



NYU  STERN

Salomon Center for the Study of Financial Institutions

Working Paper Series **CORPORATE GOVERNANCE** Research Group

THE CASH FLOW, RETURN AND RISK CHARACTERISTICS
OF PRIVATE EQUITY

Alexander Ljungqvist
Matthew Richardson

S-CG-03-01

The cash flow, return and risk characteristics of private equity ^{*} [†]

Alexander Ljungqvist
Stern School of Business
New York University
and *CEPR*

Matthew Richardson
Stern School of Business
New York University
and *NBER*

First draft: November 14, 2002

This draft: January 9, 2003

* We are grateful to an anonymous institutional investor for making the data used in this study available, and to Eric Green for many helpful discussions and suggestions. We also thank the Salomon Center at NYU Stern for generous financial assistance, and Steve Kaplan and Jeff Wurgler for many helpful comments. We are grateful to Eric Stern for excellent research assistance. All errors are our own.

† Address for correspondence: Stern School of Business, New York University, Suite 9-190, 44 West Fourth Street, New York NY 10012-1126. Fax 212-995-4233. e-mail aljungqv@stern.nyu.edu (A. Ljungqvist), mrichar0@stern.nyu.edu (M. Richardson).

The cash flow, return and risk characteristics of private equity

Abstract

Using a unique dataset of private equity funds over the last two decades, this paper analyzes the cash flow, return, and risk characteristics of private equity. Unlike previous studies, we have detailed cash flow data for each fund, rather than aggregate or accounting returns. We also know the exact timing of investments and capital returns to investors and the number and types of companies each fund invested in. We document the draw down and capital return schedules for the typical private equity fund, and show that it takes several years for capital to be invested, and over ten years for capital to be returned to generate excess returns. We provide several determining factors for these schedules, including existing investment opportunities and competition amongst private equity funds. In terms of performance, we document that private equity generates excess returns on the order of five to eight percent per annum relative to the aggregate public equity market. Moreover, while we estimate the betas of the private equity funds' portfolios to be greater than one, we show that on a risk-adjusted basis the excess value of the typical private equity fund is on the order of 24 percent relative to the present value of the invested capital. One interpretation of this magnitude is that it represents compensation for holding a 10-year illiquid investment.

1. Introduction

This paper provides the first analysis of private equity returns based on actual cash flows of venture capital and buyout funds. Despite the important role of private equity in financing and fostering innovative firms, and in reallocating capital to more productive sectors of the economy, relatively little is known about the key characteristics of private equity as an asset class: liquidity, risk, and return. Relative to other asset classes, private equity investments are illiquid, in the sense that there is no active secondary market for such investments, investors have little control over how the capital is invested, and the investment profile covers a long horizon. Our data allow us to document the degree of liquidity and the resulting compensation, if any, in terms of risk-adjusted returns provided to investors. Specifically, we empirically measure the timing and magnitude of investment decisions throughout the ten-year life of a fund, how quickly capital is returned to investors, and the overall performance of private equity as a function of various characteristics.

Our results complement an emerging literature in finance that has looked at the returns on private equity. (See, for example, Cochrane (2000), Moskowitz and Vissing-Jørgensen (2002), Kaplan and Schoar (2002), and Quigley and Woodward (2002), among others.) There is an important caveat to this literature: prior work has not had access to very precise data about private equity funds. In particular, rather than computing performance using the distribution of cash flows to and from fund investors, prior work has generally relied on fund valuation data as collected by Thomson Financial's Venture Economics service. These data have three principal shortcomings: they are available only in aggregate rather than fund-by-fund form (but see Kaplan and Schoar for an important exception); the data are self-reported and thus potentially subject to selection biases; and they are based on unrealized as well as realized investments which introduces noise and potentially biases due to subjective accounting treatment.

In contrast, we use cash flow data for a fairly large dataset of private equity funds raised over the period 1981 to 2001 to shed brighter light on the characteristics and performance of this asset class. Our dataset includes, among other items, precisely dated cash flows representing investments in portfolio companies, management fees, and distributions of capital gains to investors. We also know the portfolio compositions in terms of the number and types of companies each fund invested in. Using these data, we analyze the cash flow, return, and risk characteristics of private equity. Specifically, we ask: (i) What are the capital investment and return patterns of private equity throughout the life of the fund? (ii) What determines the speed with which funds invest their capital over time? (iii) How long does it take for returns to turn positive? (iv) What is the risk profile of

private equity funds, both in terms of systematic and unsystematic risk? And (v) are private equity returns impressive relative to their risk profile and various benchmarks? The contribution of our paper is to be able to provide the first definitive answers to these questions.

Beyond providing the first evidence of private equity returns based on actual cash flows, our paper generates some new stylized facts about private equity investing. First, we document the dynamics of both draw downs (capital investment) and capital returns over a fund's life. These suggest that knowledge of the timing of actual cash outflows and inflows is an important factor in understanding the performance of private equity funds. For example, it takes six years for 90 percent of the committed capital to be invested, so simplifying assumptions about immediate up front deployment of capital seem problematic. Moreover, we perform a cross-sectional and time-series analysis of the determining factors of how fast or how slowly a fund invests. We empirically identify two key factors, namely time variation in the availability of investment opportunities and competition for deal flow with other private equity funds. In particular, we find that it is easy to invest an existing fund when it is a good time to raise a new fund. At the same time, funds take longer to invest when their peers have more money with which to chase the same deals.

Second, the IRR of the average fund does not turn positive until the eighth year of the fund's life. Thus, once the adjustment for the cost of capital is made, it is only at the very end of a life's fund that excess returns are realized. This highlights the illiquidity of private equity investments. It also suggests that "interim" IRRs computed before a fund reaches maturity are not very informative. Measuring fund performance thus requires using precisely dated cash flows over a fund's life, rather than relying on arbitrary assumptions about the time profile of capital returns.

Third, we find that private equity has generated substantial excess returns over the past two decades. Specifically, for funds started between 1981 and 1993 (the "mature" funds in our dataset), we document internal rates of return averaging 19.81 percent, net of all fees.¹ In contrast, investment in the public stock market measured using the S&P 500 index under an identical time schedule of cash outflows yields 14.1 percent. Excess returns of close to six percent per annum no doubt compensate investors, at least in part, for the extreme illiquidity of private equity.² But they

¹ Funds raised more recently are still in operation – either actively investing or seeking to exit portfolio companies via IPOs or trade sales – and their eventual realized returns will include losses due to the internet bust and the recent recession. It is conceivable that their loss rates will be higher than has been the case historically. However, the funds in our dataset are predominantly from the buyout sector and so are less subject to the idiosyncrasies of the internet sector. We address the behavior of recently raised funds in Section 6.

² Alternatively, one could argue that the marginal investor places a zero price on illiquidity. That is, the clientele for private equity may be such that investors are liquid based on their other portfolio holdings.

may also be a reflection of the investment skill of certain fund managers and the “closed club” nature of private equity investing. These results are somewhat in contrast to the aforementioned literature, and we evaluate the extent to which this is due to our particular sample characteristics versus the higher quality of our data.

Fourth, by looking at each fund’s investments in detail, assigning industry betas to the portfolio companies, we are able to estimate fund risk. This is important because it is not possible to estimate risk at the fund level, using standard time-series correlations with the market return, as the fund’s investments are realized fully only after the fund has been liquidated, usually after ten years. We find that fund returns are still abnormally large even on a risk-adjusted basis. For example, while the return on invested capital (as measured by the Profitability Index) averages 25.07 percent assuming simplistically a beta of one with the market, it still averages 24 percent when discounting cash flows at the risk-adjusted cost of capital.

Finally, we document several key characteristics of the typical private equity fund’s portfolio of investments, including the number of companies held, the industry concentration, and portfolio beta. Our private equity funds are not well diversified: on average, they invest close to 40 percent of their capital in a single industry. Using these risk characteristics and some other standard ones from the literature, we perform a cross-sectional analysis of the determinants of private equity returns. One conclusion drawn from this analysis is that the cross-section of fund returns is unlikely to be explained by either the underlying systematic or unsystematic risk of the portfolio companies. The more important characteristics appear to be the overall amount of money flowing into the private equity sector as a whole (“money chasing deals”) and the size of the fund itself.

The paper is organized as follows. Section 2 provides a brief description of the existing literature on the risk and return of private equity investments. In Section 3, we describe our dataset. Section 4 provides the core results of the paper by documenting both the cash flow patterns of private equity funds and the returns on these funds over the last 20 years. Because a number of the funds in our dataset are still on-going concerns, special emphasis is placed on the return characteristics of the “mature” funds. Section 5 compares the performance of these mature funds to investments in public equities holding constant the time profile of investment, and adjusted for the risk of the underlying portfolio companies. In Section 6, we conjecture about private equity funds raised in the 1998-2001 period. Section 7 concludes.

2. Existing Literature

Given the volume of literature on venture capital, it may seem surprising that there are only a few papers analyzing the returns on private equity.³ The main obstacle to research has been the limited availability of data.

The main sources of data on the private equity industry are two commercial data vendors, Venture Economics and VentureOne. These are in general excellent sources of information about the investment *behavior* of private equity funds, such as which fund invests how much in which company. They are not, however, ideal for investigating the performance of private equity funds.

While Venture Economics publishes internal rates of return, it does so only at an aggregate level, such as for the median or third-quartile fund. Fund level data are not publicly available. Moreover, the Venture Economics returns represent a mixture of growth in net asset value (NAV) for unrealized investments and IRRs for realized investments.⁴ The inclusion of NAV growth rates in the Venture Economics data is potentially particularly misleading. Writing up a portfolio investment long before any actual cash distribution to investors flatters the time profile of returns and thus increases the accounting IRR. Of course, accounting valuations are not bankable and may never be followed by commensurate cash distributions. Nor are valuations subject to any type of generally accepted accounting principles. Thus, for the exact same investment, different private equity funds may assess very different values (see Blaydon and Horvath (2002)). Differences in accounting practices can occur for a number of reasons. For example, many private equity funds are conservative in their assessments and value investments at cost until the investments are realized. Other funds – particularly first-time funds – may be aggressive in their valuations by not writing down poorly performing companies or even overstating the value of ongoing ones, especially in difficult times (see Gompers (1996) and Blaydon, Horvath, and Wainwright (2002)). These differences in assessed values induce little confidence in the reported values and IRRs of private equity funds (Gompers and Lerner (1997)).

Four recent papers have attempted to quantify the returns and risk of private equity. (See Cochrane (2000), Moskowitz and Vissing-Jørgensen (2002), Kaplan and Schoar (2002), and

³ There is an extensive literature on the role of venture capital in financial markets, including the relation between venture capitalists, entrepreneurs and public markets. See, for example, Gompers and Lerner (1999) for an extensive discussion.

⁴ As funds near their liquidation date, the weight of NAV growth rates in the calculation of performance measures declines. However, unless IRRs are calculated using precisely dated cash flows, reported IRRs may still overstate or understate performance. We are told that many funds make simplifying assumptions about the timing of cash flows when reporting IRRs.

Quigley and Woodward (2002)). The general conclusion from these papers regarding private equity performance is mixed and these differences can be partly attributed to the quality of the data.

