0%	1.03	1.15	1.16	1.17	1.19	1.21	0.01
	<i>Delta</i>	-0.15	-0.03	-0.02	0.00	0.02	0.03	

This table reports the gross-of-fees performance with different assumptions regarding the fee structure for base-sample funds. Average alpha and the average *PI* are displayed. Investment phase is the first 5 years of a fund's life and the postinvestment phase is the next 5 years. Carried interest (carry) can be set to 0% or 20% and hurdle rates (hurdle) to 0% or 8%. Management fees are either (1) 0 throughout, (2) 2% of committed capital (Kcom) during the investment phase and 1% of Kcom during the postinvestment phase, (3) 2% of Kcom during the investment phase and 2% on NAVs during the postinvestment period, (4) 2% of Kcom throughout, (5) 2.5% of Kcom during the investment phase and 2.5% of NAVs during the postinvestment phase, or (6) 2.5% of Kcom throughout. Deltas show spread with respect to the base case, keeping the other dimension unchanged. The base case is shown in bold.

arrangement 3 (2% of committed capital throughout). Nowadays, arrangements 1 and 2 are the most common.

Next, we estimate carried interest. Metrick and Yasuda (2007) report that 92% of the funds have a 20% carry, 83% have it with a hurdle rate, and the carry basis is capital committed for 83% of the funds. Out of the funds with a hurdle rate, 74% set it at 8% and 92% use a catch-up provision. We thus set the carried interest at 20% and show results with and without an 8% hurdle rate (net of all fees with full catch-up). Table 6 shows the results. Column 1, row 1 (0% carry and 0% management fees) is by definition the net-of-fees performance for the base sample. We find that the consideration of management fees alone leads to a positive gross alpha (i.e., it increases from -2.26% to 1.11%). As discussed above, we focus on arrangement 3. In this case, management fees increase alpha by about 4% and carried interest by an extra 2%. We find gross alpha to average 3.8% per year and gross *PI* to average 1.16 (compared to -2.3% and 0.92 net-of-fees).⁸ A conservative estimate of total fee level represents, therefore, more than a quarter of capital invested even though performance is below that of the S&P 500.

To obtain the gross-of-fees performance on the extended sample, we use results from Table 5 specification 4, which show that funds in the extended sample

⁸ We discount all cash flows including fees at the S&P 500 rate of return. This is because we simply want to compare the performance of private equity funds to the one of the S&P 500. If we want risk-adjusted performance, then we need to assess the risk of the fees separately and discount them accordingly.

have an average gross-of-fees performance of 1.04. The final gross performance estimate is a *PI* of 1.12 and an alpha of 3% per year. This result indicates that the average fund manager may generate added value. Also, this result partly explains discrepancies between our results and reports of high gross-of-fees performance for private equity investments (e.g., Cochrane, 2005).

Our results offer two additional important contributions. The first is the finding of substantial compensation, which comes mainly from management fees and not incentive fees. The reason for large management fees is that the amount invested by a fund is much lower than capital committed—the basis for such fees. For example, Ljungqvist and Richardson (2003) report that 16% of committed capital is invested at the end of the first year (we find a similar figure with our data). A 2% fee payment at the end of the first year is thus as high as 12.5% of the amount invested.

Second is the quantification of the value of different fee arrangements. Previous literature (Gompers and Lerner, 1999; and Metrick and Yasuda, 2007) relies on simulations to gauge the total value of fees and their sensitivities to different arrangements. As assumptions about cash flow timing and amount are intimately related to the value of the fees, having *actual* cash flows represents a clear advantage. Using this feature of our data, we show that moving from 2% fees to 2.5% fees translates into a reduction of alpha by 1.3%, reducing the amount of fees in the postinvestment period (from 2% of capital committed to 1%) is only worth 0.6% per year, having fees based on NAV rather than capital committed in the postinvestment period is worth even less (0.4%), and the value of a hurdle provision is only 0.3% per year. The latter result is probably the most surprising and is due to the fact that once the hurdle rate is met, GPs get 20% of the raw profits (and not of the profits beyond the hurdle rate). Overall, we observe that variations in the fee arrangements have an impact ($\pm 1\%$ per year) that is in itself economically sizeable, but not large relative to the total fees.

