



The risk and return of venture capital[☆]

John H. Cochrane*

Graduate School of Business, University of Chicago, 5807 S. Woodlawn, Chicago, IL 60637, USA

Received 30 April 2003; received in revised form 14 July 2003; accepted 12 March 2004

Available online 14 October 2004

Abstract

This paper measures the mean, standard deviation, alpha, and beta of venture capital investments, using a maximum likelihood estimate that corrects for selection bias. The bias-corrected estimation neatly accounts for log returns. It reduces the estimate of the mean log return from 108% to 15%, and of the log market model intercept from 92% to -7% . The selection bias correction also dramatically attenuates high arithmetic average returns: it reduces the mean arithmetic return from 698% to 59%, and it reduces the arithmetic alpha from 462% to 32%. I confirm the robustness of the estimates in a variety of ways. I also find that the smallest Nasdaq stocks have similar large means, volatilities, and arithmetic alphas in this time period, confirming that the remaining puzzles are not special to venture capital. Published by Elsevier B.V.

JEL classification: G24

Keywords: Venture capital; Private equity; Selection bias

[☆]I am grateful to Susan Woodward, who suggested the idea of a selection-bias correction for venture capital returns, and who also made many useful comments and suggestions. I gratefully acknowledge the contribution of Shawn Blosser, who assembled the venture capital data. I thank many seminar participants and two anonymous referees for important comments and suggestions. I gratefully acknowledge research support from NSF grants administered by the NBER and from CRSP. Data, programs, and an appendix describing data procedures and algebra can be found at <http://gsbwww.uchicago.edu/fac/john.cochrane/research/Papers/>.

*Corresponding author.

E-mail address: john.cochrane@gsb.uchicago.edu (J.H. Cochrane).

1. Introduction

This paper measures the expected return, standard deviation, alpha, and beta of venture capital investments. Overcoming selection bias is the central hurdle in evaluating such investments, and it is the focus of this paper. We observe valuations only when a firm goes public, receives new financing, or is acquired. These events are more likely when the firm has experienced a good return. I overcome this bias with a maximum-likelihood estimate. I identify and measure the increasing probability of observing a return as value increases, the parameters of the underlying return distribution, and the point at which firms go out of business.

I base the analysis on measured returns from investment to IPO, acquisition, or additional financing. I do not attempt to fill in valuations at intermediate dates. I examine individual venture capital *projects*. Since venture funds often take 2–3% annual fees and 20–30% of profits at IPO, returns to investors in venture capital *funds* are often lower. Fund returns also reflect some diversification across projects.

The central question is whether venture capital investments behave the same way as publicly traded securities. Do venture capital investments yield larger risk-adjusted average returns than traded securities? In addition, which kind of traded securities do they resemble? How large are their betas, and how much residual risk do they carry?

One can cite many reasons why the risk and return of venture capital might differ from the risk and return of traded stocks, even holding constant their betas or characteristics such as industry, small size, and financial structure (leverage, book/market ratio, etc.). First, investors might require a higher average return to compensate for the illiquidity of private equity. Second, private equity is typically held in large chunks, so each investment might represent a sizeable fraction of the average investor's wealth. Finally, VC funds often provide a mentoring or monitoring role to the firm. They often sit on the board of directors, or have the right to appoint or fire managers. Compensation for these contributions could result in a higher measured financial return.

On the other hand, venture capital is a competitive business with relatively free (though not instantaneous; see Kaplan and Shoar, 2003) entry. Many venture capital firms and their large institutional investors can effectively diversify their portfolios. The special relationship, information, and monitoring stories that suggest a restricted supply of venture capital might be overblown. Private equity could be just like public equity.

I verify large and volatile returns *if* there is a new financing round, IPO, or acquisition, i.e., if we do not correct for selection bias. The average arithmetic return to IPO or acquisition is 698% with a standard deviation of 3,282%. The distribution is highly skewed; there are a few returns of thousands of percent, many more modest returns of “only” 100% or so, and a surprising number of losses. The skewed distribution is well described by a lognormal, but average log returns to IPO or acquisition still have a large 108% mean and 135% standard deviation. A CAPM estimate gives an arithmetic alpha of 462%; a market model in logs still gives an alpha of 92%.

The selection bias correction dramatically lowers these estimates, suggesting that venture capital investments are much more similar to traded securities than one would otherwise suspect. The estimated average log return is 15% per year, not 108%. A market model in logs gives a slope coefficient of 1.7 and a -7.1% , not $+92\%$, intercept. Mean arithmetic returns are 59%, not 698%. The arithmetic alpha is 32%, not 462%. The standard deviation of arithmetic returns is 107%, not 3,282%.

I also find that investments in later rounds are steadily less risky. Mean returns, alphas, and betas all decline steadily from first-round to fourth-round investments, while idiosyncratic variance remains the same. Later rounds are also more likely to go public.

Though much lower than their selection-biased counterparts, a 59% mean arithmetic return and 32% arithmetic alpha are still surprisingly large. Most anomalies papers quarrel over 1–2% per month. The large arithmetic returns result from the large idiosyncratic *volatility* of these individual firm returns, not from a large mean log return. If $\sigma = 1$ (100%), $e^{\mu+(1/2)\sigma^2}$ is large (65%), even if $\mu = 0$. Venture capital investments are like options; they have a small chance of a huge payoff.

One naturally distrusts the black-box nature of maximum likelihood, especially when it produces an anomalous result. For this reason I extensively check the facts behind the estimates. The estimates are driven by, and replicate, two central sets of stylized facts: the distribution of observed returns as a function of firm age, and the pattern of exits as a function of firm age. The distribution of total (not annualized) returns is quite stable across horizons. This finding contrasts strongly with the typical pattern that the total return distribution shifts to the right and spreads out over time as returns compound. A stable total return is, however, a signature pattern of a selected sample. When the winners are pulled off the top of the return distribution each period, that distribution does not grow with time. The exits (IPO, acquisition, new financing, failure) occur slowly as a function of firm age, essentially with geometric decay. This fact tells us that the underlying distribution of annual log returns must have a small mean and a large standard deviation. If the annual log return distribution had a large positive or negative mean, all firms would soon go public or fail as the mass of the total return distribution moves quickly to the left or right. Given a small mean log return, we need a large standard deviation so that the tails can generate successes and failures that grow slowly over time.

The identification is interesting. The pattern of *exits* with time, rather than the returns, drives the core finding of low mean log return and high return volatility. The distribution of *returns* over time then identifies the probability that a firm goes public or is acquired as a function of value. In addition, the high *volatility*, rather than a high mean return, drives the core finding of high average arithmetic returns.

Together, these facts suggest that the core findings of high arithmetic returns and alphas are robust. It is hard to imagine that the pattern of exits could be anything but the geometric decay we observe in this dataset, or that the returns of individual venture capital projects are not highly volatile, given that the returns of traded small

growth stocks are similarly volatile. I also test the hypotheses $\alpha = 0$ and $E(R) = 15\%$ and find them overwhelmingly rejected.

The estimates are not just an artifact of the late 1990s IPO boom. Ignoring all data past 1997 leads to qualitatively similar results. Treating all firms still alive at the end of the sample (June 2000) as out of business and worthless on that date also leads to qualitatively similar results. The results do not depend on the choice of reference return: I use the S&P500, the Nasdaq, the smallest Nasdaq decile, and a portfolio of tiny Nasdaq firms on the right-hand side of the market model, and all leave high, volatility-induced arithmetic alphas. The estimates are consistent across two basic return definitions, from investment to IPO or acquisition, and from one round of venture investment to the next. This consistency, despite quite different features of the two samples, gives credence to the underlying model. Since the round-to-round sample weights IPOs much less, this consistency also suggests there is no great return when the illiquidity or other special feature of venture capital is removed on IPO. The estimates are quite similar across industries; they are not just a feature of internet stocks. The estimates do not hinge on particular observations. The central estimates allow for measurement error, and the estimates are robust to various treatments of measurement error. Removing the measurement error process results in even greater estimates of mean returns. An analysis of influential data points finds that the estimates are not driven by the occasional huge successes, and also are not driven by the occasional financing round that doubles in value in two weeks.

For these reasons, the remaining average arithmetic returns and alphas are not easily dismissed. If venture capital seems a bit anomalous, perhaps similar traded stocks behave the same way. I find that a sample of very small Nasdaq stocks in this time period has similarly large mean arithmetic returns, large—over 100%—standard deviations, and large—53%!—arithmetic alphas. These alphas are statistically significant, and they are not explained by a conventional small-firm portfolio or by the Fama-French three-factor model. However, the beta of venture capital on these very small stocks is not one, and the alpha is not zero, so venture capital returns are not “explained” by these very small firm returns. They are similar phenomena, but not the same phenomenon.

Whatever the explanation—whether the large arithmetic alphas reflect the presence of an additional factor, whether they are a premium for illiquidity, etc.—the fact that we see a similar phenomenon in public and private markets suggests that there is little that is special about venture capital per se.

2. Literature

This paper’s distinctive contribution is to estimate the risk and return of venture capital *projects*, to correct seriously for selection bias, especially the biases induced by projects that remain private at the end of the sample, and to avoid imputed values.

Peng (2001) estimates a venture capital *index* from the same basic data I use, with a method-of-moments repeat sales regression to assign unobserved values and a

reweighting procedure to correct for the still-private firms at the end of the sample. He finds an average geometric return of 55%, much higher than the 15% I find for individual projects. He also finds a very high 4.66 beta on the Nasdaq index. Moskowitz and Vissing-Jorgenson (2002) find that a portfolio of all private equity has a mean and standard deviation of return close to those of the value-weighted index of traded stocks. However, they use self-reported valuations from the survey of consumer finances, and venture capital is less than 1% of all private equity, which includes privately held businesses and partnerships. Long (1999) estimates a standard deviation of 24.68% per year, based on the return to IPO of nine successful VC investments.

Bygrave and Timmons (1992) examine venture capital *funds*, and find an average internal rate of return of 13.5% for 1974–1989. The technique does not allow any risk calculations. Venture Economics (2000) reports a 25.2% five-year return and 18.7% ten-year return for all venture capital funds in their database as of 12/21/99, a period with much higher stock returns. This calculation uses year-end values reported by the funds themselves. Chen et al. (2002) examine the 148 venture capital funds in the Venture Economics data that had liquidated as of 1999. In these funds they find an annual arithmetic average return of 45%, an annual compound (log) average return of 13.4%, and a standard deviation of 115.6%, quite similar to my results. As a result of the large volatility, however, they calculate that one should only allocate 9% of a portfolio to venture capital. Reyes (1990) reports betas from 1.0 to 3.8 for venture capital as a whole, in a sample of 175 mature venture capital funds, but using no correction for selection or missing intermediate data. Kaplan and Schoar (2003) find that average fund returns are about the same as the S&P500 return. They find that fund returns are surprisingly persistent over time.

Gompers and Lerner (1997) measure risk and return by examining the investments of a single venture capital firm, periodically marking values to market. This sample includes failures, eliminating a large source of selection bias but leaving the survival of the venture firm itself and the valuation of its still-private investments. They find an arithmetic average annual return of 30.5% gross of fees from 1972–1997. Without marking to market, they find a beta of 1.08 on the market. Marking to market, they find a higher beta of 1.4 on the market, and 1.27 on the market with 0.75 on the small firm portfolio and 0.02 on the value portfolio in a Fama-French three-factor regression. Clearly, marking to market rather than using self-reported values has a large impact on risk measures. They do not report a standard deviation, though one can infer from $\beta = 1.4$ and $R^2 = 0.49$ a standard deviation of $1.4 \times 16 / \sqrt{0.49} = 32\%$. (This is for a fund, not the individual projects.) Gompers and Lerner find an intercept of 8% per year with either the one-factor or three-factor model. Ljungqvist and Richardson (2003) examine in detail all the venture fund investments of a single large institutional investor, and they find a 19.8% internal rate of return. They reduce the sample selection problem posed by projects still private at the end of the sample by focusing on investments made before 1992, almost all of which have resolved. Assigning betas, they recover a 5–6% premium, which they interpret as a premium for the illiquidity of venture capital investments.