Kaplan and Schoar (2002) have exclusive access to the (albeit anonymized) fund-by-fund data from which Venture Economics derive their published aggregate private equity performance data. This is a valuable improvement on the use of aggregated data, though the fund-by-fund return data remain subject to the limitations of self-reporting and accounting treatment noted earlier. These limitations can create substantial variation at the individual fund level which Kaplan and Schoar, to their credit, fully acknowledge. To this point, Kaplan and Schoar document large heterogeneity in performance across funds. These excess returns, however, do not take into account the timing of the cash flows (which is not available) or the risk profile of the portfolio companies (due to anonymity). This caveat is potentially important as we document later that the draw down and capital return schedules and portfolio risks take on complex patterns which, in turn, affect performance estimates. Most importantly, realizations of investment returns do not take place until very late in the life of the funds. Kaplan and Schoar's analysis does, however, have the significant advantage of including a large cross-section, and they present evidence suggesting performance is affected by a number of important characteristics not looked at here, such as the ability to raise follow-on funds.

As an alternative strategy, Cochrane (2000) and Quigley and Woodward (2002) focus on the individual portfolio company (rather than fund) level, and then, using various assumptions, infer the aggregate performance of private equity investing. These papers are important and document interesting facts about private equity investments. Cochrane stresses the importance of adjusting for survivorship bias, which potentially arises due to the high failure rate of private equity investments. Unfortunately, the Venture One data used there do not permit controlling for the timing of the cash flows to and from the portfolio companies, nor the actual dollar realization of the investments when taken public or sold – limitations the authors acknowledge. Assigning valuations to about three quarters of IPOs and one quarter of acquired companies using Thomson Financial SDC's new issues and M&A databases, Cochrane finds that mean log returns of individual portfolio investments are around five percent, with slightly negative alphas, though arithmetic returns are much higher. Using different assumptions, Quigley and Woodward (2002) report lower returns for private than for public equity.

Moskowitz and Vissing-Jørgensen (2002) similarly report that returns to private equity mostly fall short of those in the public equity market. Their analysis focuses on a broadly-defined notion of private equity, with data derived from the Survey of Consumer Finances and various national

income accounts. Thus, their results pertain mostly to non-intermediated (entrepreneurial) investments in non-public companies, as opposed to intermediated investments undertaken by private equity funds.

Our paper avoids these problems by using data on both the exact timing of the investments and distribution of cash flows to investors, and the types of companies contained in each fund's portfolio. The drawback of our approach is the relatively limited sample size. However, the only previous paper to have access to similar data, Gompers and Lerner (1997), looked at just one fund.⁵ In the next section, we describe the dataset and its characteristics in more detail.

3. Sample and Data

3.1 Overview of Dataset

Our dataset is derived from the records of one of the largest institutional investors in private equity in the U.S. We will refer to this investor as the "Limited Partner". As a condition for obtaining the data, we have agreed to identify neither the Limited Partner nor the names of the funds or portfolio companies in the dataset. Of particular interest, these data are not subject to survivorship bias as all investments the Limited Partner has made since 1981 are included.

Between 1981 and 1993, the Limited Partner invested in 73 funds. These funds – which we will refer to as the "mature" funds since they are around ten or more years old – form the basis for our performance analysis. The Limited Partner has continued to invest in a large number of private equity funds since 1993, but to protect the Limited Partner's identity, we have agreed not to disclose certain characteristics of the funds raised after 1993, such as their number and size. We do, however, include the underlying cash flow data for such funds in our analysis where appropriate.

Private equity firms (often called "general partners" or GPs), when successful, usually raise follow-on funds. In our dataset, 28.9 percent of funds are first-time funds, 20.6 percent are second funds, 11.6 percent are third funds, and the remaining 39.0 percent are later funds.

The extent to which the funds in our dataset are representative of the universe of private equity funds depends on the Limited Partner's investment strategy. For instance, it would be problematic if the Limited Partner only invested in follow-on funds raised by venture capitalists with proven track records. This is not the case. In part, this is because the Limited Partner's investment objectives are twofold: not only to obtain the highest risk-adjusted return, but also to increase the likelihood that

⁵ Interestingly, for the particular fund in their sample, Gompers and Lerner (1997) also report excess returns relative to public equity.

the funds will “purchase” the services our Limited Partner’s corporate parent has to offer. These services are arguably more attractive to first-time funds that have yet to build up relationships, thus the relatively high rate of first-time funds in our dataset. This issue of sample selection is further studied in Section 4.5.

We are able to match all but eight of the sample funds to funds covered in Venture Economics. This allows us to augment the information we received from the Limited Partner with publicly available information. 91.1 percent of the funds are based in the U.S., 7.4 percent in Europe, and 1.5 percent in Latin America. Within the U.S., funds are geographically concentrated: 58 percent of funds are headquartered in California or New York, 10 percent in Massachusetts, 6.7 percent in Texas, and the remainder in 11 other states.

Table 1 shows a breakdown of the sample and of several descriptive statistics by the year in which the funds were raised (the “vintage years”). For comparison, the table also shows the total number of new funds raised in each year, according to Venture Economics, by partnerships headquartered in the countries that our sample funds are located in. Every year between 1992 and 2000 saw more funds raised than the previous year, with particularly large numbers raised at the height of the internet frenzy in 1999 and 2000. Although we cannot show this in detail, the number of funds the Limited Partner invested in similarly increased throughout the 1990s, peaking in 1999-2000.

Our dataset contains both venture capital and private equity funds.⁶ For the entire period from 1981 to 2001, a quarter of funds, representing 14.8 percent of fund capital, are venture funds. This differs from the universe of funds tracked by Venture Economics, where venture funds account for 74.6 percent of funds by number and 41.5 percent by capital. Our Limited Partner thus invests disproportionately in private equity (primarily buyout) funds. In the 1981-1993 sample of “mature” funds, this pattern is even more pronounced: private equity funds account for 74 percent of funds by number and 88.2 percent by capital.

In the private equity industry, fund size is usually expressed as the sum of investors’ “capital commitments.” The capital commitment is the maximum amount of money an investor can be asked to contribute over the life of the fund. Note that when agreeing to invest in a fund, investors do not transfer the committed cash up front. Instead, general partners “draw down” cash against the

⁶ Private equity funds are mainly those flagged as “Buyout” (90.4 percent) by Venture Economics (or the Limited Partner, if not covered by Venture Economics). The remainder are flagged as “Generalist Private Equity” (3.8 percent), “Mezzanine” (4.8 percent), and “Other Private Equity” (1 percent).

commitment whenever they wish to make an investment. The rate at which cash is invested clearly affects the IRR the fund achieves.

Sample funds started between 1981 and 2001 had aggregate commitments of \$207 billion (in nominal terms), as compared to \$1.184 trillion in the Venture Economics universe over the same period (not shown). Thus, while we only have data for a relative small number of the 8,539 funds raised over the period, our funds account for a disproportionately large fraction – 17.5 percent – of capital commitments. Our coverage is even better among non-venture funds: we have data for 9.5 percent of all non-venture funds raised over the period, accounting for 29.3 percent of committed capital in those funds.

The 73 mature funds started between 1981 and 1993 had aggregate commitments of \$36.7 billion, with the average fund raising \$502.8 million. Our Limited Partner's investment in these funds is sizeable. It committed \$1.1 billion in aggregate, with an average of \$15.2 million per fund and a range between \$800,000 and \$167.4 million. As a fraction of total fund size, the Limited Partner committed 4.64 percent of the average mature fund's capital.

3.2 Cash Flows and Stock Distributions

The Limited Partner made available to us the complete cash flow records for all its private equity investments up to May 31, 2001. We subsequently obtained additional data up to September 30, 2002 for 21 funds that were close to maturity, thus increasing the number of funds that have been liquidated or are close to liquidation.

There are essentially three types of cash flows: cash flows associated with “disbursements” (investments in portfolio companies) and “exits” (receipt of cash inflows from IPOs or trade sales) as well as (occasionally) dividends paid by portfolio companies; annual management fees (typically 1-2 percent of committed capital); and (occasional) interest payments on cash held by the GP prior to making an investment. The data do not separately record the GPs' share in a fund's capital gains (usually 20 percent), called the “carried interest” or “carry”, as GPs transmit capital gains to investors net of their carry.

The cash flows involve four types of investment scenarios. 1) Cancelled transactions: a cash call followed shortly after by the return of the cash, along with bank interest. 2) Write-offs: cash outflow(s) without subsequent cash inflow, or with a subsequent accounting (non-cash) entry flagging a “capital loss”. 3) Cash distributions following successful exits (in the form of an IPO or a trade sale): cash outflow(s) followed by cash inflow(s). And 4), stock distributions following

successful exits: cash outflow(s) followed by a non-cash entry reflecting receipt of common stock. The stock would be the portfolio company's in the case of an IPO, and the buyer's in the case of a sale to a publicly traded firm. Following a stock distribution, one of two things can happen: the Limited Partner sells the stock, or it holds it in inventory. Sales show up as cash inflows. Positions that are held in inventory are marked to market periodically (usually monthly), but they are obviously not cash.

Given our focus on cash IRRs, stock distributions warrant special attention. Fortunately, they are relatively rare. Fewer than 20 percent of sample funds engage in stock distributions, so most funds distribute cash. However, this statement suffers from right-censoring, as many younger funds haven't yet started distributing anything. Among the 73 funds raised between 1981 and 1993, 31 have distributed stock.

There are 203 stock distributions in our dataset, involving 171 portfolio companies. (We estimate that there are more than 4,300 portfolio companies in our dataset of which more than 900 have been exited.) Upon receipt of distributed stock, our Limited Partner almost always sells distributed stock: there are only seven (mostly recent) cases where it has yet to sell or sell completely. Thus, we do have an objective cash valuation for virtually all stock distributions. However, the Limited Partner does not typically sell immediately. In only 65 cases does it sell within 60 days, and the average (median) time between receipt of the stock distribution and the last sale associated with that distribution is 245 (110) days.⁷

For performance evaluation purposes, we "book" cash inflows from the sale of distributed stock on the Limited Partner's actual transaction days. To the extent that our Limited Partner follows a different sale policy and time profile than other fund investors, its realized return from investing in a fund may therefore not be entirely representative. Alternatively, one might simulate the performance of a policy of selling distributed stock automatically upon receipt. However, we do not know the number of shares the Limited Partner received in the stock distributions, so we cannot estimate their market value on the distribution dates.

Potentially problematic are the cases where the Limited Partner still has unsold stock, since our focus on cash flows implies that we value stock held in inventory at zero. Luckily, there is only one such case among the "mature" funds that we focus on in our performance analysis. In this case, we

⁷ This holding period is a maximum, as positions are often unwound in multiple transactions.

obtain the Limited Partner's marked-to-market valuation and pretend it sold the position at this price.

4. The Cash Flow and Return Characteristics of Private Equity Funds

4.1 Cash Flow Patterns: Draw-downs

In evaluating the returns on private equity funds, it is common to take the capital commitment as being invested on an immediate basis.⁸ To the extent this is not the case, excess returns will be misstated for two reasons. First, the timing of the cash outflows clearly affects the return via the time value of money. The bias will necessarily be towards an understatement of the return. Second, the risk of the investment should reflect the period over which the capital is invested.

In fact, fund managers typically only draw down the limited partners' capital commitment when they are ready to invest in a portfolio company. Table 2 shows how much of the commitment was drawn down by the earlier of the end of our sampling period or a fund's liquidation date. The average fund in our sample had drawn down 67.3 percent of committed capital. However, this understates draw downs as the more recent funds are not yet fully invested. The 59 (73) funds raised between 1981 and 1992 (1993) invested on average 94.8 (94.7) percent of committed capital. Average draw downs are around 90 percent of committed capital for funds raised up to 1996, with later vintages still actively investing and so still in what is called the "commitment period".

