

Private Equity Performance: Returns, Persistence, and Capital Flows

STEVEN N. KAPLAN and ANTOINETTE SCHOAR*

ABSTRACT

This paper investigates the performance and capital inflows of private equity partnerships. Average fund returns (net of fees) approximately equal the S&P 500 although substantial heterogeneity across funds exists. Returns persist strongly across subsequent funds of a partnership. Better performing partnerships are more likely to raise follow-on funds and larger funds. This relationship is concave, so top performing partnerships grow proportionally less than average performers. At the industry level, market entry and fund performance are procyclical; however, established funds are less sensitive to cycles than new entrants. Several of these results differ markedly from those for mutual funds.

THE PRIVATE EQUITY INDUSTRY, primarily venture capital (VC) and buyout (LBO) investments, has grown tremendously over the last decade. While investors committed less than \$10 billion to private equity partnerships in 1991, they committed more than \$180 billion at the peak in 2000 (see Jesse Reyes, *Private Equity Overview and Update 2002*). Despite the increased investment in the private equity asset class and the potential importance of private equity investments for the economy as a whole, we have only a limited understanding of private equity returns, capital flows, and their interrelation. One of the main obstacles has been the lack of available data. Private equity, as the name suggests, is largely exempt from public disclosure requirements.

In this paper, we make use of a novel data set of individual fund performance collected by Venture Economics (VE).¹ The VE data set is based on voluntary reporting of fund returns by the private equity firms (or general partners (GPs)) as well as their limited partners (LPs). We study three issues with this data set that have not been closely examined before.

First, we investigate the performance of private equity funds. On average, LBO fund returns net of fees are slightly less than those of the S&P 500; VC

*Kaplan is at the University of Chicago Graduate School of Business and at the NBER; Schoar is at the Sloan School of Management at MIT, and at the NBER, and the CEPR. We thank Ken Morse at the MIT Entrepreneurship Center and Jesse Reyes from Venture Economics for making this project possible. We thank Eugene Fama, Laura Field, Josh Lerner, Alexander Ljungqvist, Jun Pan, Matt Richardson, Rex Sinquefeld, Rob Stambaugh (the editor), Rebecca Zarutskie, an anonymous referee, and seminar participants at Alberta, Arizona State, Chicago, MIT, NBER Corporate Finance, NYSE-Stanford Conference on Entrepreneurial Finance and IPOs, NYU, and USC for helpful comments. Data for this project were obtained from the VentureExpert database collected by Venture Economics.

¹ We thank Jesse Reyes from VE for making the data available.

fund returns are lower than the S&P 500 on an equal-weighted basis, but higher than the S&P 500 on a capital weighted basis.² These results combined with previous evidence on private equity fees, however, suggest that on average, both types of private equity returns exceed those of the S&P 500 gross of fees. We also find large heterogeneity in returns across funds and time.

Second, we document substantial persistence in LBO and VC fund performance. General partners (GPs) whose funds outperform the industry in one fund are likely to outperform the industry in the next and vice versa. We find persistence not only between two consecutive funds, but also between the current fund and the second previous fund. These findings are markedly different from the results for mutual funds, where persistence has been difficult to detect and, when detected, tends to be driven by persistent underperformance rather than overperformance.³ We investigate whether selection biases, risk differences, or industry differences can explain the results and conclude that they are unlikely to do so.

Third, we study the relation of fund performance to capital flows, fund size, and overall GP survival. We analyze the relation of a fund's track record to capital flows into individual GPs and the industry overall. Fund flows are positively related to past performance. In contrast to the convex relationship in the mutual fund industry, however, the relationship is concave in private equity (see Chevalier and Ellison (1997), Sirri and Tufano (1998), and Chen et al. (2003)). Similarly, new partnerships are more likely to be started in periods after the industry has performed especially well. But funds and partnerships that are raised in boom times are less likely to raise follow-on funds suggesting that these funds perform poorly. A larger fraction of fund flows during these times, therefore, appears to go to funds that have lower performance, rather than top funds. Finally, the dilution of overall industry performance in periods when many new funds enter is mainly driven by the poor performance of new entrants. The performance of established funds is less affected.

In the last section of this paper, we discuss possible explanations for our findings. Underlying heterogeneity in the skill and quality of GPs could lead to heterogeneity in performance and to more persistence if new entrants cannot compete effectively with existing funds. Several forces might make it difficult to compete with established funds. First, many practitioners assert that unlike mutual fund and hedge fund investors, private equity investors have proprietary access to particular transactions; that is, "proprietary deal flow." In other words, better GPs may be able to invest in better investments. Second, private equity investors typically provide management or advisory inputs along with capital. If high-quality GPs are scarce, differences in returns between funds

² These results and most of the analyses that follow do not explicitly adjust for differences in systematic risk or liquidity risk. We discuss this in some detail in the text.

³ See Carhart et al. (2002) for a comprehensive review of this topic and Berk and Green (2002) for a model of mutual funds returns and capital flows. Our findings on persistence also differ from those for hedge funds, which provide little or modest evidence of persistence. See Bares, Gibson, and Gyger (2002), Brown, Goetzman, and Ibbotson (1999), Edwards and Cagalyan (2001), and Kat and Menexe (2002).

could persist.⁴ Third, there is some evidence that better VCs get better deal terms (e.g., lower valuations) when negotiating with startups (see the paper by Hsu (2004)). A startup would be willing to accept these terms if some investors provided superior management, advisory, or reputational inputs.

If heterogeneity in GP skills drives the persistence results, it is puzzling that the returns to superior skill are not appropriated by the GPs through higher fees and larger funds, as has been suggested for mutual funds (see Berk and Green (2002)). From Gompers and Lerner (1999), we know that compensation was relatively homogeneous during our sample period. Most funds used a compensation scheme of a 1.5–2.5% annual management fee and a 20% carried interest or share of the profits. To the extent that there are systematic differences, Gompers and Lerner (1999) find that profit shares are higher for older and larger GPs, the GPs that tend to perform well. Alternatively, GPs could try to increase their compensation by growing the size of the fund.⁵ But we find that on average, the top performing funds grew proportionally slower than the lower performing funds in our sample period.

Our results suggest that competitive forces did not drive away persistence in our sample period. It is possible these results will not survive the late 1990s, a period of substantial growth in fund size and increases in the carried interest of top performing VC funds. Unfortunately, the net effect of these changes will not be clear for several years until the returns of those funds have been realized.⁶

I. Related Literature on Private Equity

Private equity investing is typically carried out through a limited partnership structure in which the private equity firm serves as the GP. The LPs consist largely of institutional investors and wealthy individuals who provide the bulk of the capital. The LPs commit to provide a certain amount of capital to the fund. The GP then has an agreed time period in which to invest the committed capital—usually on the order of 5 years. The GP also has an agreed time period in which to return capital to the LPs—usually on the order of 10–12 years in total. Each fund or limited partnership, therefore, is essentially a closed end fund with a finite life. When the GP exhausts a substantial portion of a fund's committed capital, the GP typically attempts to obtain commitments for a subsequent (and separate) fund.

There is a growing literature studying the economics of the private equity industry. Most of those studies have focused either on aggregate trends in private equity or on the relation between GPs and entrepreneurs. This restriction is mainly due to the difficulty of obtaining information on individual fund performance. Two recent exceptions are Jones and Rhodes-Kropf (2003) and

⁴ Please see Hellman and Puri (2002) and Kaplan and Strömberg (2004) for examples.

⁵ For related evidence for mutual funds, see Chen et al. (2003).

⁶ Interestingly, in the last several years, many of the top VCs have voluntarily returned large fractions of the committed capital in their most recent funds to their LPs, most likely because of concern over the effect of poor performance on their reputations.

Ljungqvist and Richardson (2002) who study private equity returns at the fund level. We discuss their results and the comparison to the current study in some detail in our section describing average returns.

Gompers and Lerner (1998) look at aggregate performance and capital flows. The authors find that macroeconomic factors, such as past industry performance and overall economic performance as well as changes in the capital gains tax or ERISA provisions, are related to increased capital flows into private equity.

Cochrane (2003) characterizes VC returns based on the economics of individual investments in portfolio companies. He finds that venture returns are very volatile; later stage deals have less volatility than early stage deals; returns have a market risk or beta of 1.7; and (arithmetic) returns (gross of fees) show a highly positive alpha (32% per year) over his sample period.

Papers that focus on the relation between GPs and entrepreneurs include Kaplan and Strömberg (2003), who document the structure of incentive contracts between VCs and entrepreneurs. Gompers and Lerner (2000) suggest that the valuation of individual deals is affected by overall macroeconomic conditions and the degree of competition in the VC industry.

II. Data

The data for this study have been obtained from VE. Venture Economics collects quarterly information on individual funds in the private equity industry. The data set is based on voluntary reporting of fund information by the GPs as well as by their LPs. Venture Economics claims that because they receive information from both the GPs and the LPs, there is “little opportunity for inconsistent reporting.” Given the private nature of the data, we cannot validate this statement. However, we believe that if there is a bias it would most likely take the form of underreporting by worse performing funds. If such a bias were present, this would create an upward bias on our results on average returns. As we argue later, we believe such a bias would also create a downward bias on persistence. After presenting our main results, we discuss and test for this and other potential biases in the data.

The sample covers the years 1980–2001. Because of the rapid industry growth in the 1990s, the earlier years contain relatively fewer fund observations. The VE data for each fund include the quarterly performance measures. These measures are the internal rate of return (IRR), the cumulative total value to paid-in capital (TVPI), and the distributed total value to paid-in capital (DPI). Venture Economics also collects the quarterly cash flows in and out of each fund for the life of the fund or through the end of 2001. All these performance measures, as well as the cash flows, are reported net of management fee and carried interest. We do not know the identities of the particular GPs, but we do know the sequence number of each fund, that is, if the fund is the first, second, and so forth, raised by the particular GP.

Throughout this paper, we use two samples of the data. In the main part of the analysis, we include funds: (1) that have been officially liquidated; or

(2) whose returns are unchanged for at least the final six quarters we observe. These criteria should ensure that the funds we include are largely liquidated and that the performance measures we calculate are based almost entirely on cash flows to LPs rather than subjective estimates of value by the GPs. We also exclude funds with less than \$5 million of committed capital in 1990 dollars to focus on economically meaningful funds. We obtain 746 funds that satisfy these criteria, consisting largely of funds started before 1995.

We also use a larger sample of funds that have either been officially liquidated or were started before 1997. Again we exclude funds with less than \$5 million of committed capital in 1990 dollars. Using these sample selection criteria, we obtain a sample of 1,090 funds. Because these funds are not all fully realized and we cannot reliably calculate performance for all the funds, we use the VE reported IRRs. This sample is less likely to be subject to the look ahead bias described in Carhart et al. (2002). We report most of our results using the smaller sample. Unless otherwise noted, however, our findings are qualitatively unchanged when we use the second, larger sample.