Discount rates applied by VC investors might be informative, but the contrast between high discount rates applied by venture capital investors and lower ex post average returns is an enduring puzzle in the venture capital literature. Smith and Smith (2000) survey a large number of studies that report discount rates of 35% to 50%. However, this puzzle depends on the interpretation of “expected cash flows”. If “expected” means “what will happen if everything goes as planned”, it is much larger than a conditional mean, and a larger “discount rate” should be applied.

3. Overcoming selection bias

We observe a return only when the firm gets new financing or is acquired, but this fact need not bias our estimates. If the probability of observing a return were independent of the project’s value, simple averages would still correctly measure the underlying return characteristics. However, projects *are* more likely to get new financing, and especially to go public, when their value has risen. As a result, the mean returns to projects that get additional financing are an upward-biased estimate of the underlying mean return.

To understand the effects of selection, suppose that every project goes public when its value has grown by a factor of 10. Now, every measured return is exactly 1,000%, no matter what the underlying return distribution. A mean return of 1,000% and a zero standard deviation is obviously a wildly biased estimate of the returns facing an investor!

In this example, however, we can still identify the parameters of the underlying return distribution. The 1,000% measured returns tell us that the cutoff for going public is 1,000%. Observed returns tell us about the *selection function*, not the *return distribution*. The fraction of projects that go public at each age then identifies the return distribution. If we see that 10% of the projects go public in one year, then we know that the 10% upper tail of the return distribution begins at a 1,000% return. Since the mean grows with horizon and the standard deviation grows with the square root of horizon, the fractions that go public over time can separately identify the mean and the standard deviation (and, potentially, other moments) of the underlying return distribution.

In reality, the selection of projects to get new financing or be acquired is not a step function of value. Instead, the probability of obtaining new financing is a smoothly increasing function of the project’s value, as illustrated by $\Pr(\text{IPO}|\text{Value})$ in Fig. 1. The distribution of measured returns is then the product of the underlying return distribution and the rising selection probability. Measured returns still have an upward-biased mean and a downward-biased volatility. The calculations are more complex, but we can still identify the underlying return distribution and the selection function by watching the distribution of observed returns as well as the fraction of projects that obtain new financing over time.

I have nothing new to say about *why* projects are more likely to get new financing when value has increased, and I fit a convenient functional form rather than impose




Fig. 1. Generating the measured return distribution from the underlying return distribution and selection of projects to go public.

a particular economic model of this phenomenon. It's not surprising: good news about future productivity raises value and the need for new financing. The standard q theory of investment also predicts that firms will invest more when their values rise. (MacIntosh (1997, p. 295) discusses selection.) I also do not model the fact that more projects are *started* when market valuations are high, though the same motivations apply.

3.1. Maximum likelihood estimation

My objective is to estimate the mean, standard deviation, alpha, and beta of venture capital investments, correcting for the selection bias caused by the fact that we do not see returns for projects that remain private. To do this, I have to develop a model of the probability structure of the data—how the returns we see are generated from the underlying return process and the selection of projects that get new financing or go out of business. Then, for each possible value of the parameters, I can calculate the probability of seeing the data given those parameters.

The fundamental data unit is a financing round. Each round can have one of three basic fates. First, the firm can go public, be acquired, or get a new round of financing. These fates give us a new valuation, so we can measure a return. For this discussion, I lump all three fates together under the name “new financing round”. Second, the firm can go out of business. Third, the firm can remain private at the end of the sample. We need to calculate the probabilities of these three events, and the probability of the observed return if the firm gets new financing.

Fig. 2 illustrates how I calculate the likelihood function. I set up a grid for the log of the project’s value $\log(V_t)$ at each date t . I start each project at an initial value $V_0 = 1$, as shown in the top panel of Fig. 2. (I’m following the fate of a typical dollar invested.) I model the growth in value for subsequent periods as a lognormally distributed variable,

$$\ln\left(\frac{V_{t+1}}{V_t}\right) = \gamma + \ln R_t^f + \delta(\ln R_{t+1}^m - \ln R_t^f) + \varepsilon_{t+1}; \quad \varepsilon_{t+1} \sim N(0, \sigma^2). \quad (1)$$

I use a time interval of three months, balancing accuracy and simulation time. Eq. (1) is like the CAPM, but using log rather than arithmetic returns. Given the extreme skewness and volatility of venture capital investments, a statistical model with normally distributed arithmetic returns would be strikingly inappropriate. Below, I derive and report the implied market model for arithmetic returns (alpha and beta) from this linear lognormal statistical model. From Eq. (1), I generate the probability distribution of value at the beginning of period 1, $\Pr(V_1)$ as shown in the second panel of Fig. 2.




Fig. 2. Procedure for calculating the likelihood function.

Next, the firm could get a new financing round. The probability of getting a new round is an increasing function of value. I model this probability as a logistic function,

$$\Pr(\text{new round at } t|V_t) = 1/[1 + e^{-a(\ln(V_t)-b)}]. \tag{2}$$

This function rises smoothly from 0 to 1, as shown in the second panel of Fig. 2. Since I have started with a value of \$1, I assume here that selection to go public depends on total return achieved, not size per se. A \$1 investment that grows to \$1,000 is likely to go public, where a \$10,000 investment that falls to \$1,000 is not. Now I can find the probability that the firm gets a new round with value V_t ,

$$\Pr(\text{new round at } t, \text{value } V_t) = \Pr(V_t) \times \Pr(\text{new round at } t|V_t).$$

This probability is shown by the bars on the right-hand side of the second panel of Fig. 2. These firms exit the calculation of subsequent probabilities.

Next, the firm can go out of business. This is more likely for low values. I model $\Pr(\text{out of business at } t|V_t)$ as a declining linear function of value V_t , starting from the lowest value gridpoint and ending at an upper bound k , as shown by $\Pr(\text{out}|value)$ on the left side of the second panel of Fig. 2. A lognormal process such as (1) never reaches a value of zero, so we must envision something like k if we are to generate a finite probability of going out of business.¹ Multiplying, we obtain the probability that the firm goes out of business in period 1,

$$\begin{aligned} &\Pr(\text{out of business at } t, \text{value } V_t) \\ &= \Pr(V_t) \times [1 - \Pr(\text{new round at } t|V_t)] \times \Pr(\text{out of business at } t|V_t). \end{aligned}$$

These probabilities are shown by the bars on the left side of the second panel of Fig. 2.

Next, I calculate the probability that the firm remains private at the end of period 1. These are just the firms that are left over,

$$\begin{aligned} &\Pr(\text{private at end of } t, \text{value } V_t) \\ &= \Pr(V_t) \times [1 - \Pr(\text{new round}|V_t)] \times [1 - \Pr(\text{out of business}|V_t)]. \end{aligned}$$

This probability is indicated by the bars in the third panel of Fig. 2.

Next, I again apply (1) to find the probability that the firm enters the second period with value V_2 , shown in the bottom panel of Fig. 2,

$$\Pr(V_{t+1}) = \sum_{V_t} \Pr(V_{t+1}|V_t) \Pr(\text{private at end of } t, V_t). \tag{3}$$

$\Pr(V_{t+1}|V_t)$ is given by the lognormal distribution of (1). As before, I find the probability of a new round in period 2, the probability of going out of business in

¹The working paper version of this article used a simpler specification that the firm went out of business if V fell below k . Unfortunately, this specification leads to numerical problems, since the likelihood function changes discontinuously as the parameter k passes through a value gridpoint. The linear probability model is more realistic, and results in a better-behaved continuous likelihood function. A smooth function like the logistic new financing selection function would be prettier, but this specification requires only one parameter, and the computational cost of extra parameters is high.

period 2, and the probability of remaining private at the end of period 2. All of these are shown in the bottom panel of Fig. 2. This procedure continues until we reach the end of the sample.

3.2. Accounting for data errors

Many data points have bad or missing dates or returns. Each round results in one of the following categories: (1) new financing with good date and good return data, (2) new financing with good dates but bad return data, (3) new financing with bad dates and bad return data, (4) still private at end of sample, (5) out of business with good exit date, (6) out of business with bad exit date.

To assign the probability of a type 1 event, a new round with a good date and good return data, I first find the fraction d of all rounds with new financing that have good date and return data. Then, the probability of seeing this event is d times the probability of a new round at age t with value V_t ,

$$\begin{aligned} & \text{Pr}(\text{new financing at age } t, \text{ value } V_t, \text{ good data}) \\ &= d \times \text{Pr}(\text{new financing at } t, \text{ value } V_t). \end{aligned} \quad (4)$$

I assume here that seeing good data is independent of value.

A few projects with “normal” returns in a very short time have astronomical annualized returns. Are these few data points driving the results? One outlier observation with probability near zero can have a huge impact on maximum likelihood. As a simple way to account for such outliers, I consider a uniformly distributed measurement error. With probability $1 - \pi$, the data record the true value. With probability π , the data erroneously record a value uniformly distributed over the value grid. I modify Eq. (4) to

$$\begin{aligned} & \text{Pr}(\text{new financing at age } t, \text{ value } V_t, \text{ good data}) \\ &= d \times (1 - \pi) \times \text{Pr}(\text{new financing at } t, \text{ value } V_t) \\ &+ d \times \pi \times \frac{1}{\#\text{gridpoints}} \sum_{V_t} \text{Pr}(\text{new financing at } t, \text{ value } V_t). \end{aligned}$$

This modification fattens up the tails of the measured value distribution. It allows a small number of observations to get a huge positive or negative return by measurement error rather than force a huge mean or variance of the return distribution to accommodate a few extreme annualized returns.