5. Further Analysis

5.1 Comparison with related literature

We can divide the literature on private equity performance into two groups. The first documents the gross-of-fees performance of GPs' individual VC investments (e.g., Cochrane, 2005; and Hwang, Quigley, and Woodward, 2005). Cochrane finds that log returns of VC investments have negative alphas, but arithmetic returns (and alpha) are high. However, Hwang, Quigley, and Woodward use the same dataset, but with fewer missing financing rounds. They find that average performance is close to that of the S&P 500. Studies of VC returns do not, therefore, contradict our findings in any obvious way. In addition, we note above that gross-of-fees performance in our dataset is also relatively high.

The second set of studies focuses on the cash flow stream to investors, which includes fee payments and carried interests, and contains both BO and VC investments (Ljungqvist and Richardson, 2003; and Kaplan and Schoar, 2005). The discrepancy between our results and that of Kaplan and Schoar has been decomposed above (our three corrections). The discrepancy with Ljungqvist and Richardson is stronger (they find an average *PI* of 1.25) and can be traced to two facts. First, Ljungqvist and Richardson have a disproportionate number of BO funds, which we find perform better than VC funds. Their funds are also larger, more U.S.-focused, and more experienced. These characteristics are all found to be positively related to performance (see below). Second, data are from a single investor and results in Lerner, Schoar, and Wong (2007) indicate that performance can vary dramatically across different types of investors.

5.2 Aggregation of IRRs

IRR is frequently used as performance measures for private equity funds. In this subsection, we illustrate how average IRRs are significantly upward biased. This helps us to understand why there are so many reports of high performance in the media and possibly among investors (Appendix).

The aggregation of IRRs is biased if IRR and duration are correlated. Indeed, IRR is a per period return while the object of interest to the investor is total return (i.e., $\text{duration} \times \text{IRR}$). If they are correlated, then $E(\text{duration} \times \text{IRR}) \neq E(\text{duration}) \times E(\text{IRR})$. To illustrate, assume that good performance (say 100%) occurs over 2 years on average and bad performance (say -20%) occurs over 10 years on average, both with equal probability. Expected performance is then obviously not 60%. Note that such an issue is irrelevant for *PI* since it measures excess return.

We can examine this further by empirically estimating the correlation between fund performance and duration. To compute the duration of a fund, we proceed as in the fixed-income literature. We first compute the average month at which cash flows are received where the weights are the present value of the related cash flow divided by the present value of all the cash flows. Similarly, we compute the average month when capital was paid. The difference between the two dates is the fund duration.

We run OLS regressions of fund performance on duration and the fund characteristics known to be related to performance (fund size, fund size squared, fund sequence, VC dummy, exit success). In each specification, duration is by far the most significant and robust explanatory variable. Funds with longer duration perform worse, hence the average IRR is biased upward. One way to correct for this bias is to weight each IRR by the product of the present value of investment and duration [$\text{duration} \times \text{PV}(\text{invested})$]. Thus, we obtain a sort of IRR per year and per dollar invested. Doing so decreases the average IRR from 14.64% to 12.22%, a substantial 2.42% spread (see Table 1). For the vintage year 1985, the size-weighted IRR is almost twice as large as the average IRR, which is both time and present value weighted (22.86% versus 13.88%). This is