Finally, in the analysis of persistence, we also use fund performance data that have recently been made public by several large public LPs (CALPERS, University of California, UTIMCO, and the University of Michigan). These sources yield only 150 funds that have been largely liquidated and include 42 GPs with more than one liquidated fund in the sample. (We only use funds raised before 1998 to proxy for liquidation because these sources do not provide explicit information about whether a fund has been liquidated). For these funds, IRRs are reported, but cash flows are not, making it impossible to verify the IRRs and make any market adjustments. The disadvantages of this sample are that the data quality is relatively poor, the sample size is small, and the LPs may not be representative. The advantage, however, is that there is no question of selective GP reporting.

III. Descriptive Statistics

Columns 1–3 of Table I report descriptive statistics for the subsample of 746 funds that are largely liquidated and for which we have calculated performance measures. Roughly 78% of the funds are VC funds while 22% are LBO funds. To get a sense of potential selection bias in our sample of fund returns, columns 4–6 of Table I report the same statistics for all 1,814 funds that are described in VE as having been raised before 1995. We exclude funds that have a vintage year of 1995 or later to match the sample period of the funds we use in our analyses. Our subsample covers about 40% of the funds in VE over the same time period. Roughly 50% of the funds raised do not provide performance data. The remaining 10% of the funds are not fully liquidated. The funds for which we have performance measures are larger on average than the funds in the full sample.

The average size of the funds in our sample is \$172 million (all figures in 1990 dollars), with VC funds being substantially smaller than LBO funds, \$103 million versus \$416 million. These compare to average fund size in the full

Table I
Descriptive Statistics

In columns 1–3, we restrict the sample to funds for which all relevant cash flows have been realized. This restricts the sample mainly to funds that were started before 1996. In the following, we base most of the performance analysis on this sample of funds with realized performance. Size is measured as the dollar amount of capital that is committed to a fund. The first entry in this cell is mean fund size, and the second entry is the standard deviation. Sequence is the sequence number of a fund. Fraction first, second, and third indicates the fraction of funds in the overall sample that are first-time, second-time, and third-time funds, respectively. Columns 4–6 are based on the full sample of private equity funds in the VE database over the equivalent time period. We exclude funds that are not private equity funds and those that have missing information on size and year of closing.

Sample	Funds with Performance Data			Full Sample of VE Funds		
	All Funds	VC Funds	Buyout Funds	All Funds	VC Funds	Buyout Funds
Size	172.2 (378.1)	102.9 (138.6)	415.79 (612.1)	115.5 (279.2)	53.3 (78.6)	261.5 (430.4)
Fraction First	0.41	0.38	0.50	0.40	0.39	0.42
Fraction Second	0.23	0.24	0.21	0.21	0.21	0.22
Fraction Third	0.14	0.16	0.10	0.13	0.13	0.11
No. of Observations	746	577	169	1,814	1,272	542

Standard deviations are in parentheses.

sample of \$116 million for all funds, \$53 million for VC funds, and \$262 million for LBO funds. These averages imply that our performance sample includes 88% of capital committed to VC funds and 49% of capital committed to LBO funds.⁷

Table I also documents the fraction of first-, second-, and third-time funds in the two samples. In the sample with returns, 41% of the funds are first-time funds, 23% are second-time funds, and 14% are third-time funds. The remaining 22% are funds with higher sequence numbers. The corresponding percentages for the full sample are similar: 40% are first-time funds, 21% are second-time funds, 13% are third-time funds, and 25% are funds with higher sequence numbers.

One potential bias in our returns sample, therefore, is toward larger funds. We also oversample first-time funds for buyouts. 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. We will address these issues in more detail in the next section.

⁷ As mentioned previously, we obtain similar results with the larger sample of 1,090 funds. That larger sample includes 92% of capital committed to VC funds and 54% of capital committed to LBO funds for all funds started before 1997.

A. Private Equity Performance

In this section, we describe private equity performance and compare it to the performance of the S&P 500. We do not attempt to adjust for differences in systematic risk in these basic analyses.⁸ We report performance at the fund level in three ways: (1) the IRR of the funds calculated by VE, (2) the IRR of the funds that we calculate ourselves using the funds' cash flows, and (3) the public market equivalent (PME). The PME compares an investment in a private equity fund to an investment in the S&P 500. We implement the PME calculation by investing (or discounting) all cash outflows of the fund at the total return to the S&P 500 and comparing the resulting value to the value of the cash inflows (all net of fees) to the fund invested (discounted) using the total return to the S&P 500. A fund with a PME greater than 1 outperformed the S&P 500 (net of all fees). We (not VE) perform the PME calculations using fund cash flows.

We think PME is a sensible measure for LPs as it reflects the return to private equity investments relative to public equities. For example, a private equity fund investing \$50 million in March 1997 and realizing \$100 million in March 2000 would have generated an annualized IRR of 26%. However, an LP would have been better off investing in the S&P 500 because \$50 million in the S&P 500 would have grown to \$103.5 million over that period. The PME of 0.97 (or $100/103.5$) for this investment reflects the fact that the private equity investment would have underperformed the S&P 500. Alternatively, a private equity fund investing \$50 million in March 2000 and realizing \$50 million in March 2003 would have generated an IRR of 0%. However, an LP would have been better off investing in the private equity fund because \$50 million invested in the S&P 500 would have declined to \$29.5 million over that period. The PME of 1.69 (or $50/29.5$) for this investment reflects the fact that the private equity investment would have outperformed the S&P 500.

Before proceeding, we want to come back to the issue of PME and systematic risk or beta. If private equity returns have a beta greater (less) than 1, PME will overstate (understate) the true risk-adjusted returns to private equity. In most of our measures of performance, we do not attempt to make more complicated risk adjustments than benchmarking cash flows with the S&P 500 because of the lack of true market values for fund investments until the investments are exited. Instead, in the analysis that follows, particularly the persistence regressions, we consider how differences in risk might affect our results and attempt to control for observable differences such as industry composition and the stage of investment.

Table II reports the three different performance measures for all private equity funds, VC funds only, and LBO funds for the 746 funds with largely complete cash flow data. The first number in each cell is the median return;

⁸ This is a realistic if not appropriate comparison for institutional investors who invest in private equity expecting returns to exceed public equity returns. We address issues of risk adjustment in the following section.

Table II
Private Equity Returns: Cash Flow Based

Fund IRRs and PME (Public Market Equivalent) are calculated based on the actual cash flows of the funds. For each cell, we report four different entries: the first is the median return, the second the average return, the third the standard deviation, and finally we report the distribution of returns at the 25th and 75th percentiles. The first panel of the table reports equal-weighted performance measures while the second panel reports size-weighted performance measures (where size is the amount of committed capital that a fund has). The value PME is calculated by discounting cash inflows and outflows with the returns in the public markets over the same time period. The benchmark we use here to discount funds are the returns on the S&P 500 index. Only funds that did not have a cash inflow or outflow for at least 6 quarters are included in this calculation. The variable IRR_{CF} is the IRR at the end of a fund's lifetime based on actual cash inflows and outflows. On average, this excludes funds that were started after 1994. The variable IRR_{VE} contains the IRR that are reported to VE at the end of a fund's lifetime for the sample for which we can calculate IRR_{CF} .

Sample	Equal Weighted			Size Weighted		
	All Funds	VC Funds	Buyout Funds	All Funds	VC Funds	Buyout Funds
IRR_{VE}	0.12	0.11	0.13	0.14	0.14	0.15
	0.17	0.17	0.19	0.18	0.18	0.19
	(0.32)	(0.34)	(0.27)	(0.19)	(0.19)	(0.19)
	[0.04;0.20]	[0.03;0.19]	[0.06;0.24]	[0.08;0.22]	[0.05;0.22]	[0.09;0.23]
IRR_{CF}	0.12	0.11	0.13	0.12	0.13	0.13
	0.17	0.17	0.18	0.18	0.17	0.18
	(0.31)	(0.30)	(0.22)	(0.26)	(0.31)	(0.26)
	[0.03;0.22]	[0.03;0.22]	[0.05;0.22]	[0.04;0.23]	[0.03;0.23]	[0.06;0.20]
PME	0.74	0.66	0.80	0.82	0.92	0.83
	0.96	0.96	0.97	1.05	1.21	0.93
	(0.81)	(0.69)	(0.52)	(0.70)	(0.74)	(0.65)
	[0.45;1.14]	[0.43;1.13]	[0.62;1.12]	[0.67;1.11]	[0.55;1.40]	[0.72;1.03]
No. of Observations	746	577	169	746	577	169

the next is the average return; followed by the standard deviation. The last row in each cell contains the returns at the 25th and 75th percentile.

The first three columns of Table II report the performance measures on an equal-weighted basis while the next three columns report them on a commitment value- or fund size-weighted basis. The equal-weighted median and average IRRs reported by VE over the sample period are 12% and 17%, respectively. Returns to LBO funds are slightly higher than the returns to VC funds. The IRRs that we calculate from the cash flows are virtually identical. The first three columns also indicate that the median and average funds have PMEs of 0.74 and 0.96, respectively, suggesting that private equity has returned slightly less than an investment in the S&P 500 over the sample period. The average PMEs to VC and LBO funds are roughly the same at 0.96 and 0.97.

Finally, the table is suggestive of one additional quality of private equity returns. There are large differences in the returns of individual funds. The funds at the 25th percentile show a cash flow IRR of 3% while the funds at the

75th percentile exhibit a cash flow IRR of 22% per year. The amount of variation seems qualitatively similar for all performance measures, and is greater for VC funds alone.

The value-weighted performance in the last three columns of Table II exceeds the equal-weighted performance. The VE IRRs are a median 14% and average 18% while the cash flow IRRs are a median of 12% and an average of 18%. The PME's increase to a median of 0.82 and an average of 1.05, indicating that an investment in private equity slightly outperforms the S&P 500 on average. There is a substantial difference between the average PME's for VC and LBO funds. VC funds have average PME's of 1.21 while buyout funds have average PME's of 0.93. This difference is driven by the fact that the larger VC funds of the 1990s outperformed the smaller VC funds of the 1980s while the opposite was true for LBO funds.