It is arguable when a fund is fully invested. Among the funds raised between 1981 and 1993 that have subsequently been liquidated, some never invested more than 60-70 percent of committed capital. In the overall 1981-2001 dataset, 55.6 percent of funds have invested at least 70 percent of committed capital, and 49.5 percent have invested 80 percent or more as of the end of our sampling period. These might reasonably be thought of as fully, or close to fully, invested. They include a few very recent funds that invested their committed capital very rapidly: 40 percent of the 1998 vintage funds and 10 percent of the 1999 vintage funds had already invested at least 70 percent of committed capital by May 2001.

While the magnitude of the cash outflows (i.e., the investment draw downs and the annual management fees) is clearly a key component for measuring a fund's return, the timing of these outflows is also important for the reasons stated earlier. Table 3 sheds light on the time profile. Funds are typically (but not always) ten-year limited partnerships, with possible extensions by a few

⁸ For instance, Venture Economics' "TVPI" measure, used in Kaplan and Schoar (2002), is defined as total cash returned over total cash invested, without discounting.

years subject to investor approval. The table shows average annual and cumulative draw downs for each year of a fund's life (counted from 1 to 10). The average fund draws down 16.28, 20.35, and 20.15 percent of committed capital in its first three years of operation, so it is 56.8 percent invested by the end of year 3. The draw down rate then slows down. In fact, it takes another three years to hit a 90 percent rate. By year 10, the end of its expected life, the average fund is 93.6 percent invested. While some funds remain in operation beyond year 10, there are no further draw downs.

Though not shown in the table, there is wide variation in the speed with which funds draw down committed capital. For instance, some funds draw it down immediately, while others take as long as ten years to invest 80 percent or more of their commitments. Adjusting for the fact that many of the more recent funds are right-censored, in that they drop out of our sample before they are fully invested, the average (median) fund takes 3.69 (4) years to invest 80 percent or more of their commitments.

4.2 The Determinants of Draw-downs

To shed light on the determinants of how quickly a fund invests its capital, we model the time-to-fully-invested as $\ln(t_i) = \beta\mathbf{X} + \ln(\varepsilon_i)$, where the error ε_i is assumed to follow the exponential distribution with mean β_0 . This is a standard accelerated-failure model (which can easily be rewritten as a proportional-hazard duration model), so the likelihood function has no problem correcting for the right-censoring inherent in the data (Kalbfleisch and Prentice (1980)). Therefore, we can estimate the model using all sample funds raised between 1981 and 2001. We conjecture that time-to-fully-invested varies with the (time-varying) availability of investment opportunities, competition for such investment opportunities, and the cost of funds. We also allow for potential differences between venture and buyout funds, first-time and follow-on funds, and older and more recent funds. Finally, we estimate whether larger funds take longer to invest.

To proxy for the unobserved availability of investment opportunities faced by a buyout (venture) fund in our sample, we include the annual inflows into new buyout (venture) funds (measured as the log of the real dollars raised, in March 1996 purchasing power). This assumes that new funds raise more capital, the more profitable investment opportunities are available. Note that this variable is time-varying: as inflows into new funds change over the life of a sample fund, the fund's managers can respond by accelerating or decelerating the rate at which they invest. In addition, we include a dummy that equals one during the dotcom bubble (1999-2000), on the

assumption that investment opportunities were more abundant in those years. Again, this is a time-varying covariate: over the fund's life, it equals one only in 1999-2000.

To proxy for the degree of competition faced by a buyout (venture) fund in our sample, we include the (log of the real) amount of capital committed to buyout (venture) funds *in the year the sample fund was raised*. That is, a 1990 vintage fund is assumed to be competing with other funds of that vintage. This is clearly a noisy measure of competition. Note that this variable is not time-varying. Finally, we use two measures of the cost of funds: the yield on corporate bonds, using Moody's BAA bond index estimated annually in December, and the annual return on the S&P 500 Index. Both are time-varying over the life of a sample fund.

Table 4 reports the maximum-likelihood estimation results for three different cut-offs of "fully-invested" (more than 70, 80, or 90 percent of committed capital).⁹ The results are qualitatively similar in each case. While venture funds take longer to invest than buyout funds, the difference is only marginally significant. First-time funds invest somewhat faster, but this is not significant at conventional levels. Funds raised between 1981 and 1993, on the other hand, invest significantly faster than newer funds. Note that this finding is not driven by the fact that many newer funds drop out of our sample before becoming fully invested, as we have corrected for right-censoring. We find no evidence that fund size affects the investment rate.

Our proxy for the availability of investment opportunities – the time-varying log of real fund inflows – has a strongly negative and significant effect on the time-to-fully-invested. This makes economic sense: times when it is easy to invest an existing fund are also good times to raise a new fund. The dummy for the bubble years 1999-2000 tells a similar story: funds are invested significantly faster in those two years. Our proxy for competition – the total capital raised in the fund's vintage year – has a positive and significant effect. This suggests that funds take longer to invest when their peers have more money with which to chase the same deals. Finally, the coefficient estimated for the corporate bond yield is positive and significant, indicating that funds invest more slowly as debt becomes more expensive. This is likely driven by the leverage needs of our buyout funds. The return on the S&P 500, on the other hand, has no significant effect. In other words, what happens in the public markets does not appear to affect the speed with which private equity funds deploy their committed capital.

⁹ As mentioned in the previous sub-section, a small number of the mature funds never invested more than 60-70 percent of their capital. For these, we measure time-to-fully invested as the number of years until they reached their maximum draw down.

Overall, all three specifications have good fit, as indicated by the significance of the likelihood ratio tests. We obtain the highest pseudo- R^2 , at 27.7 percent, in the specification that defines fully-invested as 80 percent or more of committed capital.

4.3 Cash Flow Patterns: Distributions

Draw downs represent just one aspect of a fund's cash flows. The other is the return of capital, and capital gains, to investors. Following liquidity events (such as an IPO), capital is returned to investors in the form of cash distributions or stock distributions. (Private equity funds typically have covenants restricting reinvestment of capital gains; see Gompers and Lerner (1996).) Distributions are net of the general partners' "carried interest", that is, the share in any capital gains (usually 20 percent) that accrues to GPs and that constitutes the bulk of their compensation. We again comment on both the magnitude and timing of these cash flows.

Table 2 shows how much of the invested and committed capital was returned to investors by the earlier of the end of our sampling period or a fund's liquidation date. The average fund distributed 106.8 percent of drawn-down capital and 94.3 percent of committed capital. Again, this understates cash flows as recent funds have yet to exit many of their portfolio holdings. The 59 (73) funds raised between 1981 and 1992 (1993) returned 2.75 (2.59) times invested capital and 2.61 (2.44) times committed capital, on average.

As in the case of draw downs, it is also important to consider the timing of these distributions. Table 3 documents the rate at which capital returns and capital gains are distributed to investors. Several observations are in order. First, as one might expect, distributions are rare in the early fund years. For example, by the end of year 3, only 16.6 percent and 12.9 percent of total invested capital and total committed capital have been distributed, respectively. Second, it takes a little under seven years for total invested capital to be returned in the average fund, and around seven years for committed capital to be returned. Much of the "capital gain" is thus generated from year 7 onwards. By year 10, the average fund has distributed 2.07 times its invested capital and 1.93 times its committed capital. Third, 52 funds experience further capital distributions beyond year 10, either because they remain in operation or due to a prolonged liquidation phase. By the time they are eventually liquidated, the 60 funds with 10 or more years of data have returned 2.78 times their invested capital and 2.62 times their committed capital, on average. It is important to note, therefore, that there is considerable payoff from private equity investments even after 10 years of operation.

In conclusion, draw downs (cash outflows) and distributions (cash inflows) are the raw inputs when assessing fund performance, but there is another ingredient: the time profile of cash flows. The later the cash outflows, and the sooner the cash inflows, the better is a fund's performance. Tables 2 and 3 show that these cash flows occur throughout the life of the fund and thus must be taken into account at the time they occur when calculating a fund's return.

4.4 The Returns of Private Equity Funds

Our primary measure of a fund's return over its life is the internal rate of return on invested capital, taking into account the exact time profile of investments and distributions. IRRs are net of carried interest and management fees and so represent actual returns to the Limited Partner. As a rule, capital gains are not reinvested in the fund, making the calculation straightforward.

We only count cash events (cash flows into and from portfolio companies and annual management fees) and ignore unrealized capital gains (including stock distributions held in inventory) or capital losses. Thus, our IRRs differ from those reported in aggregate form by Venture Economics (and used in prior studies) which represent a mixture of growth in net asset value (NAV) for unrealized investments and cash IRRs for realized investments.

Internal rates of return are calculated to the earlier of the fund's liquidation date or the last data entry (5/31/2001 for most funds, 9/30/2002 for some funds nearing maturity). As we will show, IRRs change over the life of a fund. Clearly, we have final IRRs only for liquidated funds. For funds that have completed their commitment period (are fully invested) but have not yet been liquidated, the IRR we compute is a lower bound on the eventual return: future successful exits can only increase the IRR.¹⁰ For funds that are still actively investing, future IRRs could move up or down. IRRs cannot be computed for funds that have yet to experience any positive cash flows.

Table 5 presents IRRs broken down by vintage year. We report both average and size-weighted IRRs, alongside the standard deviation and the 25th, 50th, and 75th percentile returns. In Panel A, we show IRRs for all 1981-2001 funds for which IRRs can be calculated. This mixes liquidated funds with funds still in operation. The average cash IRR is -14.6 percent, or -30.8 percent value-weighted, with a standard deviation of 60.3 percent. The median fund loses 3.7 percent. By construction, these IRRs are vastly understated as future distributions are given a value of zero.

¹⁰ This ignores future management fees. In the absence of additional exits, management fees will reduce the IRR. However, most funds charge a lower management fee after the commitment period, so the effect is small.

To rectify this problem, in Panel B, we show IRRs for all vintage years between 1981 and 1993 (that is, the mature funds).¹¹ The returns on private equity funds are high over this period (ignoring risk for the moment). In particular, the 73 funds raised between 1981 and 1993 experienced IRRs of 19.8 percent on average, with 18.7 percent for the median fund. Interestingly, the value-weighted average of 18.1 percent indicates that smaller funds performed somewhat better than larger ones.

Across the funds, the heterogeneity of returns is much lower than previously documented (see Kaplan and Schoar (2002) and – albeit at the individual investment level – Cochrane (2000)). For example, the standard deviation is 22.3 percent, with the first and third quartile funds returning 9.9 and 28.6 percent, respectively. As we will show, these funds are not anywhere near as diversified as aggregate market indices, so the consistency of their returns seems impressive. If we were willing to make the (rather dubious) assumption that the fund returns are cross-sectionally i.i.d., then the results extrapolate to volatility risk on the order of the aggregate market and Sharpe Ratios close to one.

One particular advantage of our analysis is that we know the exact timing of the cash flows in calculating the fund returns. Venture Economics reports a performance measure that is also (potentially) based on cash flows, but that does not take into account the time value of money. This measure, called TVPI, equals cash distributions over invested capital minus one. It is the main performance measure used by Kaplan and Schoar (2002). The difference between our findings and Kaplan and Schoar's, both in magnitude and cross-sectional variation, may in part be related to the lack of discounting in the Venture Economics data. As a test of this conjecture, we estimate the correlation between our IRRs and TVPIs for the 73 mature funds in our sample. The correlation coefficient is only 0.59 which shows that the ranking of fund performance would be different under the two measures. The timing of the cash flows is thus an important factor that needs to be taken into account.