Table 7
Risk adjustments

BO	VC	All	BO	VC
Panel A: Buyout versus venture capital				
S&P 500	S&P 500	0.92	0.96	0.88
S&P 500	NASDAQ	0.92	0.96	0.87
Post-IPO, 3FF	Post-IPO, 3FF	1.02	0.87	1.20
Matched-industry/size	Matched-industry/size	0.76	0.75	0.77
Panel B: Europe-focus versus US-focus				
EU	U.S.	All	BO	VC
Panel B: Europe-focus versus US-focus				
EU index, \$ return	S&P 500	0.94	0.97	0.89
EU index, euro return	S&P 500	0.92	0.95	0.88

This table reports the average *PI* found with different set of discount rates in the base sample. We also show results separately for VC funds and BO funds. In panel A, the discount rate differs for BO and VC funds. The discount rate is either S&P 500, NASDAQ (value-weighted provided by WRDS), post-IPO cost-of-capital estimate (using the three-factor model), or matched-industry/size cost-of-capital. In panel B, the discount rate differs for EU and U.S. funds. The discount rate is S&P 500 for U.S. funds and, for EU funds, is a Europe-all-market index (source: K. French website) with returns and cash flows expressed either in U.S. dollars or euros. NAVs are written off.

why our study focuses on *PI* as a measure of performance and also why certain results in the literature based on IRR may appear to contradict what we report here.

5.3 Different benchmarks

In this section, we show how results change when using different discount rates. In particular, we use discount rates that may better reflect private equity fund risk.

We begin by discussing two approaches to estimate the cost-of-capital for VCs and BOs. In the first approach, we use post-IPO cost of capital estimates provided by Brav and Gompers (1997) for VC-backed IPOs and Cao and Lerner (2006) for BO-backed IPOs. In the second approach, we use the industry of each portfolio company held by a fund to estimate an “industry/size-matched cost-of-capital.”⁹ We assume that each portfolio company has the same unlevered beta as the average publicly traded stock in the industry.¹⁰ For BO, we assume that the leverage of the portfolio company decreases over investment life from a debt-to-asset ratio of 0.75 (at entry) down to the leverage that prevails in the industry (at exit).¹¹

Panel A of Table 7 shows that the industry/size-matched cost-of-capital produces a risk-adjusted *PI* of 0.75 for BO funds and 0.77 for VC funds (before risk adjustments, it is 0.96 and 0.88, respectively). This shows that risk adjustments decrease performance substantially. According to this result,

⁹ This procedure is similar to that of Ljungqvist and Richardson (2003).
¹⁰ We assume a corporate tax rate of 35%; beta debt is 0.25 for BOs and 0 for publicly traded stocks.
¹¹ 0.75 is the (investment amount weighted) average debt-to-asset ratio that prevailed between 1980 and 2002 in the BO industry according to the S&P Leveraged Lending Review. VC investments are matched to the publicly traded companies that are in the bottom-size quintile of their industry; leverage is assumed to be the same.

investors have lost about 25% of the capital invested in either BO or VC funds. We note that using the NASDAQ instead of the S&P 500 for VC funds hardly changes performance estimates and that post-IPO cost-of-capital estimates lead to somewhat extreme results.¹²

Next, we show results with different discount rates for EU-focused funds and U.S.-focused funds. The discount rate is the S&P 500 for U.S. funds and a Europe-all-market index for EU funds (source: K. French website). We show two results for EU-focused funds. The first one is with cash flows expressed in U.S. dollars and using the European index returns in U.S. dollars as discount rates. The second one is with cash flows expressed in euros and using the European index returns in euros as discount rates.

Panel B of Table 7 shows that *PI* changes slightly when a separate benchmark is used for Europe only if returns are expressed in U.S. dollars (increases from 0.92 to 0.94). Our finding of underperformance is thus robust to using different benchmarks. In addition, when correcting for risk, results point to a much deeper underperformance.¹³

5.4 Cross-sectional results

Average performance is found to be low, but there are cross-sectional differences. It is then interesting to investigate which type of fund performs better. Kaplan and Schoar (2005) show that performance is related to size, experience, and importantly, to past performance.¹⁴ Our sample being slightly different (two more years, only mature funds, non-U.S. funds included, NAVs written off), we verify whether this holds true for our sample. In addition, we verify whether this also holds true for our extended sample after accounting for sample selection bias. Finally, we investigate whether EU-focused funds perform differently than U.S.-focused funds and whether we find evidence of money chasing deals. Money chasing deals refers to the finding of Gompers and Lerner (2000) that valuations are higher at times when capital committed is higher. The results are reported in Table 8.