It is worth pointing out that the average returns net of fees of 0.96 (equal-weighted) and 1.05 (value-weighted) suggest that the average returns to private equity gross of fees in both cases exceed the S&P 500. While we do not have information on GP compensation in individual funds, we know from Gompers and Lerner (1999) that GP compensation was fairly uniform during our sample period. They find that the carried interest or profit share for VC funds is almost always 20%. (Subsequently, some of the more successful VC funds have raised their profit share to 25% and 30%.) Our discussions with industry participants indicate that the same is true for the LBO funds in our sample. Gompers and Lerner (1999) also find that the discounted value of management fees for VC funds (discounted at 10%) equals 16–19% of committed capital. Conservatively, the management fees would reduce the denominator of PME by 8% (half of the Gompers and Lerner estimate) while adding back the private equity profit share of 20% would increase the numerator by at least 5%. The effect of these two adjustments is to increase the net PME by at least 13% leading to gross PME's well above 1, both on an equal- and value-weighted basis for both VC and LBO funds.

Table III presents performance results for the 1,090 funds in the larger sample. To put all the funds on an equal footing, we use the IRR calculated by VE 5 years after the fund began. We also report the TVPI calculated by VE 5 years after the fund began where TVPI is the ratio of cumulative total value: distributed value plus estimated residual value to paid-in capital. These results reflect a somewhat greater number of more recent funds. Relative to the results for the smaller (less recent) sample, Table III shows that the VC returns are somewhat higher and the LBO returns somewhat lower reflecting the performance of more recent funds included in this sample. The TVPI results in Table III also indicate that the average private equity fund returns roughly twice the capital committed to it.

B. Performance Correlations

Tables II and III present five different measures of performance. Table IV shows the correlations of those performance measures for the sample of 746

Table III
Private Equity Returns: Mature Funds

Fund IRRs and TVPI (Total Value to Paid-in-Capital) are obtained from VE. The variable IRR_{VE5} is the reported IRR after 5 years of fund existence. We only use funds that have at least a 5-year horizon. For each cell we report four different entries: the first is the median return, the second the average return, the third the standard deviation, and finally we report the distribution of returns at the 25th and 75th percentiles. The first panel of the table reports equal-weighted performance measures while the second panel reports size-weighted performance measures, where size is the amount of committed capital that a fund has under management.

Sample	Equal Weighted			Size Weighted		
	All Funds	VC Funds	Buyout Funds	All Funds	VC Funds	Buyout Funds
IRR_{VE5}	0.11	0.10	0.12	0.13	0.15	0.12
	0.18	0.20	0.14	0.18	0.30	0.13
	(0.36)	(0.40)	(0.24)	(0.32)	(0.46)	(0.23)
	[0.02;0.22]	[0.02;0.21]	[0.02;0.22]	[0.03;0.23]	[0.04;0.42]	[0.02;0.21]
TVPI	1.66	1.75	1.53	1.58	2.09	1.43
	2.24	2.42	1.83	1.97	2.87	1.61
	(2.30)	(2.54)	(1.55)	(1.83)	(2.83)	(1.00)
	[1.12;2.49]	[1.15;2.75]	[1.09;2.02]	[1.09;2.13]	[1.27;3.47]	[1.06;1.83]
No. of Observations	1,090	754	336	1,090	754	336

funds for which we can calculate all five measures. All five measures are highly correlated. For example, the IRR we calculate from cash flows is strongly positively correlated with the IRR calculated by VE (at 0.98). PME is strongly correlated with both the IRR calculated by VE and the IRR we calculate (at 0.88). These results suggest that our IRR and PME calculations accurately reflect the actual performance of the funds. Finally, the IRR calculated by VE

Table IV
Fund Performance Measures

Each entry reports the correlation between the different measures of fund performance based on a sample of 746 funds that have information about cash flow data. The variable IRR_{VE} is based on the realized IRR of funds started in a given period as reported by VE. The variable PME (Public Market Equivalent) is the ratio of capital outflows to capital inflows discounted by the cumulative returns on the S&P 500 during the same period. The variable IRR_{VE5} is based on the IRR reported to VE at the end of 5 years after the first closing of the funds. The variable TVPI (Total Value to Paid-in-Capital) is the sum of all cash outflows to LPs, divided by the sum of all cash contributions to the fund.

	IRR_{VE}	IRR_{CF}	PME	IRR_{VE5}	TVPI
IRR_{VE}	1				
IRR_{CF}	0.98	1			
PME	0.88	0.88	1		
IRR_{VE5}	0.92	0.89	0.86	1	
TVPI	0.74	0.75	0.65	0.60	1

for a fund after 5 years of existence also is strongly positively correlated with PME (at 0.86) and the IRR we calculate from actual cash flows (at 0.89). This suggests that calculated performance 5 years out is a strong indication of final or ultimate fund performance.

In the persistence and fund-raising analyses that follow, we use PME and the IRR that we calculate from cash flows. As we mentioned earlier, we repeat all our tests using both VE IRR measures, the final IRR, and the 5-year IRR, and obtain qualitatively and statistically similar results.

C. Industry Returns over Time

The performance of private equity overall in Tables II and III masks a great deal of time-series variation in that performance. In Table V, we detail that variation by presenting the average performance of the funds started each year from 1980 to 1997, weighted by the capital committed to each fund. We do not include returns prior to 1980 because we have fewer than three observations per vintage year in most years prior to 1980. Table V presents three measures of performance. For the 746 funds that are largely liquidated, the table presents the IRR and PME we calculate. For the 1,090 fund sample, the table presents the average VE IRR.

Column 1 of Table V shows a large increase in the number of funds in the mid-1980s as well as in the second half of the 1990s. The three measures of performance show a consistent pattern: VC funds performed relatively poorly in much of the 1980s with IRRs in the single digits and PMEs below 1. Since 1988, VC funds have had higher IRRs as well as PMEs that always exceed 1. LBO funds exhibit almost the reverse pattern with substantial IRRs and PMEs greater than 1 in the first half of the 1980s, followed by relatively poor performance in the first half of the 1990s. For funds raised from 1987 to 1994, the average PME of LBO funds exceeds 1.00 only in 1 year, 1990.

D. Relation to Other Studies and Implications

As mentioned earlier, Jones and Rhodes-Kropf (2003) use the same data set that we employ, but a different empirical methodology. They estimate quarterly private equity returns using GP estimates of value changes rather than realized returns at the end of a funds lifetime as we do. Their focus differs from ours in that they are interested in whether and how idiosyncratic risk is priced in private equity. They estimate alphas that are positive 4.68% per year for VC funds and 0.72% per year for LBO funds but not statistically significant using value-weighted regressions. They find betas of 1.80 for VC funds and 0.66 for LBO funds. The results for VC funds are qualitatively similar to ours. Our average PME of 1.21 translates into a cumulative alpha of 21% over the life of the VC fund (assuming a beta of 1). Although we cannot calculate an annual alpha given the nature of our data, the cumulative 21% is the same order of magnitude as the annual 4.68%. Our results for LBO funds are more negative with an average PME of 0.93 (again assuming a beta of 1). This may be driven

Table V
Private Equity Returns by Vintage Year

The variable IRR_{VE} is based on the realized IRR at the end of a fund's lifetime (as calculated by VE) for funds that were started in a given period. The variable IRR_{CF} is the IRR that we calculate based only on the actual cash flows of a fund. We base our IRR calculations only on the actual cash flows of funds, not reported net asset value. Therefore, we concentrate on funds that have already realized most of their investments. The variable PME (Public Market Equivalent) is the ratio of capital outflows to capital inflows discounted by the cumulative returns on the S&P 500 during the same period. Again we base our calculations only on actual cash flows. All fund performance measures are weighted by the capital under management. The variable Obs represents the number of observations per year. Since we have less than three observations per year for the buyout funds prior to 1983, we only start reporting annual performance measures for these funds in 1983.

Sample Year	All Funds				VC Funds				Buyout Funds						
	Obs	IRR _{VE}	Obs	IRR _{CF}	PME	Obs	IRR _{VE}	Obs	IRR _{CF}	PME	Obs	IRR _{VE}	Obs	IRR _{CF}	PME
1980	22	0.23	21	0.22	0.99	20	0.21	20	0.22	0.95	-	-	-	-	-
1981	24	0.13	18	0.14	0.68	23	0.09	18	0.10	0.51	-	-	-	-	-
1982	29	0.05	25	0.05	0.35	28	0.05	23	0.05	0.35	-	-	-	-	-
1983	65	0.14	54	0.15	0.71	59	0.08	48	0.07	0.53	6	0.33	6	0.30	1.06
1984	74	0.15	63	0.16	0.89	68	0.06	57	0.07	0.54	6	0.29	6	0.29	1.30
1985	59	0.23	49	0.23	1.24	47	0.09	37	0.10	0.73	12	0.34	12	0.18	1.00
1986	61	0.13	52	0.14	0.91	45	0.10	36	0.10	0.76	16	0.17	16	0.18	1.13
1987	97	0.14	85	0.16	0.84	68	0.12	63	0.12	0.98	29	0.14	22	0.16	0.84
1988	71	0.12	63	0.14	0.90	48	0.18	42	0.20	1.16	23	0.11	21	0.13	0.79
1989	79	0.18	67	0.16	1.01	54	0.18	45	0.18	1.03	25	0.18	22	0.14	1.00
1990	45	0.21	36	0.22	1.18	27	0.27	20	0.29	1.53	18	0.18	14	0.20	1.05
1991	24	0.20	15	0.16	0.95	18	0.26	11	0.22	1.13	6	0.14	6	0.14	0.87
1992	50	0.23	35	0.23	0.99	26	0.26	18	0.32	1.31	24	0.22	17	0.09	0.79
1993	67	0.23	56	0.21	1.09	39	0.30	45	0.35	1.65	28	0.20	11	0.18	0.84
1994	68	0.23	55	0.38	1.45	42	0.32	49	0.48	1.81	26	0.17	6	0.19	0.89
1995	70	0.18	52	0.20	1.14	43	0.44	-	0.54	2.05	27	0.10	7	0.07	0.62
1996	67	0.19	-	-	-	31	0.61	-	-	-	36	0.05	-	-	-
1997	109	0.13	-	-	-	58	0.41	-	-	-	51	0.06	-	-	-

by the different sample size and the fact that we restrict our sample to realized returns.

Ljungqvist and Richardson (2002) study the returns to investments in 73 venture and buyout funds by one large LP in funds raised from 1981 to 1993. They find that the funds in their sample (19 VC funds and 54 LBO funds) outperform the equity market and have positive alphas. They estimate betas of roughly 1.10 for VC and LBO funds. The results for VC funds, despite the small sample, are qualitatively similar to ours. The results for LBO funds are more positive than ours (and those of Jones and Rhodes-Kropf (2003)). The primary differences in our samples (aside from the number of funds) is that most of the funds in Ljungqvist and Richardson (2002) were raised in the 1980s, a period for which we also find higher returns for LBO funds. Moreover, they undersample first-time funds; only 29% of their funds are first-time funds. On the other hand, our LBO sample oversamples first-time funds; 50% of our LBO funds are first-time funds, relative to the VE universe of 40% first-time funds.