An additional concern is the fact that Venture Economics' TVPI (and IRR) data uses subjective net asset values provided by the funds themselves to value unrealized investments. For example, funds might report a projected value for an investment, its book value, or treat it as zero. The latter is equivalent to our IRR calculation for non-distributed investments for the "on-going funds" raised between 1994 and 2001. As one can see from Table 5, if returns were calculated by mixing

¹¹ The 1993 vintage funds have operated for just under ten years, so it is arguable whether they should be considered mature quite yet. However, excluding them does not materially change the performance estimates: the 59 funds raised between 1981 and 1992 returned about the same as the 1981-1993 cohort. To the extent that the 1993 vintage will earn further capital gains in the future, our performance estimates are conservative.

“mature” and “on-going” funds, one could reach quite different conclusions. Kaplan and Schoar’s (2002) main analysis is based on TVPIs estimated in year 5 of a fund’s life, and so combine cash with accounting data. We have accounting valuations for 28 of the 73 mature funds. Comparing their five-year TVPIs to their eventual IRRs at the end of their lives, we find a correlation of only 0.41.

Small sample problems aside, the performance of funds started in the early to mid 1980s seems to have been better than those raised in the late 1980s and early 1990s. This could be true for a number of reasons, not least randomness. Note though that there was a large surge in both the number of funds and amount of dollars raised in the mid-to-late 1980s (see Table 1). This difference in performance is, therefore, consistent with Gompers and Lerner’s (2000) article on “money chasing deals”. Related to this point, the early 1980s may have represented fundamental excess profit opportunities both in the venture capital sector (via the development of the personal computer) and the buyout sector (via corporate waste in the 1970s). For this to be true, of course, one would have to argue that the private equity sector was an under-represented asset class. Alternatively, perhaps the difference in returns reflects aggregate market movements during this period. We return to the potential effect of these characteristics in the next section.

Finally, in Panel C of Table 5, we report IRRs for funds raised after 1993. IRRs are negative, averaging -34.1 percent (VW: -45.6 percent). On the one hand, this is not surprising: these funds may not yet be fully invested, and if they are, may not yet have had a chance to exit many of their portfolio companies. On the other hand, given the frothy state of the IPO market in 1998-2000 and the much documented decline in the age of IPO issuers (Loughran and Ritter (2002); Ljungqvist and Wilhelm (2003)), we might have expected at least some of these funds to have experienced significant early exit events, especially on the venture capital side. In Section 6, we provide evidence consistent with these conjectures.

How does the IRR of a typical private equity fund evolve over the fund’s lifetime? Table 6 presents fund performance by year since a fund was raised. We ask, what is the IRR_T on the average fund in year T of its life, based on the cash flows up to T . The number of funds for which we can calculate IRR_T varies from year to year, increasing initially as more funds experience positive cash flows allowing an IRR to be computed, and then decreasing as funds of more recent vintages drop out for lack of data.

Table 6 reveals three novel stylized facts. First, and not surprisingly given the timing of outflows and inflows, average, value-weighted, and median IRRs increase with fund life. IRRs start

out negative, averaging -84.1 per cent in year 1, and increase monotonically as more portfolio companies are exited. Second, note that it takes almost until year 8 for average and median IRRs to turn positive. Value-weighted IRRs do not turn positive until year 9. Since the cost of capital is clearly not zero, excess returns – the difference between the IRR and the cost of capital – are not realized until even later. Third, and to this point, even by year 10 the exit process is still not complete: while the IRR of the cash flows received up to year 10 is 16.5 percent, post-year 10 cash flows eventually increase the IRR to 21.4 percent on average.

4.5 Sample Selection

A reasonable issue with respect to the results of Sections 4.1 to 4.4 is the degree to which they represent the private equity fund industry as a whole. There are several ways to address this sample selection question. First, one could compare the makeup of our sample to the universe of funds. In Section 3.1, we showed that our sample was skewed towards non-venture funds, i.e., buyout funds and the like. For example, whereas the universe of funds represents three-quarters venture funds, our sample is the opposite in that it represents three-quarters buyout funds. Moreover, while we have data on only a small fraction of the universe of funds, our funds account for 17.5 percent of total capital commitments, and 29.3 percent of buyout capital commitments. Thus, our sample represents a reasonable cross-section of large buyout funds and a much smaller cross-section of venture funds. This does not mean that the characteristics above cannot be applied generally to all private equity funds, only that there must be the caveat that any differences in either investment or performance patterns between these groups will not be captured.

Second, we could bring evidence to bear on this issue by comparing equivalent measures from both samples.¹² For example, Kaplan and Schoar (2002) report average performance data for the population of funds in Venture Economics, by vintage year and in aggregate, using the TVPI measure. Restricting the sample to mature funds (i.e., 1981-1993) leaves them with 692 funds and a TVPI of 2.24. In contrast, in our sample of 73 (albeit larger) funds, TVPI averages 2.59. A *t*-test of the difference, however, is not significant at conventional levels (i.e., a *p*-value of 0.10). Nevertheless, by vintage year, our sample outperforms the Venture Economics population in eight of the eleven years. Thus, there is some support for our sample of primarily larger buyout funds outperforming the Venture Economics universe. Of course, an important caveat is our earlier point,

¹² We thank Steve Kaplan for this suggestion.

made in Section 4.4 above, that the correlation between the fund's actual return and this Venture Economics ratio is tenuous at best.

Third, and finally, as in all studies with limited samples, there is the question of selection bias. There are two possibilities here. The first is that the Limited Partner had extraordinary fund-picking ability. We tend to discount this theory. As described in Section 3.1, the Limited Partner's primary motivation for investing in these funds was to build relationships for the benefit of its corporate parent. This somewhat explains the skewness towards larger buyout funds as these are more likely to provide benefits on the relationship front. Moreover, the Limited Partner itself does not have the setup one might normally observe in a professional fund-picking ("fund-of-funds") organization. The second possibility is that the Limited Partner survived these past 20 years and so we are perhaps looking at an exceptional Limited Partner in that sense. This point is also not particularly relevant as investing in private equity accounts only for a small part of the Limited Partner's overall business.

What is true, however, is that the private equity industry has survived these past 20 years, and the Limited Partner has been part of this process. Therefore, the results documented in this paper, and elsewhere, may have more to say about the historical performance of private equity than having any predictive content per se. That is, this asset class as a whole is a survivor, and it survived because of ex post good draws.¹³

5. The Relative Performance of Private Equity Funds

5.1 Measuring Relative Performance

In this section, we use two methods for measuring the relative performance of private equity funds. Our first method measures excess IRRs, that is, the difference between a fund's IRR and the return on the public equity market. At its simplest, we take as the market return the IRR of investing in an index at the time a fund is raised, and selling at the end of year 10. In some sense, however, this is an apples and oranges comparison. Fund IRRs appropriately give weight to the timing of the cash flows. This includes the fact that it takes many years for committed capital actually to be invested. A like-with-like comparison requires holding constant the time profile of the investment. Thus, we aim to compare investing a total of one dollar in a private equity fund, spread out over the fund's life, to investing a total of one dollar in the market index, spread out in the same way and

¹³ This recalls the point made in a different context about the size of the equity risk premium in the U.S. stock market relative to other countries (see, for example, Goetzmann and Jorion (1999)).

held until the end of the private equity fund's life. In this way, we compare returns over both roughly equivalent time periods and with similar durations. We consider two different draw down schedules (the average fund's, as reported in Table 3, and each fund's actual schedule) as well as two alternative "exit" valuations for the public equity investment strategy (selling at the end of year 10, or at the average index value during year 10). We also consider two possible public equity indices, the S&P 500 and the Nasdaq Composite Index.

There are two potential problems with the excess IRR measure. First, as Table 3 documents, many funds distribute capital prior to their liquidation. IRR calculations implicitly assume that these early distributions can be reinvested at the fund's IRR. This might be a reasonable assumption in our case, given the Limited Partner's multiple and ongoing investments in this asset class, but it will tend to attenuate the differences (both negative and positive) in relative performance.¹⁴ Second, IRR calculations assume one discount rate for all cash flows. One can reasonably argue that outflows (i.e., investments) should be discounted at a different, and lower, rate than inflows. If so, IRRs will tend to overstate the performance of the fund relative to the true risk profile of the cash flows.

An alternative measure that addresses these concerns is to calculate the ex post net present value (NPV) of investing in a fund, using the realized cash flows discounted at the risk-free rate for outflows and the cost of capital for inflows. As a first pass, we use the Treasury-bond rate with corresponding maturity for the outflows, and the expected return on the aggregate market for the inflows. The NPVs are scaled by the present value of the investment, giving the so-called Profitability Index. Intuitively, the Profitability Index can be thought of as the present-valued return on invested capital, that is, the excess value created for each dollar invested. Thus \$1 invested in private equity is worth one plus the Profitability Index in present value terms.

Table 7 presents the distribution of the excess IRRs and the Profitability Indices, relative to various benchmarks, for the mature funds raised between 1981 and 1993. Consider first the IRR calculations given in Panel A. Private equity generates positive excess returns relative to public equity irrespective of the assumptions underlying the draw down schedule or market benchmark. For example, assuming all the funds are invested immediately in the S&P 500, private equity produces 8.06 percent mean and 6.04 percent median excess returns on an annualized basis. While the Nasdaq benchmark reduces the relative magnitudes of these returns, the numbers are still impressive, with a 6.28 percent mean and 4.01 percent median excess return.

¹⁴ Intuitively, one might consider this a comparison between two long-term investment strategies, one in private versus the other in public equity.

Moving across the different draw down schedules, from the aforementioned immediate investment strategy to average draw downs to the fund's actual draw downs, excess returns decline, both on average and for the first quartile, median, and third quartile fund. This reduction in excess returns implies that the funds either have market timing ability, or, alternatively, make investments that perform in unison with the aggregate market. However, excess returns remain substantial, averaging at least 5.71 percent per annum.

Table 7, Panel B presents the distribution of the Profitability Index for the two benchmark indices. With respect to the S&P 500, the mean and median values are 25.07 percent and 12.18 percent, respectively. Thus, assuming the cost of capital is the S&P 500's average return over the fund's life, private equity funds create 25 cents in excess value for every dollar invested, in present value terms. Part of the 25 cents reflects, of course, the cost of liquidity. Since the NPV correctly takes into account the timing of cash flows and the different discount rates for inflows versus outflows, the 25 percent estimate may be more reflective of the excess value to private equity than the excess IRR measures reported in Panel A.

The fact that the median Profitability Index estimate is so far below the mean suggests that there is a fair amount of skewness in the distribution of possible values. In fact, the distribution suggests there is a significant downside in the form of funds performing poorly on a relative basis. For example, first-quartile funds lost 23.65 percent or more relative to an investment in the S&P 500.

Finally we note that using Nasdaq-based discount rates reduces the mean excess value to only 9.96 percent, with an even larger degree of cross-sectional variation across funds. The median excess value, for example, is now negative. The differences here are sufficiently large to explore the funds' risk characteristics in order to better understand which cost of capital is appropriate.

5.2 The Risk of Private Equity Funds

In Section 5.1, we compared the returns on private equity to a similar investment in the aggregate public equity market. Ideally, we would like to adjust for any systematic risk differences between the private equity funds and the aggregate market. This is nontrivial because of the lack of any meaningful time series. At first glance, this may seem surprising given our 20 years of data. Previous analyses, such as Kaplan and Schoar (2002), have looked at the covariation between a fund's returns (as reported to Venture Economics) and the market return. Our empirical results above suggest this is probably not a valid approach. Note that the returns on the fund investments do not actually get realized until years 9 to 10. Thus, reasonably, the only measurable return is over

a 10-year horizon. This leaves us with only two truly independent observations on the market return and thus time-series measures of betas are inappropriate. A sensible alternative is to somehow extract the information cross-sectionally.