We confirm that performance increases with fund size and is lower for first-time funds. The relation between performance and fund sequence is positive, but not always significant. We interpret this as evidence that fund sequence is

¹² For VC, (post-IPO) risk-adjusted average *PI* increases to 1.20 (from 0.88). This is because VC-backed companies resemble small growth stocks, which, according to the three-factor model, have a cost of capital equal to the risk-free rate; using estimates in Brav and Gompers (1997) and the average values of the three factors over our time period, excess cost of capital is $1.21 \times 0.62\% + 1.11 \times (-0.12\%) - 1.07 \times 0.56\% = 0.02\%$ per month. As VC funds returned more than the risk-free rate, their risk-adjusted performance is high. For BOs, the estimate is 0.87 (down from 0.96), which seems more plausible. The market beta is 1.3 and BOs resemble small value stocks. However, if we use the value-weighted risk estimates instead of the equally weighted one, the average *PI* drops to 0.64, which is due to the large alpha uncovered by Cao and Lerner (2006).

¹³ Stakes in private equity funds cannot be readily sold. This illiquidity should require an additional premium for private equity investors, similar to what Aragon (2007) finds for hedge funds.

¹⁴ Following Kaplan and Schoar (2005), if a firm raises a fund in 1990 and then a fund in 1993, then the 1993 fund performance as of 12/2003 is regressed onto the 1990 fund performance as of 12/2003 and the latter is called past performance.

Table 8
Performance and fund characteristics

	Base sample					Extended sample	
	Spec 1	Spec 2	Spec 3	Spec 4	Spec 5	Spec 6	Spec 7
Panel A: Fund characteristics and performance							
Size	0.28** (2.13)	0.26** (2.03)	0.27** (2.07)	0.26** (1.97)	0.23* (1.76)	0.14** (2.38)	0.14** (2.35)
Size-squared	−0.02 (−1.22)	−0.02 (−1.14)	−0.02 (−1.15)	−0.02 (−1.08)	−0.01 (−0.90)	−0.00 (−0.20)	−0.00 (−0.19)
Sequence	0.09 (1.32)	0.06 (1.15)		0.10* (1.70)	0.09 (1.60)	0.10*** (3.72)	0.10*** (3.87)
First-time			−0.14** (−2.35)				
Venture capital	−0.14* (−1.88)	−0.11 (−1.44)	−0.14* (−1.86)	−0.15** (2.04)	−0.14* (−1.88)	−0.04 (−1.22)	−0.05 (−1.35)
Europe-target	−0.22*** (−3.04)	−0.19*** (−2.72)	−0.22*** (−3.13)	−0.20*** (−2.94)	−0.15** (−2.17)	−0.06* (−1.86)	−0.10*** (−3.16)
Total capital committed				0.02 (0.57)	0.05 (1.23)		
Time F.E.	Yes	Yes	Yes	No	No	Yes	Yes
Erase living deads	No	Yes	No	No	Yes	No	Yes
No. of observations	852	852	852	852	852	1328	1328
R-squared	6%	7%	6%	5%	4%	16%	14%
Base sample				Extended sample			
	Spec 1	Spec 2	Spec 4	Spec 5	Spec 6	Spec 7	
Panel B: Performance persistence							
Past performance	0.20*** (3.97)	0.19*** (3.95)	0.22*** (4.62)	0.24*** (4.84)	0.16*** (3.70)	0.17*** (4.00)	
Size	0.33* (1.73)	0.33* (1.77)			0.25** (2.34)	0.26** (2.44)	
Size-squared	−0.02 (−0.92)	−0.02 (−0.96)			−0.01 (−1.10)	−0.01 (−1.17)	
Sequence	0.18 (1.40)	0.14 (1.15)			0.10 (1.58)	0.10 (1.62)	
Venture capital	0.06 (0.55)	0.08 (0.69)			−0.03 (0.52)	−0.02 (−0.35)	
Europe-target	−0.16 (−1.49)	−0.17 (−1.58)			−0.15** (−2.58)	−0.09 (−1.54)	
Time F.E.	Yes	Yes	Yes	Yes	Yes	Yes	
Write-off NAVs	No	Yes	Yes	No	Yes	No	
No. of observations	400	400	400	400	842	842	
R-squared	11%	14%	9%	7%	12%	14%	