What can we conclude from these two studies and ours? First, the results are consistent with VC funds having generated positive alphas over the estimated time period. This conclusion, however, is by no means certain, as all three studies potentially suffer from a positive selection bias and all three may understate systematic risk. At the same time, the results for LBO funds are more ambiguous. It is worth noting that all three studies assume betas of LBO funds or portfolio companies roughly equal to 1. We believe it is possible that the systematic risk for LBO funds exceeds 1 because these funds invest in highly leveraged companies.

IV. Characteristics of Fund Returns

A. Relation of Performance to Fund Size and Sequence Number

In this section, we explore how realized fund returns correlate with partnership and fund characteristics. The basic empirical specification is as follows:

$$PME_{it} = \alpha_t + \beta(FundSize_{it}) + \lambda(Sequence_{it}) + \gamma_{VC} + \epsilon_{it}, \quad (1)$$

where PME_{it} is calculated from the cash inflows and outflows of each fund, $FundSize_{it}$ is the logarithm of the capital committed to the fund, $Sequence_{it}$ is the logarithm of the sequence number of the fund (later funds of the same private equity partnership), and γ_{VC} is a dummy equal to 1 if the partnership is a venture capital firm and 0 otherwise. We also include year fixed effects in all specifications to control for the large interyear variation in returns. In the reported regressions, standard errors are corrected for heteroscedasticity and clustered at the GP level. We obtain, but do not report, lower standard errors when we cluster by year.⁹

Columns 1–4 of Table VI show the cross-sectional relations between fund performance and fund characteristics. The estimates in column 1 indicate that

⁹ We thank Gene Fama for suggesting that we do this.

Table VI
Fund Performance and Fund Characteristics

The dependent variable is realized PME (Public Market Equivalent). The variable PME is calculated by discounting the actual cash outflows and cash inflows that the fund received with the returns on the S&P 500 over the same time period and forming the ratio of the discounted cash inflows over the discounted outflows. Since we only include funds for which the majority of the cash flows have been realized, on average this restricts the sample to funds that were started before 1996. All the data are obtained from VE. Size is the amount of capital a fund has under management. Sequence is the sequence number of a fund. The value VC Dummy is equal to 1 if the fund is a venture capital fund and 0 for buyout, LBO, and mezzanine funds. First is a dummy equal to 1 if the fund is a first-time fund and 0 otherwise. Size Spline 1–4 is a piecewise regression where we allow for different slopes of the log of size.

	Dependent Variable: PME (Public Market Equivalent)									
	Full Sample					VC Only		Buyout Only		
log(Size)	0.08 (0.03)	0.53 (0.11)		0.09 (0.03)	−0.14 (0.06)	0.30 (0.20)	0.46 (0.17)	−0.14 (0.08)	0.08 (0.06)	−0.11 (0.21)
log(Size) ²		−0.05 (0.01)				−0.05 (0.03)	−0.04 (0.02)		−0.01 (0.10)	
log(Sequence)	0.14 (0.06)	0.03 (0.01)	0.16 (0.08)		−0.37 (0.25)	−0.43 (0.28)	0.02 (0.18)	−0.18 (0.32)	0.19 (0.20)	−0.26 (0.39)
log(Sequence) ²		0.07 (0.09)				0.07 (0.12)	0.09 (0.11)		−0.06 (0.11)	
First Dummy				−0.10 (0.04)						
Size Spline 1			0.10 (0.07)							
Size Spline 2			0.36 (0.17)							
Size Spline 3			0.19 (0.15)							
Size Spline 4			−0.19 (0.07)							
VC Dummy	0.24 (0.09)	0.20 (0.09)	0.28 (0.09)	0.05 (0.24)	0.03 (0.23)					
Firm F.E.	No	No	No	No	Yes	Yes	No	Yes	No	Yes
Year F.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.17	0.18	0.19	0.18	0.29	0.30	0.25	0.16	0.13	0.01
No. of Observations	746	746	746	746	398	398	577	577	169	169

Standard errors are in parentheses and are adjusted for serial correlation and heteroskedasticity.

larger funds and higher sequence numbers funds have significantly higher realized PMEs. The estimates also confirm that the VC funds in our sample have higher PMEs on average than LBO funds. The point estimate on the VC dummy is 0.24 with a standard error of 0.09. In column 2 of Table VI, we include squared terms of Fund Size and Sequence number along with the direct effects in the regression to analyze the functional form of this relation. The point estimate on the linear term of (log) Fund Size increases significantly when including the squared term, and the coefficient on the squared (log) Fund Size is negative and significant. This suggests a concave relation between Fund Size and performance. While larger funds have higher PMEs, when funds become very large, performance declines.

The relationship between fund performance and the sequence number of the fund is convex, although not significantly so. The coefficient on the squared term of Sequence Number is positive, but not significant while the coefficient on the linear term is smaller, but remains positive and significant.

To check the robustness of the size result, column 3 reports the results of a piecewise regression that allows for different slope coefficients across four different ranges of the size distribution. The results show a concave pattern similar to the quadratic specification.

To check the robustness of the sequence number relation, we estimate a different specification in column 4 in which we include a dummy variable (First-Time Fund) equal to 1 if the fund is a first-time fund. The coefficient on this dummy is negative (-0.10) and significant.

In columns 5 and 6 of Table VI, we estimate the same specifications as in columns 1 and 2, but include GP fixed effects. In the linear specification in column 5, the signs on the Fund Size and Sequence Number variables switch from positive to negative (compared to the specification without GP fixed effects). The Fund Size coefficient is significant while the coefficient on Sequence Number is not. The Fund Size result indicates that while larger funds have higher returns in the cross-section, when a given GP raises a larger fund, fund returns decline for that GP. The Sequence Number result has a similar interpretation. In the cross-section, higher sequence number funds have higher returns. However, when a given GP raises a subsequent fund, its returns decline, albeit not in a statistically significant way. When we add squared terms to the firm fixed effects regressions in column 6, we find the same concave pattern for Fund Size found in the cross-section, but with a smaller turning point.

Columns 7 and 9 estimate the quadratic regression specification of column 2 (without GP fixed effects) separately for VC and LBO funds, respectively. The concave relationship with Fund Size is present for both type of funds, but statistically significant only for VC funds. Sequence number is no longer significant for either type of fund. Columns 8 and 10 repeat the linear specification in column 5 with GP fixed effects separately for VC and LBO funds, respectively. The signs and magnitude of the coefficients are similar for both types of funds, but, again, with greater standard errors than in the regression that uses the combined data.

B. Persistence of Fund Performance

We now turn to persistence in fund performance. The results in Table VI provide an initial indication of persistence in the private equity industry. The R^2 of the regressions in columns 1 and 2 increases by roughly 13% when we include firm fixed effects, in columns 5 and 6. An F -test on the joint significance of these GP fixed effects is strongly significant (not reported). The importance of firm fixed effects suggests that GPs vary systematically in their average performance.

To test persistence more directly, we use a parameteric approach. We extend the basic specification of the previous section to include lagged performance

Table VII
Persistence of Fund Performance

The dependent variable is realized PME (Public Market Equivalent). The variable PME is calculated by discounting the actual cash outflows and cash inflows that the fund received with the returns on the S&P 500 over the same time period and forming the ratio of the discounted cash inflows over the discounted outflows. Since we only include funds for which the majority of the cash flows have been realized, on average this restricts the sample to funds that were started before 1996. All the data are obtained from VE. The variables PME_{t-1} , PME_{t-2} , and PME_{t-3} are lagged realized PMEs of a given private equity firm's previous fund, the fund before the last, and the third previous fund, respectively. Size is the amount of capital a fund has under management. Sequence is the sequence number of a fund. The variable VC Dummy is equal to 1 if the fund is a venture capital fund and 0 for buyout, LBO, and mezzanine funds.

	Dependent Variable: PME (Public Market Equivalent)								
	Full Sample				VC Only		Buyout Only		
PME_{t-1}	0.54 (0.17)	0.49 (0.21)			0.46 (0.21)	0.69 (0.21)	0.62 (0.25)	0.17 (0.08)	0.19 (0.06)
PME_{t-2}		0.28 (0.13)	0.39 (0.14)		0.28 (0.13)		0.48 (0.26)		0.07 (0.05)
PME_{t-3}				0.32 (0.24)					
log(Size)					0.09 (0.05)				
log(Sequence)					0.20 (0.20)				
VC Dummy	0.41 (0.12)	0.51 (0.16)	0.46 (0.14)	0.49 (0.20)	0.60 (0.17)				
Firm F.E.	No	No	No	No	No	No	No	No	No
Year F.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Adjusted R^2	0.19	0.21	0.20	0.25	0.27	0.12	0.12	0.23	0.10
No. of Observations	398	225	225	128	225	323	184	76	41

Standard errors are in parentheses and are adjusted for serial correlation and heteroskedasticity.

as right-hand-side variables. In Table VII, we use lagged PME of the first, second, and third previous funds raised by the GP. In Table VIII, we repeat our regressions using IRRs:

$$Performance_{it} = \alpha_t + \delta(Performance_{it-1}) + \beta(FundSize_{it}) + \lambda(Sequence_{it}) + \gamma_{VC} + \epsilon_{it}.$$

(2)

Because we include the lagged PME or IRR as a right-hand-side variable, we cannot simultaneously control for firm fixed effects in the persistence regressions. In this analysis, we implicitly assume that private equity funds have a systematic or market risk equal to 1. As described earlier, due to the illiquidity and reporting of private equity performance, it is difficult to estimate the market risk of a partnership. Therefore, we do not rely on partnership betas estimated from quarterly reported data. Instead, we present several additional tests that control for observable differences in GP risk profiles.

Table VIII
Fund Performance Based on IRR

The dependent variable is realized fund IRR at the end of the fund's existence. The variable IRR is calculated based on the actual cash inflows and outflows of a fund (see text for details). The variables IRR_{t-1} and IRR_{t-2} are lagged realized returns of a given private equity firm's previous fund and the fund before the last, respectively. All other independent variables are defined as in prior tables. Panel A, columns 1–6 is based on the sample of venture and buyout funds from VE. The last two columns of Panel A are based only on returns disclosed through public sources (see text for more details). Panel B shows results separately for venture and buyout funds.