One obvious cross-sectional distinction is whether the fund is a venture or non-venture (i.e. buyout) fund. Both the nature of the companies invested in and the degree of leverage are different enough for these types of funds that this distinction might be an important determinant of cross-sectional variation in fund performance. Thus, as a first pass, we break down the IRR results of Tables 5 and 6 into venture versus non-venture funds. The results are presented in Table 8. At least for our sample of firms, buyout funds have generally outperformed venture funds. For example, over the period 1981-1993, the IRR averages 21.83 percent for buyout funds versus only 14.08 percent for venture funds. Of course, our sample of venture funds is fairly small as the Limited Partner invests disproportionately in buyout funds. Interestingly, Panel B shows that distributions occur much more quickly for venture funds (i.e., less negative IRRs a few years out). This is as one might suspect, but in our sample venture funds still take longer to break even, consistent with the “hit or miss” nature of venture investing.

The fact that buyout funds outperform venture funds on average is consistent with the higher leverage buyout funds typically employ. We will return to leverage shortly. First, however, we ask whether the underlying systematic risk faced by these funds is different. In order to answer this question, we need to analyze each fund’s portfolio of companies.

We employ the following three-step procedure to estimate the risk of each private equity fund. First, using the information described below, we identify each portfolio company held by a fund. Second, given this identification, we assign portfolio companies to one of forty-eight broad industry groups chosen by Fama and French (1997). For each of these industries, Fama and French estimate an equity beta over a five-year period, 1989-1994. Assuming that the leverage of the private company coincides with that of the industry, we then assign the industry beta to the portfolio company. Finally, we compute the average equity beta of the fund using the capital disbursements as weights.

For reconstructing the portfolio composition of each fund, we use the following three sources:

- Venture Economics: Venture Economics' "Portfolio Companies (VIPC)" database records the holdings of most private equity funds in our dataset, albeit with some error.¹⁵
- Snapshots: In 1999 and 2000, the Limited Partner compiled a snapshot of each fund's portfolio composition, giving the name of the portfolio company and a short verbal description of its product, service, or industry, using information provided by the GP in the quarterly or annual partnership reports. The snapshots suffer from two shortcomings: (i) they do not cover funds liquidated prior to 1999-2000; and (ii) they are based on GP reports as of a certain date, so portfolio companies that have been written off are typically not included.
- Cash flow records: We augment the data with information taken from the cash flow records. However, on occasion, a company name is either not disclosed or is abbreviated in such a way as to make a positive identification impossible. Moreover, companies may have changed their names since first entered in the cash flow records, implying we may overstate the number of distinct portfolio companies.

We eliminate overlap between these sources, taking particular care to trace name changes using information provided in Venture Economics, news searches, and annual or quarterly partnership reports.

None of these sources reports SIC codes, so to assign portfolio companies to industries, we proceed as follows:

- In the case of companies that have gone public, or were public prior to being bought out, we use their CUSIPs as reported in Venture Economics to identify the primary four-digit SIC codes disclosed in S.E.C. filings.
- We look up all other companies in Dun & Bradstreet's *Million Dollar Database*, which provides summary information (including primary SIC codes) on approximately 1.6 million public and private companies in the U.S. and Canada. To avoid matching with the wrong company, we require not only a name match but also that the business description agree with Dun & Bradstreet's SIC code assignment.
- For companies that we cannot find in D&B, or cannot identify unambiguously, we use the verbal business descriptions reported by Venture Economics, and those included in the

¹⁵ There are two potential sources of error in the Venture Economics data. First, acquirer stock received as payment in the sale of a portfolio company appears as a portfolio investment. Second, and more seriously, Venture Economics frequently allocates investments to the wrong fund in a sequence of funds managed by the same general partner (say Fund I rather than Fund IV).

snapshots compiled by the Limited Partner, to manually assign companies to Fama-French industries. We fill gaps using the GPs' and company web sites and news sources.¹⁶

Table 9, Panel A describes the characteristics of the 73 mature funds' portfolio companies, broken down by buyout versus venture capital fund. Putting systematic risk aside for the moment, the table illustrates some important differences between these types of funds. Buyout funds tend to invest in far fewer companies than venture funds. For example, the mean and median number of companies are respectively 16.1 and 13 for buyout versus 37.3 and 32 for venture funds. Since investors can diversify themselves, the lack of diversification on the buyout fund front does not seem problematic. Nevertheless, it is an important factor in understanding the heterogeneity of returns across funds. In contrast, buyout and venture funds do not differ in the distribution of either the fraction of companies or the fraction of capital that is invested in a single, dominant industry. Specifically, over one-third of the companies (and 40 percent of the dollar amount) represent investments in just one industry. Thus, private equity funds tend to specialize in, or give weight to, one industry much more than the average public equity fund.

Most importantly, Panel A presents the distribution of the funds' betas using the aforementioned designation method. Three observations are in order. First, even though the buyout funds invest in relatively few companies, the distribution of their portfolio betas is fairly tight across funds. Specifically, their portfolio betas average 1.08, with first and third quartiles of 1.04 and 1.13. Thus, it is immediately apparent that the cross-sectional variation of buyout fund returns will not be fully explained by their underlying portfolio risk (assuming similar leverage across funds). Second, to the extent that buyout funds are more levered than the industries they invest in, these estimates suggest beta risk higher than the market. Third, venture funds have only slightly higher systematic risk than buyout funds, with betas averaging 1.12 and first and third quartiles of 1.09 and 1.16. Since venture funds tend to be less levered than the underlying Fama-French industries (as they invest in younger companies), there is some evidence to suggest the higher average returns in the buyout industry may partially reflect leverage differences.

To better understand the relation between a fund's return and its portfolio risk, we estimate the correlation between the fund's IRR and its portfolio beta (not reported in the table). The correlation is generally small. For example, the correlation is 0.01 overall, though this masks an important

¹⁶ Due to syndication, some companies appear in the portfolios of more than one GP. Occasionally, the GPs' product descriptions are in conflict and so would lead to different Fama-French assignments. We resolve such conflicts using the above sources.

difference between venture and non-venture funds: venture fund IRRs correlate *negatively* with portfolio betas (-0.073), while non-venture fund IRRs correlate positively (0.067). When we adjust for market returns, the correlation between a fund's excess IRR and its portfolio beta is -0.149 for venture funds and 0.059 for non-venture funds. While the portfolio betas are clearly measured with some error, these correlations suggest that cross-sectional variation in fund performance is most probably not due to risk differences. Put differently, systematic risk may not be one of the main determinants of the cross-section of private equity returns.

Table 9, Panel B reports risk-adjusted Profitability Indices for the 73 mature funds overall and broken down by venture versus buyout fund. Profitability Indices are computed as in Table 7, except that we now discount cash inflows at risk-adjusted returns defined as the riskfree rate plus the fund's estimated beta (based on its portfolio of companies) times the risk premium. We use two alternative estimates of the risk premium: (i) Fama and French's (1997) measure (which we denote the ex ante cost of capital in the table), and (ii) the mean excess return on the S&P 500 measured over the life of the fund (denoted the ex post cost of capital). As before, the riskfree rate is the yield to maturity on ten-year treasuries computed in the month the fund was raised.

Several observations are in order. First, we continue to find that private equity funds create excess value. This should not be surprising given that betas aren't all that high (leverage considerations aside). Depending on the risk premium used, the mean Profitability Index is between 24 percent (ex post) and 32.23 percent (ex ante). Second, these excess values hold for both venture and buyout funds, though, in our sample, buyout funds create more value. Third, there is substantial cross-sectional variation using these measures of risk-adjusted value creation by private equity funds. For example, using the ex post cost of capital, a quarter of the funds are worth at least 46.95 percent more than the present value of their invested capital, while a quarter lose 27.39 percent or more of their invested capital. Thus, risk adjustment does not tighten the dispersion in performance across funds.

Of course, returns on buyout funds may be high simply because such funds tend to be highly levered. While we do not have data on the funds' leverage ratios, we can do back-of-the-envelope calculations to better understand the possible impact of leverage. Specifically, we ask: how much less equity than the Fama-French industries would the funds have to have had in order for their excess value to be zero? Taking the average buyout fund's IRR of 21.83 percent and assuming the riskfree rate is eight percent, the market return is the 14.1 percent return on the S&P index, no taxes, a zero liquidity premium, and all debt is riskfree (a conservative assumption), the equity beta at

which the average fund breaks even in present-value terms is 2.267. The Fama-French equity beta for the average buyout fund in our sample is 1.08. Therefore, buyout funds create excess value as long as they don't use less than $1.08/2.267=0.476$ the amount of equity used by the firms in the Fama-French industries. For instance, if the debt-equity ratio is 30/70 in the Fama-French data, then a fund still creates excess value up to a debt-equity ratio of 66.66/33.34. If debt is not riskfree, then the debt-equity ratio consistent with excess returns is even higher.

5.3 *The Determinants of Private Equity Returns*

To investigate the determinants of private equity returns, we regress the excess IRRs from Table 7 on a number of variables.¹⁷ By using the excess IRRs, we have by construction adjusted for general market movements that match the timing of the funds' outflows and inflows. The explanatory variables we use are fund size (in logs and levels to allow for nonlinearities), a dummy for first-time funds, the log of the total amount of capital committed to buyout (venture) funds in the year the sample buyout (venture) fund was raised, the portfolio beta, and one of the four measures of portfolio diversification reported in Table 9 (the number of portfolio companies, the fraction of the portfolio that is invested in the dominant industry, by number and by invested capital, or the Herfindahl index of the portfolio weights).¹⁸ We include the latter to test Jones and Rhodes-Kropf's (2002) prediction that private equity funds that have more idiosyncratic risk earn higher returns.

We initially pool venture and buyout funds (a dummy distinguishing them is never significant), though we also report results for buyouts separately. The least-squares estimates are presented in Table 10. Overall, the explanatory power of the regressions is low, ranging from an adjusted R^2 of 3.7 percent when we use the log of the number of portfolio firms to proxy for idiosyncratic risk to 5.7 percent when we use the Herfindahl. Excess IRRs increase significantly with the log of real fund size and decrease with its level. Using the coefficient estimates, we compute that excess IRRs reach a maximum at fund sizes between \$1.1 billion and \$1.2 billion, depending on the specification. If we restrict the estimation sample to buyout funds only, on the other hand, we find no significant size effect.

¹⁷ The results carry through for the various excess IRR measures as well as the raw IRRs.

¹⁸ Kaplan and Schoar (2002) also look at the determinants of fund performance, measured using the TVPI (the ratio of cash distributions and, for unrealized investments, accounting values to total cash invested) as of fund year 5. Unlike us, they do find that venture funds perform significantly better than buyout funds. They also control for the sequence number of the fund, which we don't. Their results regarding the effect of fund size and first-time funds are similar to ours.

First-time funds perform somewhat better than follow-on funds, but the difference is not significant. Our measures of idiosyncratic or portfolio risk mostly correlate positively with fund returns, but never significantly so.

One interesting feature of the regression analysis is the difference in the systematic risk coefficients in the buyout-only sample (i.e., 0.19) versus in the overall sample (i.e., negative values ranging from -0.03 to -0.17). While the estimates are not significantly different from zero, the positive relation between portfolio beta and fund returns in the buyout sample is consistent with intuition. Economically, however, the effect is small: a one-standard deviation increase in beta would lead to only a 0.02 percent increase in the excess IRR on buyout funds.