A type 2 event, new financing with good dates but bad return data, is still informative. We know how long it takes this investment round to build up the kind of value that typically leads to new financing. To calculate the probability of a type 2 event, I sum across the vertical bars on the right side of the second panel of Fig. 2,

$$\begin{aligned} & \text{Pr}(\text{new financing at age } t, \text{ no return data}) \\ &= (1 - d) \times \sum_{V_t} \text{Pr}(\text{new financing at } t, \text{ value } V_t). \end{aligned}$$

A type 3 event, new financing with bad dates and bad return data, tells us that at some point this project was good enough to get new financing, though we know only that it happened between the start of the project and the end of the sample. To calculate the probability of this event, I sum over time from the initial round date to the end of the sample as well,

$$\begin{aligned} & \text{Pr}(\text{new financing, no date or return data}) \\ &= (1 - d) \times \sum_t \sum_{V_t} \text{Pr}(\text{new financing at } t, \text{value } V_t). \end{aligned}$$

To find the probability of a type 4 event, still private at the end of the sample, I simply sum across values at the appropriate age

$$\begin{aligned} & \text{Pr}(\text{still private at end of sample}) \\ &= \sum_{V_t} \text{Pr}(\text{still private at } t = (\text{end of sample}) - (\text{start date}), V_t). \end{aligned}$$

. Type 5 and 6 events, out of business, tell us about the *lower* tail of the return distribution. Some of the out of business observations have dates, and some do not. Even when there is apparently good date data, a large fraction of the out-of-business observations occur on two specific dates. Apparently, there were periodic data cleanups of out-of-business observations prior to 1997. Therefore, when there is an out-of-business date, I interpret it as “this firm went out of business on or before date t ”, summing up the probabilities of younger out-of-business events, rather than “on date t ”. This assignment affects the results: since one of the cleanup dates comes on the heels of a large positive stock return, using the dates as they are leads to negative beta estimates. To account for missing date data in out-of-business firms, I calculate the fraction of all out-of-business rounds with good date data c . Thus, I calculate the probability of a type 5 event, out of business with good date information, as

$$\begin{aligned} & \text{Pr}(\text{out of business on or before age } t, \text{date data}) \\ &= c \times \sum_{t=1}^t \sum_{V_t} \text{Pr}(\text{out of business at } t, V_t). \end{aligned} \tag{5}$$

Finally, if the date data are bad, all we know is that this round went out of business at some point before the end of the sample. I calculate the probability of a type 6 event as

$$\begin{aligned} & \text{Pr}(\text{out of business, no date data}) \\ &= (1 - c) \times \sum_{t=1}^{\text{end}} \sum_{V_t} \text{Pr}(\text{out of business at } t, V_t). \end{aligned}$$

Based on the above structure, for given parameters $\{\gamma, \delta, \sigma, k, a, b, \pi\}$, I can compute the probability that we see any data point. Taking the log and adding up over all data points, I obtain the log likelihood. I search numerically over values $\{\gamma, \delta, \sigma, k, a, b, \pi\}$ to maximize the likelihood function.

Of course, the ability to separately identify the probability of going public and the parameters of the return process requires some assumptions. Most important, I assume that the selection function $\Pr(\text{new round} | V_t)$ is the same for firms of all ages t . If the initial value doubles in a month, we are just as likely to get a new round as if it takes ten years to double the initial value. This is surely unrealistic at very short and very long time periods. I also assume that the return process is i.i.d. One might specify that value creation starts slowly and then accelerates, or that betas or volatilities change with size. However, identifying these tendencies without much more data will be tenuous.

4. Data

I use the VentureOne database from its beginning in 1987 to June 2000. The dataset consists of 16,613 financing rounds, with 7,765 companies and a total of \$112,613 million raised. VentureOne claims to have captured approximately 98% of financing rounds, mitigating survival bias of projects and funds. However, the VentureOne data are not completely free of survival bias. VentureOne records a financing round if it includes at least one venture capital firm with \$20 million or more in assets under management. Having found a qualifying round, they search for previous rounds. Gompers and Lerner (2000, 288pp.) discuss this and other potential selection biases in the database. Kaplan et al. (2002) compare the VentureOne data to a sample of 143 VC financings on which they have detailed information. They find as many as 15% of rounds omitted. They find that post-money values of a financing round, though not the fact of the round, are more likely to be reported if the company subsequently goes public. This selection problem does not bias my estimates.

The VentureOne data do not include the financial results of a public offering, merger, or acquisition. To compute such values, we use the SDC Platinum Corporate New Issues and Mergers and Acquisitions (M&A) databases, MarketGuide, and other online resources.² We calculate returns to IPO using offering prices. There is usually a lockup period between IPO and the time that venture capital investors can sell shares, and there is an active literature studying IPO mispricing, post-IPO drift and lockup-expiration effects, so one might want to study returns to the end of the first day of trading, or even include six months or more of market returns. However, my objective is to measure venture capital returns, not to contribute to the large literature that studies returns to newly listed firms. For this purpose, it seems wisest to draw the line at the offering price. For example, suppose that I include first-day returns, and that this inclusion substantially raises the resulting mean returns and alphas. Would we call that the “risk and return of venture capital” or “IPO mispricing”? Clearly the latter, so I stop at offering prices to focus on the former. In addition, all of these new-listing effects are small compared to the returns (and errors) in the venture capital data. Even a 10% error in

²“We” here includes Shawn Blosser, who assembled the venture capital data.

final value would have little effect on my results, since it is spread over the many years of a typical VC investment. A 10% error is only four days of volatility at the estimated nearly 100% standard deviation of return.³

The basic data consist of the date of each investment, dollar amount invested, value of the firm after each investment, and characteristics including industry and location. VentureOne also notes whether the company has gone public, been acquired, or gone out of business, and the date of these events. We infer returns by tracking the value of the firm after each investment. For example, suppose firm XYZ has a first round that raises \$10 million, after which the firm is valued at \$20 million. We infer that the VC investors own half of the stock. If the firm later goes public, raising \$50 million and valued at \$100 million after IPO, we infer that the VC investors' portion of the firm is now worth \$25 million. We then infer their gross return at $\$25\text{M}/\$10\text{M} = 250\%$. We use the same method to assess dilution of initial investors' claims in multiple rounds.

The biggest potential error of this procedure is that if VentureOne misses intermediate rounds, the extra investment is credited as a return to the original investors. For example, the edition of VentureOne I used to construct the data missed all but the seed round of Yahoo, resulting in a return even more enormous than reality. I run the data through several filters⁴ and I add the measurement error process π to try to account for this kind of error.

Venture capitalists typically obtain convertible preferred rather than common stock. (See Kaplan and Strömberg (2003). Admati and Pfleiderer (1994) have a nice summary of venture capital arrangements, especially mechanisms designed to insure that valuations are “arm's length”.) These arrangements are not noted in the VentureOne data, so I treat all investments as common stock. This approximation is not likely to introduce a large bias. The results are driven by the successes, not by liquidation values in the surprisingly rare failures, or in acquisitions that produce losses for common stock investors, where convertible preferred holders can retrieve their capital. In addition, the bias will be to understate estimated VC returns, while

³The unusually large first-day returns in 1999 and 2000 are a possible exception. For example, Ljungqvist and Wilhelm (2003, Table, II) report mean first-day returns for 1996–2000 of 17%, 14%, 23%, 73%, and 58%, with medians of 10%, 9%, 10%, 39%, and 30%. However, these are reported as transitory anomalies, not returns expected when the projects are started. We should be uncomfortable adding a 73% *expected* one-day return to our view of the venture capital value creation process. Also, I find below quite similar results in the pre-1997 sample, which avoids this anomalous period. See also Lee and Wahal (2002), who find that VC-backed firms have larger first-day returns than other firms.

⁴Starting with 16,852 observations in the base case of the IPO/acquisition sample (numbers vary for subsamples), I eliminate 99 observations with more than 100% or less than 0% inferred shareholder value, and I eliminate 107 investments in the last period, the second quarter of 2000, since the model can't say anything until at least one period has passed. In 25 observations, the exit date comes before the VC round date, so I treat the exit date as missing.

For the maximum likelihood estimation, I treat 37 IPO, acquisition, or new rounds with zero returns as out of business (0 blows up a lognormal), and I delete four observations with anomalously high returns (over 30,000%) after I hand-checking them and finding that they were errors due to missing intermediate rounds. I similarly deleted four observations with a log annualized return greater than 15 ($100 \times (e^{15} - 1) = 3.269 \times 10^8\%$) on the strong suspicion of measurement error in the dates. All of these observations are included in the data characterization, however. I am left with 16,638 data points.

the puzzle is that the estimated returns are so high. Gilson and Schizer (2003) argue that the practice of issuing convertible preferred stock to VC investors is not driven by cash flow or control considerations, but by tax law. Management is typically awarded common shares at the same time as the venture financing round. Distinguishing the classes of shares allows managers to underreport the value of their share grants, taxable immediately at ordinary income rates, and thus to report this value as a capital gain later on. If so, then the distinction between common and convertible preferred shares makes even less of a difference for my analysis.

I model the return to equity directly, so the fact that debt data are unavailable does not generate an accounting mistake in calculating returns. Firms with different levels of debt can have different betas, however, which I do not capture.

4.1. IPO/acquisition and round-to-round samples

The basic data unit is a financing round. If a financing round is followed by another round, if the firm is acquired, or if the firm goes public, we can calculate a return. I consider two basic sample definitions for these returns. In the “round-to-round” sample, I measure every return from a financing round to a subsequent financing round, IPO, or acquisition. Thus, if a firm has two financing rounds and then goes public, I measure two returns, from round 1 to round 2, and from round 2 to IPO. If the firm has two rounds and then fails, I measure a positive return from round 1 to round 2 but then a failure from round 2. If the firm has two rounds and remains private, I measure a return from round 1 to round 2, but round 2 is coded as remaining private.

One might be suspicious of returns constructed from such round-to-round valuations. A new round determines the terms at which new investors come *in* but almost never the terms at which old investors can get *out*. The returns to investors are really the returns to acquisition or IPO only, ignoring intermediate financing rounds. In addition, an important reason to study venture capital is to examine whether venture capital investments have low prices and high returns due to their illiquidity. We can only hope to see this fact in returns from investment to IPO, not in returns from one round of venture investment to another, since the latter returns retain the illiquid character of venture capital investments. More basically, it is interesting to characterize the *eventual* fate of venture capital investments as well as the returns measured in successive financing rounds.

For all these reasons, I emphasize a second basic data sample, denoted “IPO/acquisition” below. If a firm has two rounds and then goes public, I measure two returns, round 1 to IPO, and round 2 to IPO. If the firm has two rounds and then fails, I measure two failures, round 1 to failure and round 2 to failure. If it has two rounds and remains private, both rounds are coded as remaining private with no measured returns. In addition to its direct interest, we can look for signs of an illiquidity or other premium by contrasting these round-to-IPO returns with the above round-to-round returns. Different rounds of the same company overlap in time, of course, and I deal with the econometric issues raised by this overlap below.

Table 1
The fate of venture capital investments

Fate	IPO/acquisition			Round to round		
	Return	No return	Total	Return	No return	Total
IPO	16.1	5.3	21.4	5.9	2.0	7.9
Acquisition	5.8	14.6	20.4	2.9	6.3	9.2
Out of business	9.0		9.0	4.2		4.2
Remains private		45.5	45.5		23.3	23.3
IPO registered		3.7	3.7		1.2	1.2
New round				28.3	25.9	54.2

Note: Table entries are the percentage of venture capital financing rounds with the indicated fates. The IPO/acquisition sample tracks each investment to its final fate. The round-to-round sample tracks each investment to its next financing round only. “Return” indicates rounds for which we can measure a return, “No reference” indicates rounds in the given category (e.g., IPO) but for which data are missing and we cannot calculate a return.

Table 1 characterizes the fates of venture capital investments. We see that 21.4% of rounds eventually result in an IPO and 20.4% eventually result in acquisition. Unfortunately, I am able to assign a return to only about three quarters of the IPO and one quarter of the acquisitions. We see that 45.5% remain private, 3.7% have registered for but not completed an IPO, and 9% go out of business. There are surprisingly few failures. Moskowitz and Vissing-Jorgenson (2002) find that only 34% of their sample of private equity survive ten years. However, many firms go public at valuations that give losses to VC investors, and many more are acquired on such terms. (Weighting by dollars invested yields quite similar numbers, so I lump investments together without size effects in the estimation.)

I measure far more returns in the round-to-round sample. The average company has 2.1 venture capital financing rounds (16,638 rounds/7,765 companies), so the fractions that end in IPO, acquisition, out of business, or still private are cut in half, while 54.2% get a new round, about half of which result in return data. The smaller number that remain private means less selection bias to control for, and less worry that some of the still-private firms are “living dead”, really out of business.

5. Results

Table 2 presents characteristics of the subsamples. Table 3 presents parameter estimates for the IPO/acquisition sample, and Table 4 presents estimates for the round-to-round sample. Table 5 presents asymptotic standard errors.