This table reports cross-sectional regression results. The dependent variable is the *PI*. Independent variables include log of size and its square, log of sequence number, a dummy variable taking the value 1 if the fund is the first raised by the private equity firm (0 otherwise), a VC dummy variable, a Europe-focused dummy variable, log of capital committed to all funds of the same vintage year, and performance of the previously raised fund (panel B). Time-fixed effects are included in each specification. Inference is based on heteroskedasticity-consistent standard errors. *t*-statistics are reported in parentheses. *, **, *** Statistical significance at the 10%, 5%, and 1% levels, respectively.

not as good as fund size to proxy skills. A firm might be raising its fifth fund despite an inferior track record. With an inferior track record, however, it is difficult to raise a large fund.

We find significant underperformance from EU-focused funds, which echoes the findings of Hege, Palomino, and Schwienbacher (2006). We find no evidence of either a money-chasing deal effect, or a significant performance difference between VC and BO funds. In addition, and unlike Kaplan and Schoar (2005), we do not find evidence of a concave relationship between performance and size. However, the persistence result of Kaplan and Schoar is found to be large and robust. It is present in each specification on both our base sample and extended sample. When past performance is included in the regression, all the other characteristics lose their significance, including size. It appears that the unique explanatory variable for fund performance is the performance of the previously raised fund.

5.5 Performance of recently raised funds

More money and more funds were raised between 1994 and 2000 than between 1980 and 1993. Though performance figures for these are not yet definitive, it is interesting to see how these recently raised funds fare.

Using the sample of mature funds (i.e., those raised between 1980 and 1993) at different stages of their lifecycle, we can forecast the performance of the new funds (i.e., those raised between 1994 and 2000). We run a regression with final performance as a dependent variable and fund characteristics as explanatory variables. Specifically, two sets of explanatory variables are used. The first set contains time-unvarying fund characteristics: the natural logarithm of fund size and its square, the natural logarithm of the sequence number of the fund within its family, and a VC dummy variable. A second set of explanatory variables includes the characteristics of funds either in their third year, fourth year, etc., all the way up to their ninth year. These characteristics are the realized performance *at that point in time* (intermediate *PI* without NAV; e.g., in their third year), the ratio of NAV to present value of investments *at that time*, and three crossterms of the ratio of NAV to present value of investments *at that time* with fund sequence, fund size, and fund focus (dummy variable that is 1 if VC), respectively. These crossproducts aim to assess the extent to which larger funds, more mature fund families, and VC funds differ in terms of the aggressiveness or conservativeness of their reports of NAV.

To forecast the performance of funds raised in 1994 (resp. 1995, . . . , 2000), we use the regression of final performance on (1) the first set of characteristics and (2) the characteristics of the mature funds in their ninth year (resp. eighth year, . . . , third year). Regressions are estimated independently for each vintage year (1994 to 2000).