Apart from size, the only significant determinant of excess IRRs in our specifications is fund inflows: the more money was raised in the fund's vintage year, the worse is the fund's subsequent performance. This is true both in the overall sample and for buyouts. It is reminiscent of Gompers and Lerner's (2000) finding that inflows into private equity funds increase the prices funds pay for their investments. We show that "money chasing deals" tends to lead to lower excess returns. In fact, this is the one variable which seems to be a determinant for both buyout funds and venture funds.

6. The 1998-2001 Period

It would be remiss for us not to address the remarkable period of private equity investing between 1998 and 2001. As is well documented now, considerable amounts of capital were raised, invested, distributed and, most probably, lost during this time. Table 1 shows that, in terms of both the number of funds and dollars raised, this four-year period almost matches the previous fifteen years. While our Limited Partner is arguably not quite as active as the aggregate industry during this period, nevertheless, over 40 percent of all the funds it invested in were raised in 1998-2001.

Funds that could exit investments before the bubble burst achieved possibly unsustainably good returns. Funds that made portfolio investments during that time may have paid excessive valuations for companies that are now next to worthless, perhaps because too much capital was being raised that ended up "chasing too few deals" (Gompers and Lerner (2000)).

Does this mean that the good performance we have documented overstates the expected return on private equity? Possibly, but we offer three caveats. First, the differences between our estimates of excess private equity returns and those in the existing literature are not due to different treatment of this period: neither includes much by the way of late-1990s returns. (Because they are not yet

mature, funds raised in 1998-2001 have not been included in our empirical performance estimates.) Second, the majority of our results, certainly for the mature funds, refer to the buyout industry rather than to venture funds. Buyout funds generally did not invest in the internet sector, so our estimates of private equity returns for buyout funds are not necessarily “inflated”. Third, as the stock market itself fell beginning in 2000, and the economy went into recession, the relevant question is whether the portfolio companies of buyout funds fell in value by more or less than the public equity market as whole.

In order to shed more light on possible differences between recent and older funds, Table 11 reports draw down schedules (Panel A) and capital return profiles (Panels B and C) separately for pre-1998 and post-1998 funds. We also break them down by venture and non-venture funds. Consider first the non-venture funds. Note that the draw down rate is somewhat lower in the post-1998 period: by year four, recently raised funds had invested 61.72 percent of committed capital, compared to 72.47 percent for the older funds. At the same time, they had returned invested capital a little faster: cumulative distributions after four years average 43.14 percent, as compared to only 26.29 percent for funds raised pre-1998. To some extent, this is due to less capital having been invested in the first place, and so distributions of *committed* capital are no different.

Though the sample size is smaller, venture funds follow the same patterns, only more so. That is, funds were drawn down at a slower rate and capital was returned more quickly post-1998. Four years out, venture funds had invested only 68.47 percent of committed capital versus 82.08 percent for the pre-1998 funds, yet had already returned 134.42 percent of invested and 112.81 percent of committed capital on average (versus only 35.10 percent and 34.85 percent pre-1998).

In sum, recent funds have a slower draw down schedule, both among venture and non-venture funds, but they return capital more rapidly, presumably due to the favorable exit climate in the late 1990s. Such accelerated distributions are particularly evident among venture funds. This is consistent with an active IPO market during most of this period which private equity funds took advantage of. Taking the capital return numbers by year four as given, recent funds appear to have banked much of their returns earlier than usual. Given that the aggregate stock market very nearly collapsed beginning in 2000, there is little support for these funds underperforming public equities during this time. Of course, funds raised in 2000 and 2001 may be less fortunate, but it is too early to assess their performance.

7. Final Remarks

Using a unique dataset of private equity funds raised over the last two decades, we have analyzed the cash flow, return, and risk characteristics of private equity. Unlike previous studies, we have access to cash flow data that allow us to precisely estimate performance, taking into account the exact timing of investment and capital return flows, rather than relying on fund managers' subjective valuations.

We document several stylized facts that are important for understanding the dynamics of private equity investing and performance measurement. The draw down schedules and capital return distributions suggest that the timing and illiquidity of the cash flows can be important. For example, it takes over three and six years, respectively, to invest 56.9 percent and 90.5 percent of committed capital, and over eight and 10 years, respectively, for IRRs to turn positive and eventually exceed public equity returns. Moreover, these schedules are not random but depend on the underlying conditions of the market (i.e., available investment opportunities) and competition amongst firms (i.e., the origination levels of private equity funds).

In contrast to recent empirical claims, we document that private equity generates excess returns on the order of five to eight percent per annum relative to the aggregate public equity market. These excess returns are robust to assumptions about the timing of investments in the public equity market, to measures of risk across the portfolio companies, and to various measurement methodologies. For example, while we estimate the betas of the private equity funds' portfolios to be greater than one, we show that on a risk-adjusted basis the excess value of the typical private equity fund is on the order of 23.8 percent relative to the present value of the invested capital. One possible interpretation for this magnitude is that it compensates investors for holding a 10-year illiquid investment. A cross-sectional analysis of excess IRRs suggests the source of the outperformance is not necessarily compensation for systematic risk, but it may be related to the type of fund, such as buyout versus venture, and the timing of the fund (relative to the overall number of private equity funds).

References

- Blaydon, Colin, and Michael Horvath, 2002. "What's a company worth? It depends on which GP you ask." *Venture Capital Journal*, May, 40-41.
- Blaydon, Colin, Michael Horvath, and Fred Wainwright, 2002. "Venture capital survey." Unpublished working paper, Foster Center for Private Equity, Dartmouth College.
- Brealey, Richard A., and Stewart C. Myers, 1996. *Principles of Corporate Finance*. McGraw-Hill.
- Cochrane, John, 2000. "The risk and return of venture capital." Unpublished working paper, University of Chicago.
- Fama, Eugene F., and Kenneth R. French, 1997. "Industry costs of equity." *Journal of Financial Economics* 43, 153-194.
- Goetzmann, William N., and Philippe Jorion, 1999. "Global stock markets in the twentieth century." *Journal of Finance* 54, 953-980.
- Gompers, Paul A., 1996. "Grandstanding in the Venture Capital Industry," *Journal of Financial Economics* 43, 133-156.
- Gompers, Paul A., and Josh Lerner, 1996. "The use of covenants: An analysis of venture partnership agreements." *Journal of Law and Economics* 39, 463-498.
- Gompers, Paul A., and Josh Lerner, 1997. "Risk and reward in private equity investments: The challenge of performance assessment." *The Journal of Private Equity*, Winter, 5-12.
- Gompers, Paul A., and Josh Lerner, 1999. *The Venture Capital Cycle*. MIT Press.
- Gompers, Paul A., and Josh Lerner, 2000. "Money chasing deals? The impact of fund inflows on private equity valuations." *Journal of Financial Economics* 55, 281-325.
- Jones, Charles M., and Matthew Rhodes-Kropf, 2002. "The price of diversifiable risk in venture capital and private equity." Unpublished working paper, Columbia University.
- Kalbfleisch, J.D., and R.L. Prentice, 1980. *The Statistical Analysis of Failure Time Data*. John Wiley & Sons, New York.
- Kaplan, Steven N., and Antoinette Schoar, 2002. "Private equity returns: Persistence and capital flows." Unpublished working paper, University of Chicago.
- Ljungqvist, Alexander, and William J. Wilhelm, 2003. "IPO pricing in the dot-com bubble." *Journal of Finance*, 723-752.
- Loughran, Tim, and Jay R. Ritter, 2002. "Why has IPO underpricing changed over time?" Unpublished working paper, University of Florida.

Moskowitz, Toby, and Annette Vissing-Jørgensen, 2002. "The returns to entrepreneurial investment: A private equity premium puzzle?" *American Economic Review* 92, 745-778.

Quigley, John M. and Susan E. Woodward, 2002. "Private equity before the crash: Estimation of an index." Unpublished working paper, University of California at Berkeley.

Table 1. Sample overview

The sample consists of private equity and venture capital funds raised between 1981 and 2001 (the “vintage years”). 73 funds raised between 1981 and 1993 form our sample of “mature” funds. To protect the identity of the Limited Partner, we have agreed not to disclose certain characteristics of the funds raised after 1993. VC funds are those identified as “venture capital” by Venture Economics. Most non-venture funds are flagged as “Buyout” (90.4 percent); the remainder are flagged as “Generalist Private Equity” (3.8 percent), “Mezzanine” (4.8 percent), and “Other Private Equity” (1 percent). Fund size is the capital committed by investors to a fund in all closings, as reported by Venture Economics and corrected where needed using partnership reports prepared by the fund managers. Total fund size is the aggregate amount raised by all sample funds. Mean fund size is equally weighted. Commitment is the Limited Partner’s capital commitment to the funds. Total commitment is the aggregate commitment by the Limited Partner. Mean commitment is equally weighted. All monetary numbers are in nominal U.S. dollars. For comparison, we report the number of funds raised according to Venture Economics, for funds headquartered in the same countries as our sample funds (the U.S. and certain countries in Europe and Latin America).

Vintage Year	No. of funds raised acc. to Venture Economics	Sample size No. of funds in data set	Fraction VC funds		Fund Size (\$ million, nominal)				Commitment (\$ million, nominal)				Mean commitment/fund size (in %)
			by number (%)	by fund size (%)	Total	Mean	Min.	Max.	Total	Mean	Min.	Max.	
1981	94	1	0.0	0.0	75.0	75.0	75.0	75.0	5.0	5.0	5.0	5.0	6.67
1983	181	2	50.0	75.0	454.6	227.3	113.6	341.0	18.2	9.1	5.6	12.6	4.31
1984	199	5	40.0	25.7	276.5	55.3	35.0	100.0	14.4	2.9	1.2	5.5	6.10
1985	185	4	25.0	5.1	724.6	181.2	36.6	400.0	24.1	6.0	2.1	10.5	4.72
1986	159	6	16.7	51.9	2,265.6	377.6	25.0	1,175.0	36.4	6.1	1.3	10.0	5.40
1987	198	8	37.5	2.2	6,458.0	807.3	25.0	5,600.0	210.6	26.3	2.0	167.4	5.68
1988	213	12	8.3	16.0	11,122.9	926.9	100.0	2,200.0	355.5	29.6	8.8	83.0	4.42
1989	264	11	18.2	5.4	4,108.1	373.5	46.0	1,066.0	190.4	17.3	2.5	99.4	4.05
1990	236	4	50.0	11.2	1,246.6	311.7	24.3	1,015.5	12.6	3.2	0.8	6.1	2.39
1992	209	6	33.3	3.3	2,937.8	489.6	15.0	1,020.0	55.0	9.2	2.0	25.0	4.18
1993	261	14	28.6	4.6	7,033.9	502.4	27.0	1,880.0	184.8	13.2	3.2	40.7	4.59
1994	304	*	12.5	7.4	8,771.9	*	*	*	*	*	*	*	*
1995	408	*	15.4	7.3	7,261.4	*	*	*	*	*	*	*	*
1996	439	*	22.2	6.8	15,714.4	*	*	*	*	*	*	*	*
1997	655	*	29.4	7.7	19,639.8	*	*	*	*	*	*	*	*
1998	800	*	17.1	17.4	36,832.0	*	*	*	*	*	*	*	*
1999	1,087	*	20.0	14.3	32,309.2	*	*	*	*	*	*	*	*
2000	1,872	*	34.8	22.2	49,314.3	*	*	*	*	*	*	*	*
2001	829	*	100.0	100.0	464.6	*	*	*	*	*	*	*	*
Total	8,593	*	24.9	14.8	207,011.1	*	*	*	*	*	*	*	*
1981-1993	2,199	73	26.0	11.8	36,703.6	502.8	15.0	5,600.0	1,107.0	15.2	0.8	167.4	4.64

Table 2. Draw downs and capital returns by vintage year

When fund managers invest in portfolio companies, they draw down the limited partners' capital commitment. Following liquidity events (such as an IPO), capital is returned to the limited partners in the form of cash or stock distributions. In the latter case, the LP may either sell the stock directly or hold it as a public market investment. We record only stock distributions that are sold (as virtually all are in our sample). The average fund in our sample has drawn down 67.32 percent of committed capital. However, this understates draw downs as the more recent funds in the sample are not yet fully invested. Therefore, we also report draw down and distribution schedules for the 59 (73) funds raised between 1981 and 1992 (1993).