5.1. Base case results

The base case is the “All” sample in Table 3. The mean log return in Table 3 is a sensible 15%, just about the same as the 15.9% mean log S&P500 return in this

Table 2
Characteristics of the samples

	Rounds					Industries				Subsamples	
	All	1	2	3	4	Health	Info	Retail	Other	Pre 97	Dead 00
<i>IPO/acquisition sample</i>											
Number	16,638	7,668	4,474	2,453	1,234	3,915	9,190	3,091	442	5,932	16,638
Out of bus.	9	9	9	9	9	9	10	7	12	5	58
IPO	21	17	21	26	31	27	21	15	22	33	21
Acquired	20	20	21	21	19	18	25	10	29	26	20
Private	49	54	49	43	41	46	45	68	38	36	0
<i>c</i>	95	93	97	98	96	96	94	96	94	75	99
<i>d</i>	48	38	49	57	62	51	49	38	26	48	52
<i>Round-to-round sample</i>											
Number	16,633	7,667	4,471	2,453	1,234	3,912	9,188	3,091	442	6,764	16,633
Out of bus.	4	4	4	5	5	4	4	4	7	2	29
IPO	8	5	7	11	18	9	8	7	10	12	8
Acquired	9	8	9	11	11	8	11	5	13	11	9
New round	54	59	55	50	41	59	55	45	52	69	54
Private	25	25	25	23	25	20	22	39	18	7	0
<i>c</i>	93	88	96	99	98	94	93	94	90	67	99
<i>d</i>	51	42	55	61	66	55	52	41	39	54	52

Note: All entries except Number are percentages. *c* = percent of out of business with good data. *d* = percent of new financing or acquisition with good data. Private are firms still private at the end of the sample, including firms that have registered for but not completed an IPO.

period. (I report average returns, alphas and standard deviations as annualized percentages, by multiplying averages and alphas by 400 and multiplying standard deviations by 200.) The standard deviation of log return is 89%, much larger than the 14.9% standard deviation of the log S&P500 return in this period. These are individual firms, so we expect them to be quite volatile compared to a diversified portfolio such as the S&P500. The 89% annualized standard deviation might be easier to digest as $89/\sqrt{365} = 4.7\%$ daily standard deviation, which is typical of very small growth stocks.

The intercept γ is negative at -7.1% . The slope δ is sensible at 1.7; venture capital is riskier than the S&P500. The residual standard deviation σ is large at 86%. The volatility of returns comes from idiosyncratic volatility, not from a large slope coefficient. The implied regression R^2 is a very small 0.075. ($1.7^2 \times 14.9^2 / (1.7^2 \times 14.9^2 + 89^2) = 0.075$.) Systematic risk is a small component of the risk of an individual venture capital investment.

(I estimate the parameters γ, δ, σ directly. I calculate $E \ln R$ and $\sigma \ln R$ in the first two columns using the mean 1987–2000 Treasury bill return of 6.8%, and the S&P500 mean and standard deviation of 15.9% and 14.9%, e.g., $E \ln R = -7.1 + 6.8 + 1.7 \times (15.9 - 6.8) = 15\%$. I present mean log returns first in Tables 3 and 4, as the mean is better estimated, more stable, and more comparable across specifications than is its decomposition into an intercept and a slope.)

Table 3
Parameter estimates in the IPO/acquisition sample

	$E \ln R$	$\sigma \ln R$	γ	δ	σ	ER	σR	α	β	k	a	b	π
All, baseline	15	89	-7.1	1.7	86	59	107	32	1.9	25	1.0	3.8	9.6
Asymptotic σ			0.7	0.04	0.6					0.02	0.02	0.06	0.6
Bootstrap σ	2.4	6.7	1.7	0.4	7.0	5.9	11	9.4	0.4	3.6	0.08	0.28	1.9
Nasdaq	14	97	-7.7	1.2	93	66	121	39	1.4	20	0.7	5.0	5.7
Nasdaq Dec1	17	96	-0.3	0.9	92	69	119	45	1.0	22	0.7	5.4	6.3
Nasdaq <\$2M	8.2	103	-27	0.5	100	67	129	22	0.5	14	0.7	5.0	4.1
No δ	11	105				72	134			11	0.8	4.3	4.2
Round 1	19	96	3.7	1.0	95	71	120	53	1.1	17	1.0	4.2	8.0
Round 2	12	98	-1.6	0.8	97	65	120	49	0.9	16	1.0	3.6	5.0
Round 3	8.0	98	-4.4	0.6	98	60	120	46	0.7	17	0.8	3.9	2.9
Round 4	-0.8	99	-12	0.5	99	51	119	39	0.5	13	1.1	2.5	5.5
Health	17	67	8.7	0.2	67	42	76	33	0.2	36	0.7	5.1	7.8
Info	15	108	-5.2	1.4	105	79	139	55	1.7	14	0.8	4.3	4.3
Retail	17	127	11	-0.1	127	111	181	106	-0.1	11	0.4	10.0	2.9
Other	25	62	13	0.6	61	46	71	33	0.6	53	0.4	10.0	13

Note: Returns are calculated from venture capital financing round to eventual IPO, acquisition, or failure, ignoring intermediate venture financing rounds.

Columns: $E \ln R, \sigma \ln R$ are the parameters of the underlying lognormal return process. All means, standard deviations and alphas are reported as annualized percentages, e.g., $400 \times E \ln R, 200 \times \sigma \ln R, 400 \times (ER - 1)$, etc. γ, δ , and σ are the parameters of the market model in logs, $\ln\left(\frac{V_{t+1}}{V_t}\right) = \gamma + \ln R_t^f + \delta(\ln R_{t+1}^m - \ln R_t^f) + \varepsilon_{t+1}$; $\varepsilon_{t+1} \sim N(0, \sigma^2)$. $E \ln R, \sigma \ln R$ are calculated from γ, δ, σ using the sample mean and variance of the three-month T-bill rate R^f and S&P500 return R^m , $E \ln R = \gamma + E \ln R_t^f + \delta(E \ln R_t^m - E \ln R_t^f)$ and $\sigma^2 \ln R = \delta^2 \sigma^2 (\ln R_t^m)^2 + \sigma^2$. $ER, \sigma R$ are average arithmetic returns $ER = e^{E \ln R + \frac{1}{2} \sigma^2 \ln R}$, $\sigma R = ER \times \sqrt{e^{\sigma^2 \ln R} - 1}$. α and β are implied parameters of the discrete time regression model in levels, $V_{t+1}^i/V_t^i = \alpha + R_t^f + \beta(R_{t+1}^m - R_t^f) + v_{t+1}^i$. k, a, b are estimated parameters of the selection function. k is point at which firms start to go out of business, expressed as a percentage of initial value. a, b govern the selection function $\Pr(\text{IPO, acq. at } t|V_t) = 1/(1 + e^{-a(\ln(V_t) - b)})$. Given that an IPO/acquisition occurs, there is a probability π that a uniformly distributed value is recorded instead of the correct value. *Rows:* “All” includes all financing rounds. Asymptotic standard errors are based on second derivatives of the likelihood function. Bootstrap standard errors are based on 20 replications of the estimate, choosing the sample randomly with replacement. “Nasdaq”, “Nasdaq Dec1”, “Nasdaq <\$2M”, and “No δ ” use the indicated reference returns in place of the S&P500. “Round i” considers only investments in financing round i. “Health, Info, Retail, Other” are industry classifications.

The asymptotic standard errors in the second row of Table 3 indicate that all these numbers are measured with great statistical precision. The bootstrap standard errors in the third row are a good deal larger than asymptotic standard errors, but still show the parameters to be quite well estimated. The bootstrap standard errors are larger in part because there are a small number of outlier data points with very large likelihoods. Their inclusion or exclusion has a larger effect on the results than the asymptotic distribution theory suggests. The asymptotic standard errors also ignore

Table 4
Parameter estimates in the round-to-round sample

	E ln R	σ ln R	γ	δ	σ	ER	σR	α	β	k	a	b	π
All, baseline	20	84	7.6	0.6	84	59	100	45	0.6	21	1.7	1.3	4.7
Asymptotic σ			1.1	0.1	0.8					0.4	0.02	0.02	0.4
Bootstrap σ	1.1	7.2	4.7	0.5	6.4	7.5	11	5.7	0.5	3.8	0.2	0.3	0.8
Nasdaq	15	91	-4.9	1.1	87	61	110	35	1.2	18	1.5	1.5	3.4
Nasdaq Dec1	22	90	7.3	0.7	88	68	112	49	0.7	24	0.6	3.3	2.5
Nasdaq < \$2M	16	91	-4.5	0.2	90	62	111	37	0.2	16	1.6	1.4	3.5
No δ	21	85				61	102			20	1.6	1.4	4.2
Round 1	26	90	11	0.8	89	72	112	55	1.0	16	1.9	1.3	4.3
Round 2	20	83	7.5	0.6	82	58	99	44	0.7	22	1.6	1.4	3.6
Round 3	15	77	3.6	0.5	77	47	89	35	0.5	29	1.4	1.4	4.6
Round 4	8.8	84	0.1	0.2	83	46	97	37	0.2	21	1.3	1.4	3.7
Health	24	62	15	0.3	62	46	70	36	0.3	48	0.3	7.6	4.6
Info	23	95	12	0.5	94	74	119	62	0.5	19	0.7	2.9	2.2
Retail	25	121	11	0.7	121	111	171	96	0.8	14	0.5	4.1	0.5
Other	8.0	64	-3.9	0.6	63	29	70	16	0.6	35	0.5	5.2	3.6

Note: Returns are calculated from venture capital financing round to the next event: new financing, IPO, acquisition, or failure. See the note to Table 3 for row and column headings.

cross-correlation between individual venture capital returns, since I do not specify a cross-correlation structure in the data-generating model (1).

So far, the estimates look reasonable. If anything, the negative intercept is surprisingly low. However, the CAPM and most asset pricing and portfolio theory specifies arithmetic, not logarithmic, returns. Portfolios are linear in arithmetic, not log, returns, so diversification applies to arithmetic returns. The columns ER, σR , α , and β of Table 3 calculate implied characteristics of arithmetic returns.⁵ The mean

⁵We want to find the model in levels implied by Eq. (1), i.e.

$$\frac{V_{t+1}^i}{V_t^i} - R_t^f = \alpha + \beta(R_{t+1}^m - R_t^f) + v_{t+1}^i.$$

I find β from $\beta = \text{cov}(R, R^m) / \text{var}(R^m)$, and then α from $\alpha = E(R) - R^f - \beta[E(R^m) - R^f]$. The formulas are

$$\beta = e^{\gamma + (\delta - 1)(E(\ln R^m) - \ln R^f) + \sigma^2 / 2 + (\sigma^2 - 1)\sigma_m^2 / 2} \frac{(e^{\delta\sigma_m^2} - 1)}{(e^{\sigma_m^2} - 1)}, \tag{6}$$

$$\alpha = e^{\ln(R^f)} \{ (e^{\gamma + \delta(E(\ln R^m) - \ln R^f) + \delta^2\sigma_m^2 / 2 + \sigma^2 / 2} - 1) - \beta(e^{\mu_m - \ln R^f} + \sigma_m^2 / 2 - 1) \}, \tag{7}$$

where $\sigma_m^2 \equiv \sigma^2(\ln R^m)$. The continuous time limit is simpler, $\beta = \delta$, $\sigma(\epsilon) = \sigma(v)$, and

$$\alpha = \gamma + \frac{1}{2} \delta(\delta - 1)\sigma_m^2 + \frac{1}{2} \sigma^2.$$

I present the discrete time computations in the tables; the continuous time results are quite similar.