Panel A								
	Full Sample					Public Sample		
$\log(\text{Size})$	0.03 (0.01)	0.18 (0.03)						
$\log(\text{Size})^2$		-0.02 (0.00)						
$\log(\text{Sequence})$	0.43 (0.02)	-0.01 (0.03)						
$\log(\text{Sequence})^2$		0.04 (0.03)						
IRR_{t-1}			0.47 (0.13)	0.40 (0.15)		0.67 (0.15)	0.66 (0.08)	0.46 (0.09)
IRR_{t-2}				0.32 (0.15)		0.32 (0.19)		0.60 (0.29)
S&P 500					1.12 (0.34)	0.52 (0.23)		
VC Dummy	0.08 (0.03)	0.06 (0.03)	0.10 (0.03)	0.12 (0.03)	0.09 (0.03)	0.11 (0.04)	0.14 (0.05)	0.13 (0.06)
Year F.E.	Yes	Yes	Yes	Yes	No	No	Yes	Yes
Adjusted R^2	0.25	0.16	0.29	0.32	0.19	0.21	0.47	0.51
No. of Observations	746	746	398	225	746	225	126	95
Panel B								
	VC Only					Buyout Only		
$\log(\text{Size})$	0.15 (0.06)				0.01 (0.06)			
$\log(\text{Size})^2$	-0.02 (0.01)				-0.01 (0.03)			
$\log(\text{Sequence})$	-0.03 (0.05)				0.06 (0.07)			
$\log(\text{Sequence})^2$	0.04 (0.04)				-0.02 (0.04)			
IRR_{t-1}		0.36 (0.16)	0.61 (0.16)			0.53 (0.14)	0.45 (0.06)	
IRR_{t-2}		0.42 (0.21)	0.32 (0.23)			0.25 (0.18)	0.32 (0.14)	
S&P 500			1.07 (0.38)	1.23 (0.42)			0.38 (0.25)	0.41 (0.29)
Year F.E.	Yes	Yes	No	No	Yes	Yes	No	No
Adjusted R^2	0.24	0.35	0.19	0.16	0.16	0.49	0.22	0.09
No. of Observations	577	184	184	577	169	41	41	169

Standard errors are in parentheses and are adjusted for serial correlation and heteroscedasticity.

Table VII presents our basic results. We find strong persistence in fund returns across different funds for the same GP. Column 1 contains the results of a regression of PME on lagged PME, controlling for year fixed effects and a VC dummy. The coefficient on lagged PME is positive and strongly significant; the point estimate is 0.54 with a standard error of 0.17. The coefficient implies that a fund with 1% higher performance in the previous fund is associated with a 54 basis point better performance in the current fund.¹⁰ The regression in column 2 includes the performance of both the previous fund and the fund before that. Again the coefficients on both performance measures are positive and significant. The coefficients imply that a 1% increase in past performance is associated with a combined 77 basis point increase in performance in subsequent funds (the sum of the two coefficients on lagged performance).

It is possible that the current and previous fund of a particular GP have some investments in common. This could mechanically induce persistence in our sample.¹¹ To account for this possibility, column 3 of the table presents estimates of a regression that includes only the performance of the second previous fund. Because the second previous fund is typically raised 6 years before the current fund, there is likely to be little, if any overlap. The coefficient on the performance of the second previous fund is positive and significant (at the 1% level). The coefficient of 0.39 compares to that of 0.54 on the first previous fund suggesting that overlap does not drive our results. In column 4, we include only the performance of the third previous fund. The point estimate on the third previous fund performance is 0.32. Because this reduces our sample size to only 128 funds, that coefficient is not significant. It is worth noting that we do obtain a statistically significant result on the third previous fund when we use the larger sample of 1,090 funds.

An additional concern is that overlapping time periods across funds could induce some persistence. If such overlaps are important, however, persistence should decline with the amount of time that elapses between funds. In unreported regressions, we test for this possibility by interacting the PMEs for the first and second previous funds with the log of the number of years between the current fund and the respective previous fund. When we do this, the coefficients on the interaction terms are positive not negative (but statistically not significant). This result suggests that our persistence results are not caused by either investment overlap or time period overlap. If anything, this result is consistent with more persistence for GPs who invest more slowly.

In columns 6–9, we estimate the earlier persistence regressions separately for VC and LBO funds. The coefficients for the VC fund regressions are larger than those for the overall sample at 0.69, versus 0.54 for the first previous fund, and 1.10, versus 0.77 for the sum of the two coefficients on the two previous funds. The coefficients are statistically significant. The coefficients for the buyout

¹⁰ As mentioned earlier, all performance results hold if we use last reported VE IRR, cash flow IRR, and TVPI to measure performance.

¹¹ Investment across funds is likely to be more of an issue for VC funds than LBO funds because VC funded companies are more likely to require multiple equity financings than buyout funded companies.

funds, in contrast, are smaller at 0.17 for the first previous fund and 0.26 for the sum of the two previous funds. Despite the smaller magnitude and smaller sample size, the coefficient on the first previous fund is statistically significant.

Overall, the results in Table VII suggest a statistically and economically strong amount of persistence in private equity, particularly for VC funds.

C. Robustness Checks

Because the persistence results are unusual compared to mutual funds and hedge funds and because of the difficulty in controlling for systematic risk, we undertake a number of additional checks to test the robustness of our findings. In Table VIII we reestimate the persistence results measuring performance using fund IRRs based on the cash inflows and outflows to the funds. The regression in column 1 of panel A, Table VIII confirms that performance increases with fund size and with sequence number (without GP fixed effects). When squared terms are included in column 2, fund size remains concave, while sequence number becomes insignificant. The regressions in columns 3 and 4 of Table VIII show that our key persistent results hold for IRRs. The performance of the previous fund and the second previous fund are statistically significantly related to current fund performance, both individually and when included together. The coefficients are similar in magnitude to those using PME to measure performance.

Panel B estimates the regressions separately for VC and LBO funds. Columns 1 and 5 indicate that the concavity results with respect to fund size are mainly driven by the VC funds. But columns 2 and 6 of Panel B indicate that the IRRs persist both for VC funds and LBO funds. In contrast to the results for PMEs, the magnitudes of the coefficients on previous fund and second previous fund IRRs are similar and statistically indistinguishable for VC and LBO funds.

C.1. Differences in Risk

As discussed above, a concern about the persistence we have documented is that some GPs might consistently take on more systematic or market risk than others. High systematic risk GPs would have persistently higher returns in a rising stock market. We attempted to control for potential differences in systematic risk in several different ways by controlling for differences in average market risk, dividing funds according to investment stage focus and industry composition, and analyzing the cross-sectional dispersion of fund returns.

We attempt to adjust for average market risk by including the average annual return to the S&P 500 in the 5 years after a fund is raised and excluding year fixed effects. This allows us to control for variations in market returns over time in the private equity industry. As mentioned earlier, we cannot calculate true “betas” for individual funds because we choose not to rely on interim IRRs of a fund (that are necessarily based on subjective valuations by the fund’s GP).

In column 5 of panel A in Table VIII, we include the S&P 500 return, but do not include past performance. The coefficient on the S&P 500 is 1.12. In columns 3 and 7 of Panel B, we estimate the regression separately for VC and LBO funds. The coefficient on the S&P 500 is 1.23 for VC funds and 0.41 for buyout funds. The 1.23 for VC funds is higher than that found by Ljungqvist and Richardson (2002), but lower than those found by Jones and Rhodes-Kropf (2003) and (for individual deals) Cochrane (2003). The coefficient for buyout funds is low compared to that in Ljungqvist and Richardson (2002), but (with a standard error of 0.29) not statistically different from the 0.66 estimated by Jones and Rhodes-Kropf (2003).

In column 6 of Panel A, we include the returns of the two previous funds and the S&P 500 returns. This adjustment increases the measured persistence. The coefficient on the previous fund increases from 0.40 to 0.67 while the coefficient on the second previous fund of 0.32 is unchanged. In columns 4 and 8 of Panel B, we estimate coefficients for VC and LBO funds separately. The adjustment strengthens the persistence results for VC funds and leaves them essentially unchanged for LBO funds.

Next, we divide the sample of private equity funds into early stage, later stage, expansion stage, leveraged buyout, and mezzanine funds. If these different stages are correlated with differences in market risk, we would expect to see decreases in persistence after controlling for the differences. We obtain (but do not report in the table) qualitatively and statistically similar persistence results controlling for the different types of private equity funds.

To control for industry, we asked VE to construct measures of industry focus for the funds in our sample. The industry classes VE uses are biotech, communications and media, computer hardware, computer software and services, consumer-related, industrial/energy, internet, medical/health, and semiconductors and other electronics. Any fund that has more than 60% of its investments in one industry is classified as focused. These data were available for 412 funds in our sample.¹² We estimated the regressions in columns 6–8 controlling for industry fixed effects. We also estimated the regressions in these tables including only a dummy for whether the partnership is focused or not. In either case, the persistence results are qualitatively unchanged.

Finally, we also considered the possibility that differences in total risk might drive our results. To do so, we look at the dispersion in fund returns conditional on having been in the top, medium, or bottom tercile of performance in the previous fund. If differences in investment risk that GPs take on explain persistence, funds with high returns in the first period (which would imply they took on a lot of risk and had good return realizations) should show more dispersion in returns in the follow-on funds. In contrast, low return funds should have relatively little return dispersion in their next funds, since they make more conservative investment choices. When we look at the raw PME's, we find that the dispersion increases slightly for funds in the middle and high performance

¹² Because of disclosure concerns by VE, we could not obtain more precise information about the industry composition of the portfolio firms.

tercile. However, when we use residual PME (after controlling for size and year fixed effects), we find no increase in dispersion; if anything we see a decrease.

Overall, while we cannot rule out the possibility that differences in market and total risk drive the persistence results, we believe that such an explanation is unlikely.

C.2. Selection Bias

The last issue we consider in this section is whether selection bias and data collection issues drive our persistence results. Several mutual fund studies have found that return persistence can be affected by sample selection. Carhart et al. (2002) provide an in-depth discussion of this issue in mutual funds. Our sample differs from the mutual fund studies because we only use realized returns at or near the end of a fund's lifetime. Observations are not dropped from the data set if returns in a given period fall below a certain threshold as is sometimes the case in mutual fund data sets. In an interview in the Asset Alternative Newsletter (Asset Alternative 2002), Jesse Reyes from VE states that VE observes very few incidents of funds that stop reporting when returns worsen. Furthermore, VE does not rely solely on the GPs for performance information, but also on LPs who are not prone to this type of selection bias.