The results are reported in Table 9. The first striking observation is that the amount of capital raised by private equity funds in our cash flow dataset from 1994 to 2000 is almost three times that raised from 1980 to 1993. The number

Table 9
Expected performance of young funds

Year	Universe N funds	Universe size	Sample N funds	Sample size	PV inv.	PI w/t NAV	PI with NAV	Exp. PI
2000	1,748	27,824	280	14,190	84,609	0.14	0.86	0.85
1999	1,134	166,549	227	96,598	69,188	0.25	0.87	0.74
1998	859	145,333	198	103,186	65,996	0.41	1.06	0.93
1997	680	109,501	170	64,606	33,266	0.65	1.24	1.10
1996	471	63,871	108	30,019	14,883	0.93	1.36	1.28
1995	435	51,597	109	29,926	12,424	0.95	1.27	1.12
1994	351	46,148	102	25,381	13,543	0.95	1.24	1.09
Total Mean	5,678	610,823	1194	363,906	293,909	0.39	1.01	0.92

This table shows the expected performance (Exp. PI) of funds raised between 1994 and 2000. Expectation is based on an OLS regression described in the text. The number of funds raised and the sum of committed capital each year is also displayed (both for the universe and for the TVE sample), as well as the average PI as of December 2003, with NAVs treated either as market values or written off. The mean is computed with the present value of investments (PV inv.) as weight.

of funds, however, is similar, with 1,171 recently raised funds compared to our base sample of 852 mature funds.

If we consider NAV to be the current market value of the new funds, their average *PI* is 1.01. This figure is similar to the 0.99 found for mature funds. If NAVs are written off, the average *PI* is 0.42. For such young funds, writing off NAV is obviously unwarranted, but it gives an informative lower bound for their performance. Finally, our extrapolation exercise described above leads to an average performance of 0.92, which is that observed for the mature funds before sample bias correction. Our findings thus suggest that new funds have very similar expected performance as the mature funds investigated in this paper.

6. Discussion

Given the finding of low performance, an interesting area for further research is to understand why investors allocate large amounts to this asset class given such poor historic performance. We discuss three sets of “nonmutually exclusive” explanations as directions for future research: learning, mispricing, and positive externalities.

6.1 Learning and performance persistence

Running a private equity fund requires skill, and learning may play an important role. In addition, we found that inexperienced funds have lower performance. It is thus possible that by participating in inexperienced and hence poorly performing funds, investors tacitly obtain the right to participate in future better performing funds. This would be a first “learning channel” that could explain observed average poor performance for this relatively young industry.

Investing in private equity equally requires skill. Lerner, Schoar, and Wong (2007) argue that large differences in skills explain differences in performance

across investor types. This would be a second learning channel explaining poor performance by investors' initial learning costs.

These learning-based explanations may be tempered by our finding that recently raised funds have an equally low performance and that there is no general upward trend in the time series of fund performance. Nonetheless, performance disclosure has been rare in the past and might become more frequent in the future. It is thus possible that, as a result, learning will be faster and future performance better than that observed over the last 25 years.

6.2 Mispricing

Certain investors might have misvalued this asset class. We note that payoffs are highly skewed and investors might attribute too much weight to the performance of a few successful investments. Along these lines, we note that entrepreneurial investment in nonpublic companies, whose performance distribution resembles that of private equity funds, is also found to have relatively low performance (e.g., Moskowitz and Vissing-Jorgensen, 2002).

Another possibility is that investors have a biased view of performance because performance is generally reported gross-of-fees and, as shown above, fees larger than for other asset classes. Along the same lines, we note that it is basically impossible for investors to benchmark the past performance of funds with information reported in prospectuses. These documents contain only multiples and IRRs (see also CalPERS online report at <http://www.calpers.ca.gov/index.jsp?bc=/investments/assets/equities/aim/private-equity-review/aim-performance-review/home.xml>). We show above that averages IRRs give upward-biased performance estimates. In addition, IRRs cannot be directly compared to the performance of, say, the S&P 500 over the same period. Similarly, it makes little sense to compare multiples with S&P 500 returns. In our conversations with investors, we note that a recurring argument is that they are satisfied with past performance because they "doubled" their money. Funds in our base sample indeed offer a multiple of 1.9 net-of-fees. However, how long it took to obtain such a multiple must be considered, and as investors face a continuous stream of inflows and outflows, this is not trivial to determine. We find that the average fund duration is 75 months (i.e., 6.25 years). The stock-market portfolio has returned on average 1% per month from 1980 to 2003, which means that over 6.25 years, investors would have more than doubled their money ($\times 2.1$) with the S&P 500.