Table 5
Asymptotic standard errors for Tables 3 and 4

	IPO/acquisition (Table 3)							Round to round (Table 4)						
	γ	δ	σ	k	a	b	π	γ	δ	σ	k	a	b	π
All, baseline	0.7	0.04	0.6	0.02	0.02	0.06	0.6	1.1	0.08	0.8	0.4	0.02	0.02	0.4
Bootstrap	1.7	0.37	7.0	3.57	0.08	0.28	1.9	4.7	0.46	6.4	3.8	0.20	0.33	0.8
Nasdaq	1.0	0.05	1.1	0.81	0.02	0.13	0.5	0.7	0.01	0.9	0.4	0.03	0.03	0.4
Nasdaq Decl	1.0	0.04	1.1	0.62	0.02	0.15	0.6	1.1	0.04	1.1	1.0	0.01	0.02	0.4
Nasdaq <\$2M	1.7	0.02	0.8	0.35	0.01	0.08	0.5	1.2	0.01	0.5	0.2	0.03	0.02	0.3
No δ	0.7		1.0	0.15	0.02	0.11	0.6	0.7		0.8	0.6	0.03	0.03	0.3
Round 1	1.2	0.05	1.8	0.94	0.04	0.11	1.0	1.4	0.09	1.3	0.7	0.06	0.03	0.5
Round 2	2.4	0.20	2.3	1.18	0.06	0.16	1.2	2.3	0.16	1.8	1.4	0.07	0.05	0.8
Round 3	3.2	0.24	2.5	1.16	0.08	0.31	1.1	2.4	0.16	1.7	1.3	0.08	0.08	0.9
Round 4	5.3	0.38	3.3	1.57	0.08	0.20	1.8	4.0	0.26	2.9	2.0	0.12	0.14	1.1
Health	1.7	0.14	1.4	2.06	0.03	0.19	1.2	1.9	0.15	1.5	2.2	0.01	0.20	0.8
Info	1.7	0.13	1.6	0.69	0.01	0.06	0.8	1.7	0.13	1.3	0.8	0.01	0.04	0.4
Retail	1.9	0.08	3.5	1.27	0.00	0.00	1.2	5.5	0.30	3.8	1.7	0.02	0.10	0.5
Other	3.5	0.26	4.6	7.52	0.01	0.14	4.6	6.8	0.46	5.0	6.1	0.10	1.07	4.0

arithmetic return ER in Table 3 is a whopping 59%, with a 107% standard deviation. Even the 1.9 arithmetic β and the large S&P500 return in this period do not generate a return that high, leaving a 32% arithmetic α .

The large mean arithmetic returns and alphas result from the *volatility* rather than the *mean* of the underlying lognormal return distribution. The mean arithmetic return is $E(R) = e^{E \ln R + (1/2) \sigma^2 \ln R}$. With $\sigma^2 \ln R$ on the order of 100%, usually negligible $\frac{1}{2} \sigma^2$ terms generate 50% per year arithmetic returns by themselves. Venture capital investments are like call options; their arithmetic mean return depends on the mass in the right tail, which is driven by volatility more than by drift. I examine the high arithmetic returns and alphas in great detail below.

The out-of-business cutoff parameter k is 25%, meaning that the chance of going out of business rises to $\frac{1}{2}$ at 12.5% of initial value. This is a low number, but reasonable. Venture capital investors are likely to hang in there and wait for the final payout despite substantial intermediate losses.

The parameters a and b control the selection function. b is the point at which there is a 50% probability of going public or being acquired per quarter, and it occurs at a substantial 380% log return. Finally, the measurement error parameter π is about 10% and statistically significant. The estimation accounts for a small number of large positive and negative returns as measurement error rather than treat them as extreme values of a lognormal process.

The round-to-round sample in Table 4 gives quite similar results. The average log return is slightly higher, 20% rather than 15%, with quite similar volatility, 84% rather than 89%. The average log return splits in to a lower slope, 0.6, and thus a higher intercept, 7.6%. As we will see below, IPOs are more sensitive to market conditions than new rounds, so an estimate that emphasizes new rounds sees a lower slope. As in the IPO/acquisition sample, the average arithmetic returns, driven by large idiosyncratic volatility, are huge at 59%, with 100% standard deviation and 45% arithmetic α . The selection function parameter b is much lower, centering that function at 130% growth in log value. The typical firm builds value through several rounds before IPO, so this is what we expect. The measurement error π is lower, showing the smaller fraction of large outliers in the round to round valuations. The asymptotic standard errors in Table 5 are quite similar to those of the IPO/acquisition sample. Once again, the bootstrap standard errors are larger, but the parameters are still well estimated.

5.2. Alternative reference returns

Perhaps the Nasdaq or small-stock Nasdaq portfolios provide better reference returns than the S&P500. We are interested in comparing venture capital to similar traded securities, not in testing an absolute asset pricing model, so a performance attribution approach is appropriate. The next three rows of Tables 3 and 4 address this case.

In the IPO/acquisition sample of Table 3, the slope coefficient declines from 1.7 to 1.2 using Nasdaq and to 0.9 using the CRSP Nasdaq decile 1 (small) stocks. We expect betas nearer to one if these are more representative as reference returns. However, the residual standard deviation actually goes up a little bit, so the implied R^2 s are even smaller. The mean log returns are about the same, and the arithmetic alphas *rise* slightly.

Nasdaq <\$2M is a portfolio of Nasdaq stocks with less than \$2 million in market capitalization, rebalanced monthly. I discuss this portfolio in detail below. It has a 71% mean arithmetic return and a 62% S&P500 alpha, compared to the 22% mean arithmetic return and statistically insignificant 12% alpha for the CRSP Nasdaq decile 1, so a β near one on this portfolio would eliminate the arithmetic alpha in venture capital investments. This portfolio is a little more successful. The log intercept declines to -27% , but the slope coefficient is only 0.5 so it only cuts the arithmetic alpha down to 22%. In the round-to-round sample of Table 4, there are small changes in the slope and log intercept γ from changing the reference return, but the 60% mean arithmetic return and 45% arithmetic alpha are basically unchanged.

Perhaps the complications of the market model are leading to trouble. The “No δ ” rows of Tables 3 and 4 estimate the mean and standard deviation of log returns directly. The mean log returns are just about the same. In the IPO/acquisition sample of Table 3, the standard deviation is even larger at 105%, leading to larger mean arithmetic returns, 72% rather than 59%. In the round-to-round sample of Table 4, all means and standard deviations are just about the same with no δ .

5.3. Rounds

The “Round i ” subsamples in Tables 3 and 4 break the sample down by investment rounds. It’s interesting to see whether the different rounds have different characteristics, i.e., whether later rounds are less risky. It’s also important to do this for the IPO/acquisition sample, for two reasons. First, the model taken literally should not be applied to a sample with several rounds of the same firm, since we cannot normalize the initial values of both first and second rounds to a dollar and use the same probability of new financing as a function of value. Applying the model to each round separately, we avoid this problem. The selection function is rather flat, however, so mixing the rounds might not make much practical difference. Second, the use of overlapping rounds from the same firm induces cross-correlation between observations, ignored by my maximum likelihood estimate. This should affect standard errors and not bias point estimates. When we look at each round separately, there is no overlap, so standard errors in the round subsamples will indicate whether this cross-correlation in fact has any important effects.

Table 2 already suggests that later rounds are slightly more mature. The chance of ending up as an IPO rises from 17% for the first round to 31% for the fourth round in the IPO/acquisition sample, and from 5% to 18% in the round-to-round sample. However, the chance of acquisition and failure is the same across rounds.

In the IPO/acquisition sample of Table 3, later rounds have progressively lower mean log returns, from 19% to -0.8% , steadily lower slope coefficients, from 1.0 to 0.5, and steadily lower intercepts, from 3.7% to -12% . All of these estimates paint the picture that later rounds are less risky—and hence less rewarding—investments. These findings are consistent with the theoretical analysis of Berk et al. (2004). The asymptotic standard error of the intercept γ (Table 5) grows to five percentage points by round 4, however, so the statistical significance of this pattern that γ declines across rounds is not high. The volatilities are huge and steady at about 100%, so we still see large average arithmetic returns and alphas in all rounds. Still, even these decline across rounds; arithmetic mean returns decline from 71% to 51% and arithmetic alphas decline from 53% to 39% from first-round to fourth-round investments. The cutoff for going out of business k declines for later rounds, the center point of the selection function b declines from 4.2 to 2.5, and the measurement error π declines, all of which suggest less risky and more mature projects in later rounds.

These patterns are all confirmed in the round-to-round sample of Table 4. As we move to later rounds, the mean log return, intercept, and slope all decline, while volatility is about the same. The mean arithmetic returns and alphas are still high, but means decline from 72% to 46% and alphas decline from 55% to 37% from the first to fourth rounds.

In Table 5, the standard errors for round 1 (with the largest number of observations) of the IPO/acquisition sample are still quite small compared to economically interesting variation in the coefficients. The most important change is the standard error of the intercept γ which rises from 0.67 to 1.23. Thus, even if there

is *perfect* cross-correlation between rounds, in which case additional rounds give no additional information, the coefficients are well measured.

5.4. Industries

Venture capital is not all dot-com. Table 2 shows that roughly one-third of the sample is in health, retail, or other industry classifications. Perhaps the unusual results are confined to the special events in the dot-com sector during this sample. Table 2 shows that the industry subsamples have remarkably similar fates, however. Technology (“info”) investments do not go public much more frequently, or fail any less often, than other industries.

In Tables 3 and 4, mean log returns are quite similar across industries, except that “Other” has a slightly larger mean log return (25% rather than 15–17%) in the IPO/acquisition sample, and a much lower mean log return (8% rather than 25%) in the round-to-round sample. However, the small sample sizes mean that these estimates have high standard errors in Table 5, so these differences are not likely to be statistically significant.

In Table 3, we see a larger slope $\delta = 1.4$ for the information industry, and a correspondingly lower intercept. Firms in the information industry went public following large market increases, more so than in the other industries.

The main difference across industries is that information and retail have much larger residual and overall variance, and lower failure cutoffs k . Variance is a key parameter in accounting for success, especially early successes, as a higher variance increases the mass in the right tail. Variance together with the cutoff k accounts for failures, as both parameters increase the left tail. Thus, the pattern of higher variance and lower k is driven by the larger number of early and highly profitable IPOs in the information and retail industries, together with the fact that failures are about the same across industries.

Since the volatilities are still high, we still see large mean arithmetic returns and arithmetic alphas, and the pattern is confirmed across all industry groups. The retail industry in the IPO/acquisition sample is the champion, with a 106% arithmetic alpha, driven by its 127% residual standard deviation and slightly negative beta. The large arithmetic returns and alphas occur throughout venture capital, and do not come from the high tech sample alone.

6. Facts: fates and returns

Maximum likelihood gives the appearance of statistical purity, yet it often leaves one unsatisfied. Are there robust stylized facts behind these estimates? Or are they driven by peculiar aspects of a few data points? Does maximum likelihood focus on apparently well-measured but economically uninteresting moments in the data, at the expense of capturing apparently less well-measured but more economically important moments? In particular, the finding of huge arithmetic returns and alphas sits uncomfortably. What facts in the data lie behind these estimates?

As I argue earlier, the crucial stylized facts are the pattern of exits—new financing, acquisition, or failure—with project age, and the returns achieved as a function of age. It is also interesting to contrast the selection-biased, direct return estimates with the selection-bias-corrected estimates above. So, let us look at the observed returns, and at the speed with which projects get a return or go out of business.