Vintage Year	No. of funds	Average draw downs / committed capital	Fraction of funds that are 70% invested	Fraction of funds that are 80% invested	Average distributions / capital invested	Average distributions / committed capital
1981	1	0.9991	1.000	1.000	3.2780	3.2751
1983	2	0.8973	1.000	0.500	3.2168	2.9249
1984	5	0.9688	1.000	1.000	3.0794	2.9797
1985	4	1.0112	1.000	1.000	5.1357	5.1416
1986	6	1.0003	1.000	1.000	3.8571	3.8577
1987	8	0.8654	0.750	0.625	2.6453	2.3634
1988	12	0.9780	1.000	1.000	2.0259	1.9661
1989	11	0.9517	1.000	0.909	2.6084	2.4332
1990	4	0.9217	1.000	0.750	1.9637	1.7902
1992	6	0.9027	0.833	0.833	1.8777	1.6396
1993	14	0.9462	1.000	0.929	1.9346	1.7836
1994	*	0.9313	0.938	0.875	1.3123	1.1882
1995	*	0.9101	1.000	0.923	1.2377	1.1478
1996	*	0.9016	0.944	0.889	0.8367	0.7804
1997	*	0.7632	0.618	0.441	0.5130	0.4348
1998	*	0.6511	0.400	0.400	0.5966	0.4377
1999	*	0.4119	0.100	0.025	0.1995	0.0918
2000	*	0.1906	0.000	0.000	0.1187	0.0130
2001	*	0.1831	0.000	0.000	0.0001	0.0000
1981-2001	*	0.6732	0.556	0.495	1.0683	0.9434
1981-1992	59	0.9477	0.949	0.881	2.7471	2.6102
1981-1993	73	0.9474	0.959	0.890	2.5913	2.4517

Table 3. Draw downs and capital returns by fund year

Funds are typically ten-year limited partnerships, with possible extensions by a few years subject to the limited partners' approval. At the end of its life, the general partner "liquidates" the fund by selling all remaining assets and distributing the cash to the limited partners. The liquidation phase can potentially take a few years. The table shows average annual and cumulative draw downs for each year of a fund's life, divided by committed capital, and the corresponding average cumulative distributions divided by invested and by committed capital. The average fund draws down 16.28, 20.35, and 20.15 percent of committed capital in its first three years of operation. At the end of its fourth year, it is 72.64 percent invested, and at the end of its expected life (year ten) it is 93.62 percent invested. There are no further draw downs beyond year ten. It takes just under seven years for total invested capital to be returned in the average fund, and around seven years for committed capital to be returned. There are further capital returns beyond year 10, so the last row reports cumulative distributions up to the liquidation date.

Years since first closing (T)	Annual draw downs / committed capital	Cumulative draw downs / committed capital, up to year T	Cumulative distributions / total capital invested, up to year T	Cumulative distributions / committed capital, up to year T
1	0.1628	0.1628	0.0097	0.0014
2	0.2035	0.3662	0.0372	0.0210
3	0.2015	0.5677	0.1660	0.1292
4	0.1587	0.7264	0.3393	0.2913
5	0.1113	0.8377	0.5370	0.4925
6	0.0622	0.8999	0.8323	0.7688
7	0.0235	0.9234	1.0859	1.0086
8	0.0095	0.9329	1.4433	1.3458
9	0.0050	0.9379	1.8587	1.7456
10	-0.0003	0.9362	2.0663	1.9317
10 or more		0.9362	2.7817	2.6247

Table 4. The determinants of draw down rates

The dependent variable is the log of the time (in years) between a fund being raised and it having drawn down at least X% of its committed capital. We use three cutoffs for X: 70, 80 and 90 percent. The explanatory variables are listed in the table. We estimate accelerated time-to-failure models using maximum likelihood estimators that are corrected for the right-censoring caused by funds leaving our sample before they are fully invested. We thus include all sample funds raised between 1981 and 2001 in the analysis. The error is assumed to have an exponential distribution with mean β_0 (the constant). This model is identical to a proportional-hazard duration model, and coefficients can easily be converted into hazard ratios. Standard errors are shown in italics. We use ***, **, and * to denote significance at the one percent, five percent, and 10 percent level (two-sided), respectively.

	time-varying?	Time to investing at least ...		
		70% of committed capital	80% of committed capital	90% of committed capital
dummy=1 if venture fund	no	0.373 <i>0.259</i>	0.478* <i>0.260</i>	0.487* <i>0.289</i>
dummy=1 if first-time fund	no	-0.021 <i>0.190</i>	-0.197 <i>0.195</i>	-0.280 <i>0.214</i>
dummy=1 if raised in 1981-1993	no	-0.612** <i>0.293</i>	-0.904*** <i>0.284</i>	-0.668** <i>0.282</i>
log real fund size (in \$m)	no	-0.073 <i>0.078</i>	-0.051 <i>0.086</i>	0.032 <i>0.096</i>
log real fund inflows (in \$m), p.a.	yes	-1.065*** <i>0.193</i>	-0.961*** <i>0.201</i>	-0.593*** <i>0.187</i>
dummy=1 if year is 1999 or 2000	yes	-0.967*** <i>0.274</i>	-0.834*** <i>0.291</i>	-0.684** <i>0.289</i>
log real fund inflows, same vintage year	no	1.735*** <i>0.201</i>	1.459*** <i>0.163</i>	1.117*** <i>0.163</i>
BAA corporate bond yield, p.a. (in %)	yes	0.317*** <i>0.107</i>	0.475*** <i>0.120</i>	0.342*** <i>0.130</i>
return on S&P 500 index, p.a.	yes	0.933 <i>0.647</i>	0.915 <i>0.676</i>	0.598 <i>0.695</i>
constant		-6.597** <i>3.194</i>	-5.932* <i>3.039</i>	-5.401 <i>3.455</i>
Pseudo- R^2		25.3 %	27.7 %	24.0 %
Likelihood ratio test: all coeff. = 0 (χ^2)		121.16***	124.43***	99.86***
Number of observations (time at risk)		938	1020	1126

Table 5. Fund performance by vintage year

Internal rates of return are calculated to the earlier of the fund's maturity or the last data entry (5/31/2001 for most funds, 9/30/2002 for some funds). They are based on precisely dated cash flows. Unrealized capital gains (including stock distributions held in inventory) or capital losses are ignored. IRRs are net of carried interest and management fees and so represent actual returns to the Limited Partner. Funds without positive cash flows are excluded as IRRs cannot be computed. The mean IRR is simple weighted. The weighted average IRR is weighted by fund size (as defined in Table 1). "Mature" funds are those for which we have at least nine years of data post-first closing. These have either been liquidated or are likely to have earned most of their returns already. "Ongoing" funds are younger funds that may still earn returns in the future.

Vintage Year	No. of funds	No. of funds with IRR data	Internal rate of return (in percent)					
			weighted mean	weighted average	st. dev.	first quartile	median	third quartile
Panel A: Whole sample								
1981-2001	*	*	-14.59	-30.78	60.25	-56.68	-3.73	18.66
Panel B: "Mature" funds								
1981-1993	73	73	19.81	18.13	22.29	9.85	18.66	28.59
1981-1992	59	59	20.46	17.67	22.42	10.60	18.73	28.59
1981	1	1	23.43	23.43				
1983	2	2	27.82	21.19	18.74	14.57	27.82	41.07
1984	5	5	33.67	32.21	37.72	13.45	16.74	32.52
1985	4	4	33.69	31.19	21.36	22.14	24.43	45.24
1986	6	6	21.54	20.61	10.91	13.27	20.52	26.19
1987	8	8	16.87	7.69	12.10	6.90	18.69	27.10
1988	12	12	12.65	13.55	10.25	9.72	14.13	17.73
1989	11	11	20.38	21.35	12.10	18.09	21.10	28.93
1990	4	4	21.41	24.02	24.75	1.48	17.96	36.18
1992	6	6	16.33	39.48	50.10	15.77	20.04	57.77
1993	14	14	17.08	20.07	22.35	-5.98	12.15	33.41
Panel C: "On-going" funds								
1994-2000	*	*	-34.06	-45.63	66.06	-80.03	-41.58	-0.24
1994	*	*	-1.57	8.92	34.87	-10.60	2.78	20.14
1995	*	*	-0.24	-8.60	37.53	-26.97	-4.45	26.49
1996	*	*	-22.14	-48.25	47.77	-58.04	-19.47	8.88
1997	*	*	-37.85	-32.82	41.30	-65.33	-42.93	-14.21
1998	*	*	-32.98	-58.84	118.35	-88.72	-52.06	-31.33
1999	*	*	-69.53	-60.23	28.43	-98.04	-77.45	-41.42
2000	*	*	-60.72	-55.21	45.82	-98.97	-80.69	2.51
2001	*	*						

Table 6. Performance by fund year

Internal rates of return are calculated to the end of year T since the first closing, with $T = 1 \dots 20$ years. Weighted averages are based on fund size (total committed capital).

Years since first closing (T)	No. of funds with IRR data	Internal rate of return (in percent)					
		weighted mean	weighted average	st. dev.	first quartile	median	third quartile
1	18	-84.09	-67.81	34.39	-99.87	-97.77	-93.23
2	58	-82.38	-81.34	33.86	-99.23	-95.71	-83.13
3	89	-64.91	-71.97	71.18	-97.70	-86.41	-47.34
4	121	-55.30	-64.58	65.44	-94.07	-72.60	-37.91
5	115	-42.24	-50.74	45.45	-83.81	-48.08	-13.05
6	107	-22.09	-24.86	43.29	-53.80	-20.40	6.29
7	99	-9.18	-14.63	37.23	-27.37	-11.81	13.88
8	83	2.63	-6.48	31.58	-9.26	1.53	17.70
9	68	13.20	4.12	25.93	-1.19	10.03	26.09
10	60	16.54	10.60	25.07	2.68	15.10	28.44
11	50	18.01	11.30	19.32	6.19	16.05	26.03
12	43	18.70	11.59	18.35	9.45	17.46	23.15
13	33	18.74	11.86	18.82	9.70	17.95	23.08
14	25	17.80	11.40	12.78	12.69	18.73	24.88
15	14	17.16	10.43	11.21	12.07	15.61	21.94
16	9	19.66	20.17	9.83	14.24	17.49	22.60
17	4	23.87	20.71	12.02	16.20	19.93	27.61
18	2	32.24	34.04	12.49	27.82	32.24	36.65
20	1	23.43	23.43				
10 or more	60	21.39	18.63	22.89	10.40	18.78	30.62

Table 7. The relative performance of private equity funds

Panel A reports excess IRRs between private and public equity investments for the 73 funds raised between 1981 and 1993. Excess IRRs are measured as the fund's IRR minus the return on a market index under five different assumptions regarding the public-market investment schedule. Panel B reports Profitability Index estimates for the 73 funds. The Profitability Index is a measure of the Limited Partner's return on invested capital, present-valued using various estimates of the cost of capital. It is computed as the ratio of the NPV of the fund's cash flows to the present value of the investment (i.e. the cash outflows or "draw downs"); see Brealey and Myers (1996), pp. 101f. All cash outflows are assumed riskfree and so discounted at the riskfree rate. Our estimates are thus conservative. Cash inflows are discounted at two different measures of the cost of capital: the annualized return on the S&P 500 or the Nasdaq Composite Index between the date the fund was raised and the date the fund matured. Thus, the holding period for the index investment matches the lifespan of the fund.