6.3 Side benefits of investing in private equity funds

A potential explanation for the low performance of private equity funds is that investors' objective may not only be to maximize returns. Ljungqvist and Richardson (2003) recount that the investor who provided them with data invests in private equity funds to establish a commercial relation with GPs: "...the Limited Partner's twin investment objectives (are) not only to obtain

the highest risk-adjusted return, but also to increase the likelihood that the funds will purchase the services our Limited Partner's corporate parent has to offer." These side benefits include consulting work (e.g., for M&As) and underwriting securities for debt or equity issues. A recent study by Hellmann, Lindsey, and Puri (2005) corroborates this view. The authors argue that banks are strategic investors in the VC market as they use their VC investments to build relations for their lending activities.

In addition, certain investors, most notably pension fund managers and government related agencies, use private equity to "stimulate" local economies (see Lerner, Schoar, and Wong, 2007). For example, the European Investment Fund (EIF) invested in over 200 private equity funds with the objective to "commit to the development of a knowledge-based society, centered on innovation, growth and employment, the promotion of entrepreneurial spirit, regional development and the cohesion of the Union."

We cannot estimate whether LPs are satisfied ex post with the total outcome (investment performance and additional benefits). Nor do we know how much these side benefits explain the current puzzle. It is, nonetheless, important to bear in mind that there may be positive externalities of investing in private equity for, at least, certain investors.

7. Conclusion

This paper sheds light on the performance of private equity funds. Our findings show the importance of accounting for sample selection bias and of treating NAVs of old funds with circumspection for evaluating private equity fund performance. We also argue that performance should be evaluated with appropriately weighted profitability indices. Furthermore, we show that using average IRRs biases performance upward. We find that an average net-of-fees performance is lower than that of the S&P 500 by 3% per year, but gross-of-fees performance is above that of the S&P 500 by 3% per year. Adjusting for risk decreases performance by about 3% per year, bringing alpha net-of-fees to -6% per year. It is important to acknowledge, however, that performance estimates are only reliable for mature funds and that the majority of private equity funds have not reached maturity. Hence our understanding of this asset class will develop over the years as a more comprehensive sample of returns data becomes available. Our results, however, indicate that newly raised funds have a performance similar to that of our sample funds at the same age. It is also important to keep in mind that there is a wide dispersion in performance and that performance is predictable to a certain extent.

Appendix: Press Accounts of Past Performance

April 10, 2006, *Financial Times*:

"Mark O'Hare, managing partner of Private Equity Intelligence, a consultancy, said: 'Private equity has delivered good returns net of fees ...'"

April 3, 2006, *Financial Times*:

"Private equity can seemingly do no wrong in investors' eyes. The industry is raising record amounts and returns have outstripped those from equity markets in the past few years. (...) Antoine Drean, managing partner of Triago, a private equity placement firm based in Paris, says: 'For people looking in the rear mirror, buyouts look great performance wise.'"

September 26, 2005, *Financial Times*:

(Echoes of the survey of big U.K. investors from the Center for Management Buyout Research)

"The main reason for investing in private equity is to boost returns, with diversification a secondary reason. The survey shows investors hope to make an average annual net return of 12.8 percent from their private equity investments. They also expect private equity to return 4.2 percentage points more than public equity investments. (...) Their expectations have mostly been met. Nearly two-thirds said their actual returns were in line with expectations, and 13 percent did better than expected. But 23 percent were disappointed, saying the return was lower than expected."