6.1. Fates

Fig. 3 presents the cumulative fraction of rounds in each category—new financing or acquisition, out of business, or still private—as a function of age, for the IPO/acquired sample. The dashed lines give the data, while the solid lines give the predictions of the model, using the baseline estimates from Table 3.

The data paint a picture of essentially exponential decay. About 10% of the remaining firms go public or are acquired with each year of age, so that by five years after the initial investment, about half of the rounds have gone public or been acquired. (The pattern is slightly speeded up in later subsamples. For example,




Fig. 3. Cumulative probability that a venture capital financing round in the IPO/acquired sample will end up IPO or acquired, out of business, or remain private, as a function of age. Dashed lines: data. Solid lines: prediction of the model, using baseline estimates from Table 3.

projects that start in 1995 go public and out of business at a slightly faster rate than projects that start in 1990. However, the difference is small, so age alone is a reasonable state variable.)

The model replicates these stylized features of the data reasonably well. The major discrepancy is that the model seems to have almost twice the hazard of going out of business seen in the data, and the number remaining private is correspondingly lower. However, this comparison is misleading. The data lines in Fig. 3 treat out-of-business dates as real, while the estimate treats data that say “out of business on date t ” as “went out of business on *or before* date t ”, recognizing VentureOne’s occasional cleanups. This difference means that the estimates recognize failures about twice as fast as in the VentureOne data, and that is the pattern we see in Fig. 3. Also, the data lines characterize only the sample with good date information, while the model estimates are chosen to fit the entire sample, including firms with bad date data. And, of course, maximum likelihood does not set out to pick parameters that fit this one moment as well as possible.

Fig. 4 presents the same picture for the round-to-round sample. Things happen much faster in this sample, since the typical investment has several rounds before going public, being acquired, or failing. Here roughly 30% of the remaining rounds go public, are acquired, or get a new round of financing each year. The model provides an excellent fit, with the same understanding of the out-of-business lines.

6.2. Returns

Table 6 characterizes observed returns in the data, i.e., when there is a new financing or acquisition. The column headings give age bins in years. For example, the “1–2 year” column summarizes all investment rounds that went public or were acquired between one and two years after the venture capital financing round, and for which I have good return and date data. The average log return in all age categories of the IPO/acquisition sample is 108% with a 135% standard deviation. This estimate contrasts strongly with the selection-bias-corrected estimate of a 15% mean log return in Table 3. Correcting for selection bias has a huge impact on estimated mean log returns.

Fig. 5 plots smoothed histograms of log returns in age categories. (The distributions in Fig. 5 are normalized to have the same area; they are the distribution of returns *conditional* on observing a return in the indicated time frame.) The distribution of returns in Fig. 5 shifts slightly to the right and then stabilizes. The average log returns in Table 6 show the same pattern: they increase slightly with horizon out to 1–2 years, and then stabilize. These are total returns, not annualized. This behavior is unusual. Log returns usually grow with horizon, so we expect five-year returns five times as large as the one-year returns, and $\sqrt{5}$ times as spread out. Total returns that stabilize are a signature of a selected sample. In the simple example that all projects go public when they have achieved 1,000% growth, the distribution of measured total returns is the *same*—a point mass at 1,000%—for all horizons. Fig. 5 dramatically makes the case that we should regard venture capital

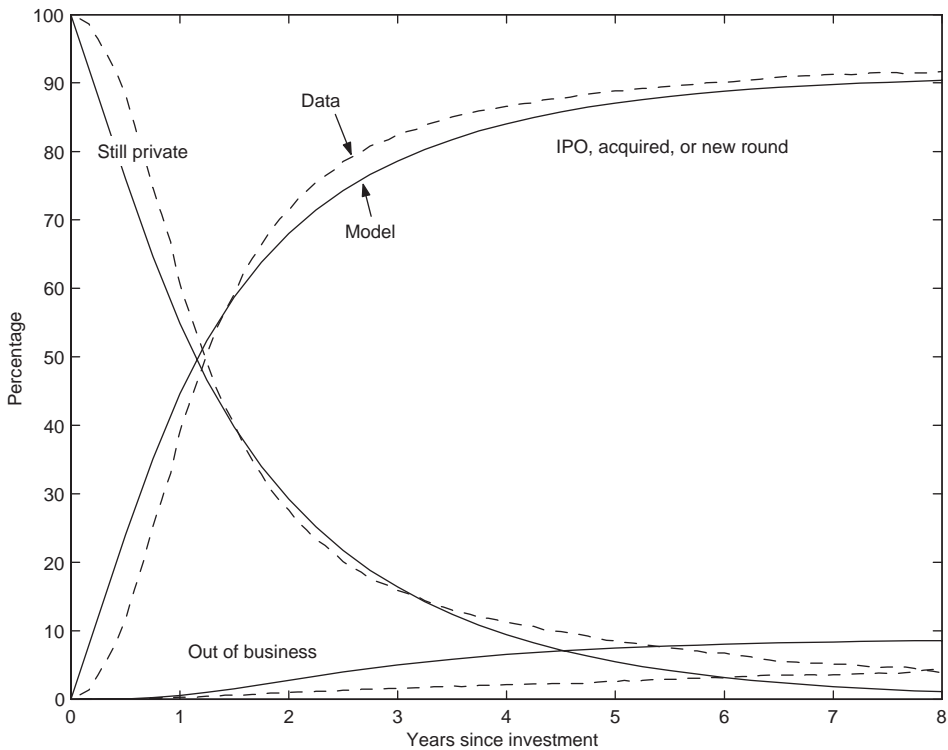


Fig. 4. Cumulative probability that a venture capital financing round in the round-to-round sample will end up IPO, acquired, or new round; out of business; or remain private, as a function of age. Dashed lines: data. Solid lines: prediction of the model, using baseline estimates from Table 4.

projects as a selected sample, with a selection function that is stable across project ages.

Fig. 5 shows that, despite the 108% mean log return, a substantial fraction of projects go public or are acquired at valuations that generate losses to the venture capital investors, even on projects that go public or are acquired soon after the venture capital investment (0–1 year bin). Venture capital has a high mean return, but it is not a gold mine.

Fig. 6 presents the histogram of log returns as predicted by the model, using the baseline estimate of Table 3. The model captures the return distributions of Fig. 5 quite well. In particular, note how the model return distributions settle down to a constant at five years and above.

Fig. 6 also includes the estimated selection function, which shows how the model accounts for the pattern of observed returns across horizons. In the domain of the 3 month return distribution, the selection function is low and flat. A small fraction of projects go public, with a return distribution generated by the lognormal with a small mean and a huge volatility, and little modified by selection. As the horizon increases, the underlying return distribution shifts to the right, and starts to run in to the

Table 6
 Statistics for observed returns

	Age bins							
	1 month–∞	1–6 month	6–12 month	1–2 year	2–3 year	3–4 year	4–5 year	5 year–∞
(1) IPO/acquisition sample								
Number	3,595	334	476	877	706	525	283	413
<i>(a) Log returns, percent (not annualized)</i>								
Average	108	63	93	104	127	135	118	97
Std. dev.	135	105	118	130	136	143	146	147
Median	105	57	86	100	127	131	136	113
<i>(b) Arithmetic returns, percent</i>								
Average	698	306	399	737	849	1067	708	535
Std. dev.	3,282	1,659	881	4,828	2,548	4,613	1,456	1,123
Median	184	77	135	172	255	272	288	209
<i>(c) Annualized arithmetic returns, percent</i>								
Average	3.7e+09	4.0e+10	1,200	373	99	62	38	20
Std. dev.	2.2e+11	7.2e+11	5,800	4,200	133	76	44	28
<i>(d) Annualized log returns, percent</i>								
Average	72	201	122	73	52	39	27	15
Std. dev.	148	371	160	94	57	42	33	24
(2) Round-to-round sample								
<i>(a) Log returns, percent</i>								
Number	6,125	945	2,108	2,383	550	174	75	79
Average	53	59	59	46	44	55	67	43
Std. dev.	85	82	73	81	105	119	96	162
<i>(b) Subsamples. Average log returns, percent</i>								
New round	48	57	55	42	26	44	55	14
IPO	81	51	84	94	110	91	99	99
Acquisition	50	113	84	24	46	39	44	–0

Note: The “IPO/acquisition” sample consists of all venture capital financing rounds that eventually result in an IPO or acquisition in the indicated time frame and with good return data. The “round-to-round” sample consists of all venture capital financing rounds that get another round of financing, IPO, or acquisition in the indicated time frame and with good return data.

steeply rising part of the selection function. Since the winners are removed from the sample, the measured return distribution then settles down to a constant. The risk facing a venture capital investor is as much *when* his or her return will occur as *how much* that return will be.




Fig. 5. Smoothed histogram of log returns by age categories, IPO/acquisition sample. Each point is a normally weighted kernel estimate.

The estimated selection function is actually quite flat. In Fig. 6, it only rises from a 20% to an 80% probability of going public as log value rises from 200% (an arithmetic return of $100 \times (e^2 - 1) = 639\%$) to 500% (an arithmetic return of $100 \times (e^5 - 1) = 14,741\%$). If the selection function were a step function, we would see no variance of returns conditional on IPO or acquisition. The smoothly rising selection function is required to generate the large variance of observed returns.

6.3. Round-to-round sample

Table 6 presents means and standard deviations in the round-to-round sample; Fig. 7 presents smoothed histograms of log returns for this sample, and Fig. 8 presents the predictions of the model, using the round-to-round sample baseline estimates. The average log returns are about half of their value in the IPO/acquisition sample, though still substantial at about 50%. Again, we expect this result since most firms have several venture rounds before going public or being acquired. The standard deviation of log returns is still substantial, around 80%. As the round to round means are about half the IPO/acquisition means, the round to




Fig. 6. Distribution of returns conditional on IPO/acquisition, predicted by the model, and estimated selection function. Estimates from “All” subsample of IPO/acquisition sample.

round variances are about half the IPO/acquisition variances, and round to round standard deviations are lower by about $\sqrt{2}$. The return distribution is even more stable with horizon in this case than in the IPO/acquisition sample. It does not even begin to move to the right, as an unselected sample would do. The model captures this effect, as the model return distributions are even more stable than in the IPO/acquisition case.

6.4. Arithmetic returns

The second group of rows in the IPO/acquisition part of Table 6 presents arithmetic returns. The average arithmetic return is an astonishing 698%. Sorted by age, it rises from 306% in the first six months, peaking at 1,067% in year 3–4 and then declining a bit to 535% for years 5+. The standard deviations are even larger, 3,282% on average and also peaking in the middle years.

Clearly, arithmetic returns have an extremely skewed distribution. Median net returns are half or less of mean net returns. The high average reflects the small possibility of earning a truly astounding return, combined with the much larger probability of a more modest return. Summing squared returns really emphasizes the




Fig. 7. Smoothed histogram of log returns, round-to-round sample. Each point is a normally weighted kernel estimate. The numbers give age bins in years.

few positive outliers, leading to standard deviations in the thousands. These extreme arithmetic returns are just what one would expect from the log returns and a lognormal distribution: $100 \times (e^{1.08+(1/2)1.35^2} - 1) = 632\%$, close to the observed 698%. To make this point more clearly, Fig. 9 plots a smoothed histogram of log returns and a smoothed histogram of arithmetic returns, together with the distributions implied by a lognormal, using the sample mean and variance. This plot includes all returns to IPO or acquisition. The top plot shows that log returns are well modeled by a normal distribution. The bottom plot shows visually that arithmetic returns are hugely skewed. However, the arithmetic returns coming from a lognormal with large variance are also hugely skewed, and the fitted lognormal captures the right tail quite well. The major discrepancy is in the left tail, but kernel density estimates are not good at describing distributions in regions where they slope a great deal, and that is the case here.