Selective reporting could create an upward bias in measured persistence if funds initially report good performance, but if performance declines, they stop reporting. We test whether GPs stop reporting performance of a fund after large (negative) changes in quarterly reported performance in the previous period. To do so, we construct a panel of quarterly reported performance for all funds in our sample from the first quarter of the fund to the final quarter in which they are liquidated. Given the nature of the data, it is impossible for us to know if a fund was really liquidated or just stopped reporting. Therefore, we construct a dummy variable equal to 1 for the last quarter in which a fund reported performance and 0 otherwise. We regress this dummy variable on quarterly reported IRR (or changes in IRR) including controls for log of fund size, sequence number, and fund fixed effects. The coefficient on IRR (change in IRR) is economically equal to 0, indicating that funds do not stop reporting in periods after an abnormally large change in performance in the prior quarter.

A second concern is that partnerships do not report the performance of subsequent funds after a particularly poor or particularly good fund. To investigate this, we test whether the likelihood of reporting the performance of a follow-on fund is a function of the previous fund's performance. We find that GPs of better performing funds are more likely to report the performance of a follow-on fund. This is not surprising since better performing GPs also have a higher probability of raising a subsequent fund. We rerun these regressions with firm fixed effects and again obtain a positive coefficient on the performance measure. Funds that experience an increase in performance from one fund to the other are more likely to report the performance of the next fund. These results do not support the hypothesis that once funds become very good they stop reporting their performance to VE.

Table IX
Transition Probabilities: Fund Performance

We sort all funds for which we have follow-on funds into performance terciles and calculate the conditional probability that a partnership's next fund will either stay in the same performance tercile, or move into one of the other two terciles. The results in Panel A are based on PME (Public Market Equivalent over the entire lifetime of the fund). For this calculation, we use 398 funds that have at least one follow-on fund in our sample of realized funds. In Panel B, we use IRR_{VE5} , the IRR of a fund after 5 years of a fund's existence. This allows us to employ a larger sample of 639 funds that have at least one follow-on fund. Again we calculate conditional Markov probabilities as in Panel A.

	Lower Tercile (%)	Medium Tercile (%)	Upper Tercile (%)
Panel A: PME			
Lower tercile	44	37	19
Medium tercile	24	34	42
Upper tercile	11	34	55
Panel B: IRR_{VE5}			
Lower tercile	49	31	20
Medium tercile	30	38	32
Upper tercile	21	31	48

It is worth pointing out that these stories of selection or survivorship biases would predict that persistence of returns is driven either by the positive end of the performance distribution or by the negative end. To test for this in Table IX, we sort all the funds for which we have follow-on funds into performance terciles. We then calculate the conditional probability that a partnership's next fund will either stay in the same performance tercile or move into one of the other two terciles. The results in Panel A measure fund performance using PME. The results in Panel B use the VE IRR at the end of 5 years. We find persistence at both ends of the distribution. In both panels, funds in the top (and bottom) terciles have at least a 44% chance of remaining in those terciles and at most a 21% chance of moving to the bottom (and top) terciles. In both panels, a chi-squared test rejects the equality of all cells at the 1% level.

Finally, we estimate the persistence regressions using publicly available data reported by public pension funds. The advantage of this sample is that funds cannot stop reporting to their LPs. As we noted earlier, while this sample is smaller and likely less reliable than the VE sample, it does provide an additional independent sample with no reporting bias. We restrict our analysis to mature funds, those funds raised before 1997. (Given the limited disclosure by LPs, we do not have information about whether a fund was truly liquidated or not.) We are able to estimate regressions using 126 subsequent funds for more than 40 different GPs. The last two columns of Table VIII show that the observed persistence in this data set is stronger than that in the VE sample. If we include only one lag of IRR, the coefficient on the previous fund IRR is 0.66 with a standard error of 0.08 (compared to 0.47 in the VE sample). When we include two lags, in column 8, the coefficient on the second lag is large and significant

(the point estimate is 0.60 and the standard error 0.29). The sum of the two lags is 1.06 compared to 0.72 in the VE sample.

While we cannot rule out that risk differences and selection biases affect our results, we believe that our evidence favors the existence of persistence in private equity.

V. Capital Flows and Fund Performance

We now analyze the relationship between past performance and the flow of capital into subsequent funds. Using mutual funds as a benchmark again, studies by Sirri and Tufano (1998) and Chevalier and Ellison (1997) indicate that funds that outperform the market experience increased capital inflows. This relationship tends to be convex; mutual funds with above average performance increase their share of the overall mutual fund market.

We estimate the relationship between fund size and performance using Tobit regressions that control for the left censoring in the size variable. If poorly performing GPs are unable to raise follow-on funds, a simple OLS estimator will be biased because poor first-time funds will drop out of the sample. Yet not being able to raise a fund at all is clearly an (extreme) outcome of poor initial performance.

In Table X, we find that capital flows into private equity funds are positively and significantly related to past performance. Column 1 of Table X shows the basic specification controlling for sequence number and VC Dummy. (The logarithm of) Fund size is positively and significantly related to the performance of the previous fund. Fund size also increases with sequence number controlling for performance. This suggests that the sequence number may also be picking up some aspect of fund quality and/or past performance. The sequence number result also indicates that GPs of higher sequence number funds are better able to survive the poor performance of one particular fund. Finally, VC funds are smaller than LBO funds. In column 2, we include lagged fund size to control for partnership specific effects, given that we cannot use firm fixed effects in the specification with lagged return variables. The only effect of this control is to render the VC dummy insignificant. The coefficient on previous performance is still positive and significant.

Parallel to our analysis of persistence in returns, we also consider the relationship between fund size and performance in the second previous fund. Column 3 shows that current fund size is positively and significantly related to the performance of each of the two previous funds. These findings indicate that funds with persistently good performance are especially favored in the fund-raising process. This timing effect makes sense given that returns take some time to realize in the private equity industry. When a GP raises the first follow-on fund, investors may not have learned completely about the true performance of the previous fund. By the time of the second follow-on fund, LPs are able to evaluate a longer track record of the GP.¹³

¹³ We also estimated, but do not report, probit regressions of the likelihood a GP raises a next fund. The likelihood of raising a next fund is negatively related to the performance of the previous fund.

Table X
Fund Size and the Firm's Track Record

The dependent variable (fund size) is the logarithm of the amount of capital committed to the next fund of a partnership. We estimate a Tobit regression, since the size variable is censored at 0. If a partnership does not raise a follow-on fund, the size of the next fund is 0. The variables PME_{t-1} and PME_{t-2} are the Public Market Equivalent (for construction of PME please see Table VI) of the previous fund and the one prior to this, respectively. We only include funds that are raised prior to 1996, and only one observation per fund. Lagsize is the amount of capital a fund has under management in the fund before the current one. Sequence is the sequence number of a fund. VC Dummy is equal to 1 if the fund is a venture capital fund and 0 for buyout, LBO, and mezzanine funds.

	Dependent Variable: Logarithm of Fundsize								
	Full Sample				VC Only		Buyout Only		
PME _{t-1}	0.27 (0.09)	0.20 (0.08)	0.23 (0.09)	0.88 (0.18)	0.49 (0.19)	0.17 (0.08)	0.89 (0.18)	1.17 (0.44)	2.32 (0.94)
PME _{t-2}			0.34 (0.16)		0.65 (0.35)				
PME ² _{t-1}				-0.11 (0.02)	-0.06 (0.02)		-0.10 (0.02)		-0.36 (0.25)
PME ² _{t-2}					-0.10 (0.05)				
log(Lagsize)		0.78 (0.09)	0.73 (0.10)	0.72 (0.09)	0.67 (0.11)	0.80 (0.09)	0.71 (0.10)	0.86 (0.20)	0.82 (0.20)
log(Sequence)	1.07 (0.14)	0.73 (0.14)	0.43 (0.21)	0.71 (0.18)	0.52 (0.21)	0.53 (0.15)	0.53 (0.15)	1.04 (0.34)	1.04 (0.34)
VC Dummy	-1.20 (0.22)	-0.25 (0.24)	-0.48 (0.27)	-0.29 (0.23)	-0.55 (0.27)				
Year F.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Pseudo R ²	0.16	0.22	0.25	0.22	0.25	0.07	0.08	0.08	0.09
No. of Observations	746	746	399	746	399	577	577	169	169

Standard errors are in parentheses and are adjusted for serial correlation and heteroskedasticity.

In column 4 of Table X, we test whether the relation between fund size and prior performance is best characterized by a linear relation. To do this, we include a squared term of the PME_{it-1} . The relationship between fund size and performance is positive but concave. The coefficient on the squared term is negative and significant. This result differs from that for the mutual fund industry where researchers find a convex relation between fund size and excess returns. In column 5, we repeat this exercise by including a squared term for the PMEs of the two previous funds along with the linear terms. We find a concave relation between fund raising and returns for both previous funds.

In columns 6 and 7, we estimate regressions for VC funds only; in columns 8 and 9, for LBO funds only. The results are economically and qualitatively identical to the results for the entire sample. The only difference is that the coefficient on the squared term for PME is not significant in the LBO fund regression, although it is negative and economically larger than the corresponding coefficient in the VC funds regression.

The findings in Table X suggest that the top performing funds in the private equity industry grow less than proportionally with the increase in performance than do the lower performers. Given that most LPs claim that the top funds are all highly oversubscribed, it seems likely that the better funds voluntarily choose to stay smaller. This result, in turn, might help explain the persistence in performance that we find. By growing relatively less rapidly than the market on a performance-adjusted basis, top funds are able to avoid moving into regions of diminishing returns.

There are at least two reasons why superior GPs might choose to do so. On the demand side, it is possible that the number of good deals in the economy is limited at each point in time. If GPs believe that they face diseconomies of some sort even on their inframarginal deals when moving down the quality curve, it could be in their interest to grow slowly. On the supply side, better funds might face constraints if GP human capital is not easily scalable and new, qualified individual GPs are scarce. Under either of these constraints, superior GPs have to trade off staying small (and having high returns) or growing at the same speed as the market (or at a higher speed), but moving down the marginal returns curve. In the next section, we look at the flow of funds and the entry and exit decisions of GPs in more detail.

VI. Market Dynamics and Entry of Funds

In this section, we analyze the overall dynamics of performance and capital flows in the private equity industry. So far, our results indicate that there is strong persistence in performance across funds. But, at the same time, mediocre performing GPs grow proportionally faster than the top GPs. This raises the question of how capital is allocated to poorly performing funds and whether this has a spillover effect on the industry overall.

A. Timing of Entry and Cyclicalities of Returns

We first consider the entry and exit of GPs into the private equity market and the fund-raising activities of existing GPs. To this end, we turn to the general database collected by VE. The benefit of this data set is that it is more comprehensive in coverage of funds than the performance data set. Venture Economics estimates that this data set covers 70% of the overall private equity market. The drawback is that we do not observe performance for all of the funds. As a result, we rely on aggregate measures of industry returns as well as information on fund size and fund sequence number (which we found to be positively related to performance in the previous analyses).