July 25, 2005, *Financial Times*:

"Despite years of good performance, private equity is still regarded as a risky asset and sidelined by many pension funds."

References

- Aragon, G. 2007. Share Restrictions and Asset Pricing: Evidence from the Hedge Fund Industry. *Journal of Financial Economics* 83:33–58.
- Brav, A., and P. Gompers. 1997. Myth or Reality? The Long-Run Underperformance of Initial Public Offerings: Evidence from Venture and Non-Venture Capital-Backed Companies. *Journal of Finance* 52:1791–821.
- CalPERS online report, "AIM Program Fund Performance Review," <http://www.calpers.ca.gov/index.jsp?bc=/investments/assets/equities/aim/private-equity-review/aim-perform-review/home.xml> (accessed March 5, 2008).
- Cao, J. X., and J. Lerner. 2006. The Performance of Reverse Leveraged Buyouts. Working Paper, Boston College and Harvard Business School.
- Cochrane, J. 2005. The Risk and Return of Venture Capital. *Journal of Financial Economics* 75:3–52.
- Fung, W., D. A. Hsieh, N. Y. Naik, and T. Ramadorai. 2006. Hedge Funds: Performance, Risk, and Capital Formation. Working Paper No. 5565, CEPR.
- Gompers, P., and J. Lerner. 1999. An Analysis of Compensation in the U.S. Venture Capital Partnership. *Journal of Financial Economics* 51:3–44.
- Gompers, P., and J. Lerner. 2000. Money Chasing Deals? The Impact of Fund Inflows on Private Equity Valuations. *Journal of Financial Economics* 55:281–325.
- Hege, U., F. Palomino, and A. Schwienbacher. 2006. Venture Capital Performance: The Disparity Between Europe and the United States. Working Paper No. 001, RICAPE.
- Hellmann, T., L. Lindsey, and M. Puri. 2005. Building Relationships Early: Banks in Venture Capital. Working Paper, Duke University.
- Hochberg, Y., A. Ljungqvist, and Y. Lu. 2007. Venture Capital Networks and Investment Performance. *Journal of Finance* 62:251–301.
- Hwang, M., J. Quigley, and S. E. Woodward. 2005. An Index for Venture Capital, 1987–2003. Contributions to Economic Analysis and Policy 4, Article 13.
- Jones, C. and M. Rhodes-Kropf. 2003. The Price of Diversifiable Risk in Venture Capital and Private Equity. Working Paper, Columbia University.
- Kaplan, S. N., and A. Schoar. 2005. Private Equity Performance: Returns, Persistence, and Capital Flows. *Journal of Finance* 60:1791–823.
- Lerner, J., F. Hardyman, and A. Leamon. 2004. *Venture Capital and Private Equity: A Casebook*, 3rd ed. New York: John Wiley & Sons.

Lerner, J., A. Schoar, and W. Wong. 2007. Smart Institutions, Foolish Choices? The Limited Partner Performance Puzzle. *Journal of Finance* 62:731–64.

Ljungqvist, A. and M. Richardson. 2003. The Cash Flow, Return, and Risk Characteristics of Private Equity. Working Paper No. 9495, NBER.

Metrick, A., and A. Yasuda. 2007. Economics of Private Equity Funds. Working Paper, Wharton Business School.

Moskowitz, T., and A. Vissing-Jorgensen. 2002. The Private Equity Premium Puzzle. *American Economic Review* 92:745–78.

Sørensen, M. 2007. How Smart Is Smart Money? A Two-Sided Matching Model of Venture Capital, *Journal of Finance* 62:2725–2762.

Wermers, R. 2000. Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transactions Costs, and Expenses. *Journal of Finance* 55:1655–95.

Copyright of Review of Financial Studies is the property of Society for Financial Studies and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.