Though the estimated 59% mean and 107% standard deviation of arithmetic returns in Table 3 might have seemed surprisingly high, they are nothing like the 698% mean and 3,282% standard deviation of arithmetic returns with no sample




Fig. 8. Distribution of returns conditional on new financing predicted by the model, and selection function. “All” estimate of the round-to-round sample.

selection correction. The sample selection correction has a dramatic effect on estimates of the arithmetic mean return.

6.5. Annualized returns

It might seem strange that so far I have presented total returns without annualizing. The next two rows of Table 6 show annualized returns. The average annualized return is 3.7×10^9 percent, and the average in the first six months is 4.0×10^{10} percent. These must be the highest average returns ever reported in the finance literature, which just dramatizes the severity of selection bias in venture capital. The mean and volatility of annualized returns then decline sharply with horizon.

The extreme annualized returns result from a small number of sensible returns that occur over very short time periods. If you experience a moderate (in this dataset) 100% return, but it happens in two weeks, the result is a $100 \times (2^{24} - 1) = 1.67 \times 10^9$ percent annualized return. Many of these outliers were checked by hand, and they appear to be real. There is some question whether they represent




Fig. 9. Smoothed histograms (kernel density estimates) and distributions implied by a lognormal. The top panel presents the smoothed histogram of all log returns to IPO or acquisition (solid), using a Gaussian kernel and $\sigma = 0.20$, together with a normal distribution using the sample mean and variance of the log returns (dashed). The bottom panel presents a smoothed histogram of all arithmetic returns to IPO or acquisition, using a Gaussian kernel and $\sigma = 0.25$, together with a lognormal distribution fitted to the mean and variance of log returns.

arm's-length transactions, however. Ebay is a famous story (though not in the dataset). Dissatisfied with the offering price, Ebay got one last round of venture financing at a high valuation, and then went public a short time later at an even larger value. More typically, the dataset contains seed financings quickly followed by first-stage financings involving the same investors. It appears that in many cases, the valuation in the initial seed financing is a matter of little consequence, as the overall allocation of equity will be determined at the time of the first round, or the decision could be made not to proceed with the start-up. (See for instance, the discussion in [Halloran, 1997](#).) While not data errors per se, huge annualized returns from seed to first round in such cases clearly do not represent the general rate of return to venture capital investments. (This is analogous to the “calendar time” vs. “event time” issue in IPO returns.) Below, I check the sensitivity of the estimates to these observations in several ways.

However, the log transformation again gives sensible numbers, so the large average annualized returns are fundamentally a story of extreme volatility, not a story about outliers or data errors. Average annualized log returns also decline roughly with the inverse of the horizon. Again, this is what we expect from a selected sample. For an unselected sample, we expect annualized returns to be stable across horizon, and total returns to grow with horizon. In a selected sample, total returns are stable with horizon, so annualized returns decline with horizon.⁶

In the round-to-round sample, arithmetic returns and annualized returns (not shown) behave in the same way: arithmetic returns are large and very skewed with huge standard deviations; annualized arithmetic returns are huge for short horizons, and annualized returns decline quickly with horizon.

There is no right and wrong here. Statistics are just statistics. Skewed arithmetic returns are what one expects from roughly lognormal returns with extreme variance. The constant total returns and declining annualized returns with horizon are what one expects from a roughly constant total log return distribution, generated by a selection function of value and not of horizon. It is clear from this analysis that one cannot do much of anything with the observed returns without correcting for selection effects.

6.6. Subsamples

How different are returns to a new round, IPO, or acquisition? In addition to the direct interest in these questions, I lumped outcomes together in the estimation, and it's important to check that this procedure is not unreasonable. The final rows of Table 6 present mean log returns across horizons for these subsamples of the round-to-round sample, and Fig. 10 collects the distribution of returns for different outcomes, summing over all ages.

The mean log returns to IPO are a bit larger (81%) than returns to a new round or acquisition (50%). Except for good returns to acquisitions in the first six months, and poor returns to new rounds and acquisitions after five years, each category is reasonably stable over horizon. Fig. 10 shows that the modal return to acquisition is about the same as the modal return to IPO; the lower mean return to acquisition comes from the larger left tail of acquisitions. The largest difference is the surprisingly greater volatility of acquisition returns, and the much lower volatility of new round returns. I conclude that lumping the three outcomes together is not a gross violation of the data, and not worth fixing at the large cost of adding parameters to the already complex ML estimation. Most important, the figure confirms that IPOS are similar to other fundings and revaluations, and not a qualitatively different jackpot as in popular perception.

⁶Also, we should not expect the average annualized arithmetic returns of Table 6 to be stable across horizons, even in an unselected sample. In such a sample, the *annualized average* return is independent of horizon, not the *average annualized* return. $E(R_{0 \rightarrow t}^{1/t}) = E(e^{1/t \ln R_{0 \rightarrow t}}) = e^{(\mu)/t + (t\sigma^2)/2t^2} = e^{\mu + \sigma^2/(2t)}$, while $[E(R_{0 \rightarrow t})]^{1/t} = [e^{\mu + (1/2)\sigma^2 t}]^{1/t} = e^{\mu + \sigma^2/2}$. Small σ and large t approximations do not work well in a dataset with huge σ and occasionally very small t .




Fig. 10. Smoothed histogram of returns, all ages, subsamples of the round-to-round sample. “New round”, “IPO”, and “Acquired” are the returns for all rounds whose next financing is a new round, IPO, or acquisition, from initial investment to the indicated event.

7. How facts drive the estimates

Having seen estimates and a collection of stylized facts, it is time to see how the stylized facts drive the estimates. This discussion can give us confidence that the estimates not driven by a few data points or by odd and untrustworthy aspects of the data.

7.1. Stylized facts for mean and standard deviation

Table 6 finds average log returns of about 100% in the IPO/acquisition sample, stable across horizons, and Fig. 3 shows about 10% of financing rounds going public or being acquired per year in the first few years. These facts allow us to make a simple back-of-the-envelope estimate of the mean and variance of venture capital returns, correcting for selection bias. The same general ideas underlie the more realistic, but hence more complex, maximum likelihood estimation, and this simple

calculation shows how some of the rather unusual results are robust features of the data.

Consider the very simple selection model: we see a return as soon as the log value exceeds b . We can calibrate b to the average log return, or about 100%. Once again, returns identify the selection function. The fraction of projects that go public by year t is given by the right tail of the normal $\Phi\left(\frac{b-t\mu}{\sqrt{t}\sigma}\right)$, where μ and σ denote the mean and standard deviation of log returns. The 10% right tail of a standard normal is 1.28, so the fact that 10% go public in the first year means $\frac{1-\mu}{\sigma} = 1.28$.

A small mean $\mu = 0$ with a large standard deviation $\sigma = \frac{1}{1.28} = 0.78$ or 78% would generate the right tail. However, a small standard deviation $\sigma = 0.1$ or 10% and a huge mean $\mu = 1 - 0.1 \times 1.28 = 0.87$ or 87% would also work. Which is it? The second year separately identifies μ and σ . With a zero mean and a 78% standard deviation, we should see that by year 2, $\Phi\left(\frac{1-2 \times 0}{0.78 \times \sqrt{2}}\right) = 18\%$ of firms have gone public, i.e., an additional 8% in year 2, which is roughly what we see. With a huge mean $\mu = 87\%$ and a small standard deviation $\sigma = 10\%$, we predict that by year 2, essentially all ($\Phi\left(\frac{1-2 \times 0.86}{0.10 \sqrt{2}}\right) = \Phi(-5.2) = (100 - 8 \times 10^{-6})\%$) firms have gone public. This is not at all what we see—more than 80% are still private at the end of year two.

To get rid of the high mean arithmetic returns, despite high variance, we need a strongly negative mean log return. The same logic rules out this option. Given $\frac{1-\mu}{\sigma} = 1.28$, the lowest value of $\mu + \frac{1}{2}\sigma^2$ we can achieve is given by $\mu = -64\%$ and $\sigma = 128\%$ (min $\mu + \frac{1}{2}\sigma^2$ s.t. $\frac{1-\mu}{\sigma} = 1.28$), leading to $\mu + \frac{1}{2}\sigma^2 = 0.18$ and a reasonable mean arithmetic return $100 \times (e^{0.18} - 1) = 20\%$. But a strong negative mean implies that IPOs quickly cease and practically every firm goes out of business in short order as the distribution marches to the left. With $\mu = -64\%$ and $\sigma = 128\%$, we predict that $\Phi\left(\frac{1-2 \times (-0.64)}{1.28 \sqrt{2}}\right) = \Phi(1.26) = 10.4\%$ go public in two years. But 10% go public in the first year, so only 0.4% more go public in the second year. After that, things get worse. $\Phi\left(\frac{1-3 \times (-0.64)}{1.28 \sqrt{3}}\right) = \Phi(1.32) = 9.3\%$ go public by year 3. Since 10% went public already in year 1, this number reveals a distribution moving quickly to the left and the oversimplification of this back-of-the-envelope calculation that ignores intermediate exits.

To see the problem with failures, start with the fact from Fig. 3 that a steady small percentage—roughly 1%—fail each year. The simplest failure model is a step function at k , just like our step function at b for going public. The 1% tail of the normal is 2.33 standard deviations from the mean, so to get 1% to go out of business in one year, we need $\frac{k-\mu}{\sigma} = -2.33$. Using $\mu = -64\%$ and $\sigma = 128\%$, that means $k = -2.33 \times 1.28 - 0.64 = -3.62$. A firm goes out of business when value declines to $100 \times e^{-3.62} = 2.7\%$ of its original value, which is both sensible and close to the formal estimates in Tables 2 and 5. (The latter reports essentially twice this value, since the selection for out of business is a linearly declining function of value rather than a fixed cutoff.) But at these parameters, in two years, $\Phi\left(\frac{k-2\mu}{2\sqrt{\sigma}}\right) =$

$\Phi\left(\frac{-2.34+2\times 0.64}{2\sqrt{1.28}}\right) = \Phi(-0.47) = 32\%$ fail, and by 3 years $\Phi\left(\frac{k-3\mu}{3\sqrt{\sigma}}\right) = \Phi\left(\frac{-2.34+3\times 0.64}{3\sqrt{1.28}}\right) = \Phi(-0.12) = 45\%$ fail!

In sum, the fact that firms *steadily* go public and fail, as seen in Fig. 3, means we must have a log return distribution with a small mean—no strong tendency to move to the left or right—and a high variance. Then the tails, which generate firms that go public or out of business, grow gradually with time. Alas, a mean log return near zero and large variance imply very large *arithmetic* returns. This logic is compelling, and suggests that these central findings are not specific to the sample period.

The round-to-round sample has lower average returns, about 50% in Table 6. We also see more frequent new financings, about 30% per year in Fig. 7. The 30% right tail is 0.52 standard deviations above the mean. Thus, we know from the first year that $\frac{0.5-\mu}{\sigma} = 0.52$. With a mean $\mu = 0$, this implies $\sigma = 0.50/0.52 = 96\%$. The lower observed returns and greater probability of seeing a return are offsetting, giving about the same estimate of standard deviation as for the IPO/acquisition sample.

It is comforting to see and understand the same underlying mean and standard deviation parameters in the two samples, despite their quite different observed means, standard deviations, and histories. This simple calculation shows why, and why it is a robust feature of natural stylized facts.