In Table XI, we look at the timing of funds raised by new private equity partnerships. We regress the logarithm of the total number of partnerships started in a given year on different measures of market returns in the current and previous year. We have 26 years of data for this exercise. In column 1 of Table XI, we relate the number of partnerships that are started each year to the returns on the Nasdaq Composite index in the current and the prior year. Lagged

Table XI
Entry of Private Equity Funds into the Industry

The dependent variable in columns 1–3 is the aggregate number of new partnerships that are started in a given year from 1975 to 2000. The dependent variable in columns 4–6 is the logarithm of the total amount of capital raised by first-time funds in a given year, again from 1975 to 2000. Nasdaq and Nasdaq_{*t*-1} are the annual returns on the Nasdaq Composite index in the current and the prior year, respectively. Similarly, S&P and S&P_{*t*-1} are the annual returns on the S&P 500 index in the current and the prior year, respectively. And finally, VC Returns and Buyout Returns are the aggregate annual returns of the venture capital and the buyout industry, respectively, as reported by VE. Since we only have venture industry returns since 1980, the number of observations in columns 3, 4, 7, and 8 drops to 20. Moreover, entering firms in columns 3 and 7 are restricted to venture capital firms. Similarly, entering firms in column 4 and 8 are restricted only to buyout firms.

Dependent Variable	Number of New PE Firms				Capital Raised by New PE Firms			
Nasdaq _{<i>t</i>}	0.47				1.53			
	(0.88)				(1.66)			
Nasdaq _{<i>t</i>-1}	1.46				3.20			
	(0.90)				(1.81)			
S&P _{<i>t</i>}		2.60				5.11		
		(1.41)				(2.66)		
S&P _{<i>t</i>-1}		2.41				5.01		
		(1.41)				(2.60)		
VC Returns _{<i>t</i>}			0.02				-0.08	
			(0.07)				(0.13)	
VC Returns _{<i>t</i>-1}			0.20				0.45	
			(0.06)				(0.13)	
Buyout Returns _{<i>t</i>}				0.00				-0.03
				(0.01)				(0.02)
Buyout Returns _{<i>t</i>-1}				-0.01				-0.01
				(0.01)				(0.02)
Adjusted <i>R</i> ²	0.02	0.11	0.55	0.21	0.08	0.13	0.56	0.21
No. of Observations	26	26	20	20	26	26	20	20

Standard errors are in parentheses and are adjusted for serial correlation and heteroskedasticity.

Nasdaq returns have a coefficient of 1.46 and a standard error 0.9 while current Nasdaq returns have a smaller coefficient of 0.47 with a standard error of 0.88. Similarly, in column 2 of Table XI, we include current and lagged returns on the S&P 500. These have a positive and significant relation (at the 10% level) to the number of partnerships that are started each year as well. The point estimate is 2.6 with a standard error of 1.4 for the current S&P returns, and the coefficient on the lagged returns is 2.41 with a standard error of 1.4. Finally, we repeat this exercise for the aggregate returns of the venture capital industry. Column 3 of Table XI shows that there is an increase in partnership starts when lagged venture capital returns are high, while the contemporaneous relation between industry returns and partnership starts is positive but not significant.

We also repeat these estimates using the aggregate amount of capital that is raised by first-time partnerships in a given year. We report the results in

columns 4–6 of the same Table XI. The results are consistent with the previous findings based on the number of new entrants. In fact, the relationship between lagged returns and partnership entry becomes stronger. This suggests not only do more partnerships decide to start up after a period in which the industry performed well, but also that first-time funds tend to raise bigger amounts of capital when the private equity industry performed well. Gompers and Lerner (1999) find similar results for aggregate industry returns.

B. Which Type of Funds Are Raised in Boom Times?

We now consider how the allocation of capital flows across funds varies in boom versus bust times. In particular, do funds that are started in boom times have systematically lower performance than those started in downturns? As mentioned before, we do not observe individual fund performance for all of the funds in the full VE data set. However, we do observe whether a GP raises a follow-on fund. Our earlier analysis indicates that the ability to raise a follow-on fund is a rough proxy for fund performance.

In Table XII, we report the results of a linear probability model in which we relate the likelihood of raising a follow-on fund to market conditions at the time the initial fund was raised. We construct a dummy variable equal to 1 if a given GP raises a follow-on fund and 0 if the current fund is the last fund. Because we do not want to bias the results for GPs that only recently raised a fund (and, therefore, have not had enough time pass to need to raise a next fund), we drop any fund that was started after 1998. We regress this dummy variable on the measures of market performance we have used throughout this paper: S&P 500, Nasdaq Composite Index, and venture industry returns. We include contemporaneous performance at the time the current fund was raised, and market performance 1 year before the current fund started.

Columns 1 and 4 of Table XII report that funds raised when market returns are higher are less likely to raise a follow-on fund. This negative relation is significant for all contemporaneous and lagged performance measures for Nasdaq and S&P. This suggests that funds raised in boom years are more likely to perform poorly and, therefore, are unable to raise a follow-on fund. Column 7 uses the annual return to venture funds as calculated by VE. The results of this regression are more ambiguous. The coefficient on contemporaneous venture performance is positive and of the same magnitude as the negative coefficient on lagged performance.

In columns 2, 5, and 8 of Table XII, we also include a measure of market returns in the third year after the current fund was raised. From anecdotal evidence, we know that GPs tend to raise new funds about every 3 years. Therefore, we include this 3-year leading market indicator to capture market conditions at the time the partnership is most likely to be trying to raise a next fund. The coefficient on the 3-year leading market index is positive and strongly significant for all measures of market returns. Again, this finding is consistent with private equity firms being able to raise capital more easily when overall market conditions are good.

Table XII
Probability of Raising a Follow-on Fund

We estimate a linear probability model of the likelihood that an existing fund raises a follow-on fund. The dependent variable is a dummy equal to 1 if the fund raises a next fund and 0 otherwise. Nasdaq is the annual return on the Nasdaq Composite Index in the year the fund was raised, $Nasdaq_{t-1}$ and $Nasdaq_{t+3}$ are the returns in the 1 year lagged and the 3 year leading Nasdaq Composite index, respectively. The former captures the market conditions in which the current fund was raised, and the latter captures the market conditions in which the next fund will be raised. Similarly, $S\&P_t$ is the annual return on the S&P 500, and VC Returns is the annual return of the venture capital industry as reported by VE. To avoid bias for funds that were only raised very recently, we drop the last 3 years of data, 1999–2001. Size is the amount of capital under management in the current fund. Sequence is the sequence number of the current fund. The variable VC Dummy is equal to 1 if the fund is a venture capital fund and 0 for buyout, LBO, and mezzanine funds. Columns 3, 6, and 9 only include the subsample of first-time funds.

Dependent Variable: Does Fund Raise a Follow-on Fund?									
Nasdaq _t	−0.26 (0.03)	−0.15 (0.04)	−0.16 (0.09)						
Nasdaq _{t−1}	−0.12 (0.04)	−0.06 (0.04)	−0.03 (0.09)						
Nasdaq _{t+3}		0.10 (0.04)	0.12 (0.05)						
S&P _t				−0.30 (0.05)	−0.23 (0.05)	−0.24 (0.11)			
S&P _{t−1}				−0.44 (0.05)	−0.25 (0.06)	−0.31 (0.11)			
S&P _{t+3}					0.15 (0.05)	0.17 (0.10)			
VC Returns _t							0.04 (0.00)	0.01 (0.01)	0.01 (0.01)
VC Returns _{t−1}							−0.04 (0.00)	−0.01 (0.00)	−0.01 (0.01)
VC Returns _{t+3}								0.00 (0.00)	0.02 (0.00)
log(Sequence)	0.09 (0.01)	0.07 (0.01)		0.09 (0.01)	0.07 (0.01)		0.09 (0.01)	0.07 (0.01)	
log(Size)	−0.01 (0.00)	−0.01 (0.00)	−0.01 (0.00)	−0.01 (0.00)	−0.02 (0.00)	−0.00 (0.00)	0.00 (0.01)	0.00 (0.01)	0.04 (0.02)
Adjusted R ²	0.07	0.04	0.03	0.07	0.04	0.01	0.06	0.02	0.03
No. of Observations	2,789	2,467	751	2,831	2,503	756	1,608	1,147	481

Standard errors are in parentheses and are adjusted for serial correlation and heteroskedasticity.

Finally, we repeat the above estimation for follow-on funds by new GPs or first-time funds (as opposed to follow-on funds by an existing GP). We report the results in columns 3, 6, and 9 of Table XII. The same pattern we observe for individual funds holds true for GP starts. GPs that enter the market in boom times are less likely to raise a follow-on fund. However, if the market conditions are positive 3 years after the initial funds were raised, the likelihood of being able to raise a follow-on fund improves significantly.

In sum, it appears that the marginal dollar invested in boom times goes toward financing funds that are less likely to be able to raise a subsequent fund (which we interpret as a proxy for poor performance).

C. Are Fund Returns Affected by the Number of New Entrants?

Finally, we consider the effects of the entry of new GPs and funds on overall industry performance and on the performance of particular types of funds. In Table XIII, we regress the performance of individual funds on the number of new funds entering the industry in the year the fund was started as well as controls for fund size, sequence number, and VC Dummy. For this table, we use the larger sample of private equity funds, and we measure performance using the VE IRR after 5 years. The variable Entry is the logarithm of the aggregate number of new private equity funds in a given year. We also control for the returns on the Nasdaq Composite Index in the year a fund was started, as we know from the results in the previous table that funds are more likely to get started in boom years.¹⁴

We report results for the overall industry in columns 1 and 2 of Table XIII. Column 1 shows the correlation between fund returns and the logarithm of the number of new entrants. The point estimate is negative (-0.14), but statistically insignificant. However, when we include an interaction term between the entry variable and the logarithm of the sequence number of a fund, the coefficient on the interaction term is positive and statistically significant (the point estimate is 0.76 with a standard error of 0.28). At the same time, the coefficient on the direct effect of entry is now negative and statistically significant as well (point estimate of -0.61 and a standard error of 0.20). These results suggest that in periods of increased entry of funds into the industry overall, we observe a larger negative effect on the young funds (those with lower sequence numbers) than on the older, more established funds.