Public equity returns	Relative to S&P 500 (in %)					Relative to Nasdaq Composite (in %)				
	Mean	St.dev.	First quartile	Median	Third quartile	Mean	St.dev.	First quartile	Median	Third quartile
Panel A: Excess IRRs										
Excess IRRs = Fund IRR – R_m , where R_m is the IRR of ...										
... buy index at fund closing, sell at end of year 10	8.06	22.54	-2.34	6.04	15.01	6.28	23.01	-4.16	4.01	13.73
... invest \$1 according to the average fund's draw down schedule, hold till end of year 10	7.17	22.93	-2.51	4.74	15.17	4.78	23.74	-6.34	2.85	11.59
... invest \$1 according to the average fund's draw down schedule, sell at average index value in year 10	7.48	22.69	-2.23	5.63	16.03	5.39	23.11	-5.02	3.32	13.31
... invest \$1 according to the fund's actual draw down schedule, hold till end of year 10	5.71	22.67	-4.21	4.06	13.82	2.63	23.34	-7.24	0.92	9.69
... invest \$1 according to the fund's actual draw down schedule, sell at average index value in year 10	5.93	22.57	-4.49	3.94	14.23	2.62	23.16	-8.90	1.51	8.59
Panel B: Profitability Index	25.07	80.65	-23.65	12.18	48.03	9.96	79.39	-39.18	-9.93	38.37

Table 8. Performance by fund type

This table breaks down performance by type of fund: venture funds and non-venture funds, as identified in Table 1. Internal rates of return are calculated to the end of the fund's life (in Panel A) or the end of year T since the first closing, with $T = 1 \dots 20$ years (in Panel B). Among funds at or near maturity, non-venture funds perform substantially better. As Panel B illustrates, venture funds take about a year longer to break even, suggesting a longer commitment period and later exits.

	Venture funds						Non-venture funds					
	No. of funds	Internal rate of return (in percent)					No. of funds	Internal rate of return (in percent)				
		mean	st. dev.	first quartile	median	third quartile		mean	st. dev.	first quartile	median	third quartile
Panel A: Vintage year												
All years	*	1.44	90.54	-56.47	6.61	23.68	*	-19.19	47.65	-56.68	-8.02	17.96
1981-1992	15	11.45	26.73	10.06	17.49	24.08	44	23.53	20.18	12.60	19.24	28.76
1981-1993	19	14.08	26.88	6.61	17.49	26.91	54	21.83	20.33	9.97	19.00	27.99
Panel B: Years since first closing (T)												
1	2	-30.41	54.66	-49.74	-30.41	-11.08	16	-90.80	26.62	-99.90	-98.64	-94.94
2	7	-60.87	32.95	-77.39	-65.81	-55.29	51	-85.33	33.21	-99.35	-96.56	-86.60
3	13	-22.06	161.47	-95.93	-80.27	-23.34	76	-72.24	37.42	-97.78	-87.22	-48.63
4	25	-27.09	120.82	-85.20	-67.72	-17.73	96	-62.65	38.19	-94.12	-73.66	-41.16
5	28	-36.08	46.83	-81.19	-38.18	-6.81	87	-44.22	45.09	-85.36	-50.87	-13.27
6	27	-28.69	41.38	-55.51	-26.25	-5.22	80	-19.87	43.94	-53.34	-17.90	10.76
7	23	-13.18	37.07	-34.63	-13.16	8.15	76	-7.97	37.44	-24.46	-10.81	14.89
8	21	-9.72	35.37	-33.75	-3.64	7.98	62	6.81	29.33	-8.52	4.73	18.64
9	17	4.63	30.38	-8.19	12.44	28.06	51	16.05	23.93	2.87	10.06	24.46
10	15	5.41	29.19	-4.40	9.22	26.41	45	20.25	22.70	6.18	16.18	29.24
10 or more	15	12.56	27.59	10.06	17.49	26.91	45	24.34	20.61	10.60	19.31	31.96

Table 9. Portfolio risk and risk-adjusted performance

Panel A reports measures of portfolio risk for the 73 funds raised between 1981 and 1993. We assign portfolio companies to the 48 Fama-French (1997) industries (see the text for further information regarding our classification procedure). We report the fraction of the portfolio held in the dominant Fama-French industry by number of companies and invested capital. This is a measure of the degree of diversification/concentration of a fund's portfolio. In addition, we compute a Herfindahl index of investment weights by industry. A larger Herfindahl indicates a more concentrated portfolio. Using Fama and French's five-year CAPM estimates of industry betas and industry risk premia (that is, industry betas times the equity risk premium; see their Table 8), we estimate funds' portfolio betas and portfolio risk premia using invested capital as weights. Panel B reports risk-adjusted Profitability Index estimates for the 73 funds. The Profitability Index is computed as in Table 7, except that we now discount cash inflows at risk-adjusted returns. The first risk-adjusted return measure is defined as the riskfree rate plus the fund's weighted industry risk premium. The riskfree rate is the yield to maturity on ten-year treasuries computed in the month the fund was raised (i.e., there is one riskfree rate per fund). The other risk-adjusted discount rate is defined as the riskfree rate plus the fund's weighted average Fama-French beta times the difference between the index return and the riskfree rate, where the index return is the annualized return on the S&P 500 between the date the fund was raised and the date it matured.

	All funds (1981-1993)					Venture funds (N=19)					Non-venture funds (N=54)				
	Mean	St.dev.	First quartile	Median	Third quartile	Mean	St.dev.	First quartile	Median	Third quartile	Mean	St.dev.	First quartile	Median	Third quartile
Panel A: Portfolio risk															
No. of portfolio companies	21.6	21.5	9.0	14.0	26.0	37.3	33.0	12.0	32.0	48.0	16.1	11.7	9.0	13.0	23.0
% of companies in dominant industry	32.07	16.38	18.83	28.57	44.95	32.73	14.99	21.43	29.63	43.90	31.84	16.98	18.75	28.57	45.45
% of invested capital in dominant industry	39.32	20.46	25.62	34.16	47.11	36.94	14.71	23.95	35.99	47.71	40.16	22.19	26.02	32.25	46.11
Portfolio concentration (Herfindahl)	0.280	0.208	0.144	0.225	0.350	0.243	0.142	0.133	0.184	0.392	0.293	0.226	0.151	0.234	0.332
Weighted portfolio betas	1.09	0.10	1.05	1.10	1.14	1.12	0.06	1.09	1.12	1.16	1.08	0.11	1.04	1.09	1.13
Weighted FF industry risk premia (%)	5.55	0.52	5.38	5.63	5.82	5.70	0.31	5.56	5.73	5.93	5.50	0.57	5.32	5.58	5.78
Panel B: Risk-adjusted Profitability Index															
Ex ante cost of capital (%)	32.23	83.01	-13.83	17.68	52.26	28.08	80.84	-34.66	20.06	72.77	33.69	84.46	-8.91	16.32	52.26
Ex post cost of capital, S&P 500 (%)	24.00	84.72	-27.39	10.94	46.95	15.11	73.16	-34.98	-1.45	46.29	27.13	88.86	-27.39	12.02	55.05

Table 10. The determinants of private equity returns

The dependent variable is excess IRR, measured as the fund's IRR minus the return on the S&P 500 index (assuming \$1 is invested in the index according to the fund's actual draw down schedule and held till the end of year 10). The mean excess return measured this way is 5.71 percent (see Table 7). The results are not qualitatively different if we use any of the other excess IRR measures from Table 7, or if we use raw IRRs. The explanatory variables are as defined in Tables 4 and 9. The regression is estimated using OLS. Standard errors are shown in italics. We use ***, **, and * to denote significance at the one percent, five percent, and 10 percent level (two-sided), respectively.

	All mature funds (1981-1993)				Buyout funds
log real fund size (in \$m)	0.099** <i>0.039</i>	0.103*** <i>0.036</i>	0.106*** <i>0.036</i>	0.107*** <i>0.036</i>	0.012 <i>0.048</i>
real fund size (in \$bn)	-0.084** <i>0.042</i>	-0.089** <i>0.040</i>	-0.091** <i>0.040</i>	-0.090** <i>0.040</i>	-0.021 <i>0.043</i>
dummy=1 if first-time fund	0.024 <i>0.064</i>	0.022 <i>0.059</i>	0.017 <i>0.060</i>	0.015 <i>0.060</i>	-0.037 <i>0.065</i>
log real fund inflows in vintage year	-0.101** <i>0.049</i>	-0.093* <i>0.047</i>	-0.098** <i>0.044</i>	-0.100** <i>0.044</i>	-0.097** <i>0.046</i>
weighted portfolio beta	-0.167 <i>0.488</i>	-0.044 <i>0.266</i>	-0.033 <i>0.260</i>	-0.046 <i>0.260</i>	0.186 <i>0.273</i>
<i>Portfolio diversification:</i>					
log no. of portfolio companies	-0.007 <i>0.037</i>				
% of companies in dominant industry		0.002 <i>0.006</i>			
% of invested capital in dominant industry			0.065 <i>0.131</i>		
portfolio concentration (Herfindahl)				0.095 <i>0.128</i>	0.031 <i>0.132</i>
constant	0.685 <i>0.657</i>	0.430 <i>0.458</i>	0.437 <i>0.449</i>	0.463 <i>0.446</i>	0.750 <i>0.452</i>
Adjusted R^2	3.7 %	5.0 %	5.3 %	5.7 %	0.5 %
Number of observations	73	73	73	73	54

Table 11. Draw downs and distributions of recent funds

The table shows average cumulative draw downs and distributions for each year of a fund's life T , computed as in Table 3, for recent (1998-2001) and older (1981-1997) funds. Recent funds have a slower draw down schedule, both among venture and non-venture funds, but they return capital more rapidly, presumably due to the favorable exit climate in the late 1990s. Such accelerated distributions are particularly evident among venture funds.

Years since first closing (T)	All funds		Venture funds		Non-venture funds	
	1998-2001	1981-1997	1998-2001	1981-1997	1998-2001	1981-1997
Panel A: Average cumulative draw downs / committed capital, up to year T						
1	0.1477	0.1748	0.1600	0.2379	0.1434	0.1548
2	0.3263	0.3953	0.3356	0.4216	0.3233	0.3870
3	0.5218	0.5885	0.6563	0.6309	0.4906	0.5754
4	0.6295	0.7473	0.6847	0.8208	0.6172	0.7247
Panel B: Average cumulative distributions / total capital invested, up to year T						
1	0.0198	0.0017	0.0328	0.0043	0.0153	0.0008
2	0.0616	0.0195	0.0969	0.0230	0.0503	0.0184
3	0.2878	0.1111	0.6339	0.1065	0.2074	0.1125
4	0.5973	0.2836	1.3442	0.3510	0.4314	0.2629
Panel C: Average cumulative distributions / total capital committed, up to year T						
1	0.0014	0.0013	0.0021	0.0031	0.0012	0.0008
2	0.0257	0.0176	0.0375	0.0212	0.0220	0.0164
3	0.1891	0.1022	0.5273	0.1050	0.1106	0.1013
4	0.4311	0.2611	1.1281	0.3485	0.2762	0.2343