In columns 3 and 4, we repeat the above analysis only for the subsample of VC funds. In this specification, we base the aggregate number of new entrants only on new VC funds that enter the industry in a given year. We find that the effect is stronger in the VC industry. In column 3, we only include the direct effect of the number of new entrants. The coefficient on this variable is negative and barely significant (coefficient of -0.34 with a standard error of 0.18). In column 4, we again include the interaction term between the entry variable and the logarithm of the sequence number of a fund. The coefficient on this term is positive and significant, with a coefficient of 1.13 and a standard error of 0.42 . When we combine the direct effect and the interaction effect, the returns of older funds are relatively unaffected by the inflow of new funds. The returns

¹⁴ We also repeat the regressions in Table XIII when measuring entry as the number of new *partnerships* entering the industry in a given year. The results are qualitatively unchanged. Similarly, we also include annual returns on the S&P 500 or the venture capital index as measures of market performance. Again, the results are qualitatively the same.

Table XIII
Market Entry and Fund Performance

The dependent variable is the realized fund IRR after 5 years of existence. On average this restricts the sample to funds that were started before 1998, and only one observation per fund. Entry is the logarithm of the number of private equity funds overall that are started in the same year as a given fund. This variable is calculated based on the full VE database. In columns 1 and 2, we use the sample of venture capital and buyout funds. In columns 3 and 4, we use only the subsample of venture capital funds. The entry variable in these columns is based only on the number of venture capital funds entering the market in a given year. In columns 5 and 6, we use only the subsample of buyouts funds. Parallel to before, we use only the number of buyout funds entering the market in a given year to calculate the entry variable. Nasdaq_{*t*} is the annual return on the Nasdaq Composite Index. Size is the amount of capital a fund has under management in the current fund. Sequence is the sequence number of the fund. VC Dummy is equal to 1 if the fund is a venture capital fund and 0 for buyout, LBO, and mezzanine funds.

Dependent Variable: IRR _{VE5} (IRR after 5 Years)						
	All Funds		VC Funds		Buyout Funds	
Entry	-0.14 (0.14)	-0.61 (0.20)	-0.34 (0.18)	-0.89 (0.36)	-1.35 (0.25)	-1.54 (0.36)
Entry * log(Sequence)		0.76 (0.28)		1.13 (0.42)		0.33 (0.28)
log(Sequence)	0.44 (0.15)	-3.36 (1.34)	0.57 (0.21)	-4.86 (1.86)	0.12 (0.08)	-1.21 (1.21)
log(Size)	0.09 (0.05)	0.09 (0.05)	0.20 (0.08)	0.15 (0.08)	0.01 (0.04)	0.01 (0.05)
Nasdaq _{<i>t</i>}	0.28 (0.11)	0.29 (0.11)	0.79 (0.31)	0.50 (0.16)	0.17 (0.11)	0.17 (0.11)
VC Dummy	0.66 (0.19)	0.63 (0.19)				
Adjusted <i>R</i> ²	0.05	0.07	0.06	0.10	0.22	0.23
No. of Observations	1,090	1,090	754	754	336	336

Standard errors are in parentheses and are adjusted for serial correlation and heteroskedasticity.

of the entering funds, however, are significantly lower in these periods. The overall effect on the industry returns, therefore, is negative in periods with a large number of new entrants.

Finally, in columns 5 and 6, we repeat these regressions for the subsample of buyout funds. We base the aggregate number of new entrants only on new buyout funds that enter the industry in a given year. In column 5, we find that the direct effect of the number of new entrants is negative and statistically significant for the buyout industry. The coefficient on the entry variable is -1.35 with a standard error of 0.25. Overall returns in the buyout industry are significantly diluted in periods where many new funds enter the market. In column 6, we again include the interaction term between the entry variable and the logarithm of the sequence number of a fund. The coefficient on this term is positive, but smaller than for the sample of VC funds and it is not statistically significant. This result suggests that unlike in the VC industry, the returns of older funds in the buyout industry (those with higher sequence numbers) are

less isolated from the entry of new funds.¹⁵ However, since our sample of buyout funds is relatively small and less comprehensive than the sample of venture capital funds, we do not want to overstate the inference we can draw from the differences between VC and buyout funds.

VII. Summary and Implications

In this paper, we investigate the performance of private equity partnerships using a data set of individual fund returns and cash flows collected by VE. Over the entire sample period (1980–1997), average fund returns net of fees are roughly equal to those of the S&P 500. Weighted by committed capital, venture funds outperform the S&P 500 while buyout funds do not. Our estimates also suggest that gross of fees, both types of private equity partnerships earn returns exceeding the S&P 500. We acknowledge, however, that the average return results are potentially biased because we do not control for differences in market risk and because of possible sample selection biases.

We find that returns persist strongly across funds raised by individual private equity partnerships. We observe persistence for all measures of fund performance and at both ends of the performance distribution. We also observe persistence in a different sample of funds with publicly available data. We present a number of analyses that indicate that these results are unlikely to be induced by differences in risk or sample selection biases. We also perform a number of robustness checks that suggest that industry or investment stage differences, and positive performance biases do not drive the results.

We also document that performance increases (in the cross-section) with fund size and with the GP's experience. The relation with fund size is concave, suggesting decreasing returns to scale. Similarly, a GP's track record is positively related to the GP's ability to attract capital into new funds. In contrast to the convex relationship in the mutual fund industry, this relationship is concave in private equity. Finally, the fact that past performance (measured both as PME or IRR) is strongly related to a GP's ability to raise future funds and the amount of those funds is consistent with our hypothesis that the persistence of PMEs and IRRs measures persistence in performance rather than differences in risk.

Finally, we find some evidence that funds and partnerships that are raised in boom times are less likely to raise follow-on funds, suggesting that these funds perform worse. In conjunction with our results on average returns, this suggests a boom and bust cycle in which positive market-adjusted returns encourage entry, which leads to negative market-adjusted returns, and so forth.

These empirical relations between performance and capital flows differ substantially from the ones found for mutual funds. We think the most likely explanation for these results is a model of underlying heterogeneity in the skills

¹⁵ One could speculate about the reasons why we observe such a difference between the returns in the buyout versus the venture capital industry. There are a number of practitioners as well as academics who suggest that buyout returns are more sensitive to market timing and herding effects; see, for example, Kaplan and Stein (1993). In contrast, it is often stated that VC returns are crucially driven by the specific human capital or networks of a fund's GPs.

of GPs, and concavity in the production function. Successful GPs might choose not to grow their funds until the excess returns have been diluted if there are strong diseconomies from scale and scope. This could be true for a number of reasons. First, successful GPs might not easily scale up investments by putting more money in any particular deal or investing in more companies because they provide other inputs that are difficult to scale, such as time and advice. Similarly, it may be difficult to hire partners of the same quality as the existing partners. Second, one could imagine that top GPs choose to raise less capital than they could because the number of good startups in the economy is limited at each point in time. However, passing up less profitable (but potentially still positive NPV projects) could only be an optimal choice for the GP if there are negative spillover effects on the *inframarginal* deals from engaging in these investments.

If, indeed, the persistence results are driven by heterogeneity in GP skill and limited scalability of human capital, it remains puzzling that these returns to superior skills are not appropriated by the scarce input (i.e., the GP) in the form of higher fees. Moreover, we need to understand why during boom times capital flows disproportionately to funds with lower performance (which subsequently have very low returns) instead of flowing to the best GPs.

One could conjecture that better performing GPs have better governance structures or LPs who provide more value added (see, for example, Lerner and Schoar (2004) for a possible mechanism). However, we do not have a way to test these theories in the current paper. Our findings highlight the need for future work that aims to better understand the organizational structure of the private equity industry.

REFERENCES

- Bares, Pierre-Antoine, Rajna Gibson, and Sebastien Gyger, 2002, Performance in the hedge fund industry: An analysis of short-term and long-term persistence, Working paper, Swiss Banking Institute, University of Zurich.
- Berk, Jonathan, and Richard Green, 2002, Mutual fund flows and performance in rational markets, Working paper no. 9275, National Bureau of Economic Research.
- Brown, Stephen J., William N. Goetzman, and Roger J. Ibbotson, 1999, Offshore hedge funds: Survival and performance, *Journal of Business* 72, 91–119.
- Carhart, Mark M., Jennifer Carpenter, Andrew Lynch, and David Musto, 2002, Mutual fund survivorship, *Review of Financial Studies* 15, 1439–1463.
- Chen, Joseph, Harrison Hong, Ming Huang, and Jeffrey Kubik, 2004, Does fund size erode mutual fund performance? The role of liquidity and organization, *American Economic Review* 94, 1276–1302.
- Chevalier, Judith, and Glenn Ellison, 1997, Risk taking by mutual funds as a response to incentives, *Journal of Political Economy* 105, 1167–1200.
- Cochrane, John, 2003, The risk and return of venture capital, Working paper, University of Chicago.
- Edwards, Franklin, and Mustafa Cagalyan, 2001, Hedge fund performance and manager skill, Working paper, Columbia University.
- Gompers, Paul, and Joshua Lerner, 1998, What drives venture capital fundraising? *Brookings Papers on Economic Activity and Microeconomics* 1998, 149–192.
- Gompers, Paul, and Joshua Lerner, 1999, An analysis of compensation in the U.S. venture capital partnership, *Journal of Financial Economics* 51, 3–44.

- Gompers, Paul, and Joshua Lerner, 2000, Money chasing deals? The impact of fund inflows on private equity valuation, *Journal of Financial Economics* 55, 281–325.
- Hellman, Thomas, and Manju Puri, 2002, Venture capital and the professionalization of start-up firms: Empirical evidence, *Journal of Finance* 57, 169–197.
- Hsu, David, 2004, What do entrepreneurs pay for venture capital affiliation? *Journal of Finance* 59, 1805–1844.
- Jones, Charles, and Matthew Rhodes-Kropf, 2003, The price of diversifiable risk in venture capital and private equity, Working paper, Columbia University.
- Kaplan, Steven, and Jeremy Stein, 1993, The evolution of buyout pricing and financial structure in the 1980s, *Quarterly Journal of Economics* 27, 215–246.
- Kaplan, Steven, and P. Strömberg, 2003, Financial contracting theory meets the real world: Evidence from venture capital contracts, *Review of Economic Studies* 70, 281–315.
- Kaplan, Steven, and Per Strömberg, 2004, Characteristics, contracts, and actions: Evidence from venture capitalist analyses, *Journal of Finance* 59, 2177–2210.
- Kat, Harry, and Faye Menexe, 2002, Persistence in hedge fund performance: The true value of a track record, Working paper, University of Reading.
- Lerner, Joshua, and Antoinette Schoar, 2004, The illiquidity puzzle: Theory and evidence from private equity, *Journal of Financial Economics* 72, 3–40.
- Ljungqvist, Alexander, and Matt Richardson, 2002, The cash flow, return and risk characteristics of private equity, Working paper, New York University.
- Sirri, Erik, and Peter Tufano, 1998, Costly search and mutual fund flows, *The Journal of Finance* 53, 1589–1622.