7.2. Stylized facts for betas

How can we identify and measure betas? In the simple model that all firms go public at b value, we would identify β by an increased *fraction* that go public following a large market return, not by any change in *return*, since all observed returns are the same (b). With a slowly rising selection function, we will see increased returns as well, since the underlying value distribution shifts to the right. The formal estimate also relies on more complex effects. For example, after a runup in the market, many firms will go public, so the distribution of remaining project values will be different than it would have been otherwise. These dynamic effects are harder to characterize as stylized facts.

We can anticipate that these tendencies will be difficult to measure, so that beta estimates might not be precise or robust. With 100% per year idiosyncratic risk, a typical 15% (1σ) rise in the market is a small risk, and shifts the distribution of returns only a small amount. In the simple model, a 15% rise in the market raises the fraction of firms that go public in one year from $\Phi\left(\frac{1-0}{0.78}\right) = \Phi(1.28) = 10\%$ to $\Phi\left(\frac{1-0.15}{0.78}\right) = \Phi(1.09) = 14\%$. The actual selection functions rise slowly, so moving the return distribution to the right 15% will push even fewer firms over the border. Similarly, the huge residual standard deviation means that the R^2 s are low, so market model return regression estimates will be imprecise.

Still, let us see what facts can be documented about returns and fates conditional on index returns. Table 7 presents regressions of observed returns on the S&P500 index return. With arithmetic returns, the intercepts (alpha) are huge. At 462%, this is probably the largest alpha ever claimed in a finance paper, though it surely reflects the severe selection bias in this sample rather than a golden-egg-laying goose. The

Table 7
Market model regressions

	α (%)	$\sigma(\alpha)$	β	$\sigma(\beta)$	R^2 (%)
IPO/acq. arithmetic	462	111	2.0	0.6	0.2
IPO/acq. log	92	3.6	0.4	0.1	0.8
Round to round, arithmetic	111	67	1.3	0.6	0.1
Round to round, log	53	1.8	0.0	0.1	0.0
Round only, arithmetic	128	67	0.7	0.6	0.3
Round only, log	49	1.8	0.0	0.1	0.0
IPO only, arithmetic	300	218	2.1	1.5	0.0
IPO only, log	66	4.8	0.7	0.2	2.1
Acquisition only, arithmetic	477	95	-0.8	0.5	0.3
Acquisition only, log	77	9.8	-0.8	0.3	2.6

Note: Market model regressions are $R_{t \rightarrow t+k} = \alpha + \beta R_{t \rightarrow t+k}^m + \varepsilon_{t \rightarrow t+k}$ (arithmetic) and $\ln R_{t \rightarrow t+k} = \alpha + \beta \ln R_{t \rightarrow t+k}^m + \varepsilon_{t \rightarrow t+k}$ (log). For an investment made at date t and a new valuation (new round, IPO, acquisition) at $t+k$, I regress the return on the corresponding S&P500 index return for the period $t \rightarrow t+k$. Standard errors are plain OLS ignoring any serial or cross correlation.

32% arithmetic alpha in the selection-bias-corrected Table 3 pales by comparison. Once again, though the selection-bias-corrected estimates leave some puzzle, the selection bias correction has dramatic effects on the uncorrected estimates.

The beta for arithmetic returns is large at 2.0. There is a tendency for market returns to coincide with even larger venture capital returns. Log returns trim the outliers, however, and produce a lower beta of 0.4. The R^2 values in these regressions are tiny, as expected. For this reason, betas are poorly measured, despite the huge sample and optimistic plain-vanilla OLS standard errors.

The round-to-round regressions produce lower betas still, suggesting that much of the measured beta comes from a tendency to go public at high market valuations rather than a tendency for new rounds to be more highly valued when the market is high. Splitting into new round, IPO, and acquisition categories we see this pattern clearly. The positive betas come from the IPOs.

Fig. 11 graphs the time series of the fraction of outstanding firms in the IPO/acquisition sample that go public each quarter, along with the previous year's S&P500 returns. (The fraction that goes public is a two-quarter moving average.) If you look hard, you can see that IPOs increase following good market returns in 1992–1993, 1996–1997, and 1999–2000. (There was a huge surge in IPOs in the last two years of the sample. However, there was also a huge surge of new projects, so the fraction of outstanding firms that go public only rises modestly as shown.) 1992 and 1996–1997 also show a modest correlation between average IPO returns and the S&P500 index. For the IPOs, increased numbers rather than larger returns drive the estimated betas.

Fig. 12 graphs the same time series for firms in the IPO/acquisition sample that are acquired. Here we see no tendency at all for the frequency of acquisitions to rise following good market returns. However, we do see that the returns to acquired




Fig. 11. Percentage of outstanding projects going public, percent average log returns for projects going public (right scale), and previous year's percent log S&P500 return. Percentage of projects going public and their returns are two-quarter moving averages. IPO/acquisition sample.

firms track the S&P500 index well, with a scale factor of about two or three. This graph suggests that returns rather than greater frequency of acquisitions drive a beta estimate among acquisitions. However, this picture is not confirmed in Table 7, which found negative betas. There are more observations in later years, so the regression and this graph weight observations differently.

A similar figure for new rounds in the round-to-round sample shows no tendency for an increased frequency of financing, and a barely discernible tendency towards higher values on the tail of stock market rises. The maximum likelihood beta estimates in the round-to-round sample are correspondingly lower and less precise, and are driven by the acquisition and IPO outcomes in that sample.

In sum, the correlation of observed returns with market returns, and the correlation of the frequency of observed new financing or acquisition with market returns, form the basic stylized facts behind beta estimates. The stylized facts are there: the frequency of IPOs rises when the market rises, and the valuation of acquisitions rises when the market rises. However, the stylized facts are much weaker than those that drive average returns and the variance of returns. This weakness explains why the intercept and beta estimates of the formal model are not




Fig. 12. Percent of outstanding projects acquired, average log returns of acquisitions (right scale), and previous year's S&P500 return, IPO/acquisition sample. Percent acquired is a two-quarter moving average.

particularly well estimated or stable across subsamples or variations in technique, while the average and standard deviation of log returns are quite stable. This weakness also explains why I have not extended the estimation, for example to three-factor betas or other risk corrections.

8. Testing $\alpha = 0$

An arithmetic return of 59% and a 32% arithmetic alpha are still uncomfortably large. We have already seen that they result from a mean log return near zero, the large volatility of log returns, and $e^{\mu + \frac{1}{2}\sigma^2}$. We have seen in a back-of-the-envelope way that $\mu = -50\%$ would produce IPOs that cease after a few years and all firms soon failing. But perhaps the more realistic model and formal estimate do not speak so strongly against $\alpha = 0$. What if we change all the parameters? In particular, can we accept the high mean arithmetic return, but imagine a β of three to five so that the high mean return is explained? The stylized facts behind high volatility are compelling, but those driving us to a small beta are not so convincing. Can we imagine that the data are wrong in simple ways that would overturn the finding of a

Table 8
Additional estimates and tests for the IPO/acquisition sample

	E ln R	σ ln R	γ	δ	σ	ER	σR	α	β	k	a	b	π	χ^2
All, baseline	15	89	-7.1	1.7	86	59	107	32	1.9	25	1.0	3.8	9.6	
$\alpha = 0$	-0.9	82	-30	2.5	73	34	93	0.0	2.6	23	0.9	3.9	15	1,428
ER = 15%	-3.3	60	-3.3		60	15	64			28	1	3.4	28	2,523
Pre-1997	11	81	11	-0.8	80	46	94	48	-0.8	9.6	1.0	3.6	4.4	
Dead 2000	36	59	27	0.3	59	58	69	48	0.3	150	0.7	4.9	31	
No π	11	115	-4.0	0.9	114	85	152	67	1.1	11	0.6	5.8		170

Note: “ $\alpha = 0$ ” imposes $\alpha = 0$ on the estimation, by always choosing γ so that, given the other parameters, the arithmetic α calculation is zero. “ER = 15%” imposes that value on the no- δ estimation, choosing γ so that the arithmetic average return calculation is always 15%. χ^2 gives the likelihood ratio statistic for these parameter restrictions. Each statistic is $\chi^2(1)$ with a 5% critical value of 3.84. “Pre-1997” limits the data sample to January 1 1997, treating as “still private” any exits past that date. “Dead 2000” assumes that any project still private at the end of the sample goes out of business. “No π ” removes the measurement error.

Table 9
Additional estimates for the round-to-round sample

	E ln R	σ ln R	γ	δ	σ	ER	σR	α	β	k	a	b	π	χ^2
All, baseline	20	84	7.6	0.6	84	59	100	45	0.6	21	1.7	1.3	4.7	
$\alpha = 0$	-3.6	77	-27	1.9	72	27	86	0.0	1.9	21	1.5	1.4	6.8	1,807
ER = 15%	-8.9	69	-8.9		69	15	74			19	2.2	1.0	9.9	3,060
Pre-1997	21	75	10	0.4	75	52	87	40	0.4	19	0.4	5.1	2.2	
Dead 2000	32	76	16	0.9	74	65	91	47	1.0	108	0.3	6.4	7.2	
No π	16	104	1.6	0.9	103	77	133	60	1.0	11	1.2	1.8		864

Note: See note to Table 8.

high α ? All these questions point naturally to an estimate with restricted parameters such that $\alpha = 0$, and a likelihood ratio test.

Table 8 presents additional estimates for the IPO/acquisition sample, starting with a test of $\alpha = 0$. (I solve Eq. (7) for the value of γ that, given the other parameters, results in $\alpha = 0$, and I fix γ at that value in the estimation.) Table 10 collects asymptotic standard errors. Imposing $\alpha = 0$ lowers the mean log return from 15% to -0.9%. Together with a slightly lower standard deviation, the mean arithmetic return is cut in half, from 59% to 34%. However, imposing $\alpha = 0$ changes the decomposition of mean log return, lowering the log model intercept from -7.1% to -30%, and raising the slope coefficient δ from 1.7 to 2.5. Interestingly, the estimate

Table 10
Asymptotic standard errors for Tables 8 and 9 estimates

	IPO/acquisition sample							Round-to-round sample						
	γ	δ	σ	k	a	b	π	γ	δ	σ	k	a	b	π
$\alpha = 0$		0.06	0.7	0.59	0.03	0.13	0.8		0.01	0.6	0.4	0.04	0.03	0.4
ER = 15%			0.6	0.65	0.01	0.01	1.1			0.6	0.3	0.02	0.01	0.6
Pre-1997	1.2	0.11	1.1	0.42	0.04	0.12	0.8	1.3	0.12	1.1	0.9	0.01	0.06	0.4
Dead 2000	0.7	0.05	1.2	0.08	0.03	0.16	1.1	1.0	0.06	1.1	1.1	0.00	0.02	0.5
No π	1.1	0.08	1.1	0.37	0.02	0.17		1.2	0.08	0.8	0.2	0.02	0.03	

does not just raise beta. It achieves half of the alpha decline via the difficult route of lowering mean returns. Apparently, there is strong sample evidence against the high-beta parameterization, despite the apparent weakness of stylized facts seen in the last section. The estimate also increases measurement error, to try to handle observations that now cause trouble. Alas, the statistical evidence against this parameterization is strong. Imposing $\alpha = 0$, the log likelihood declines by 1,428/2. Compared to the 5% $\chi^2(1)$ critical value of 3.84, the $\alpha = 0$ restriction is spectacularly rejected.

The round-to-round sample in Table 9 behaves similarly. The average log return declines from 20% to -3.6%, and the average arithmetic return is cut in half from 59% to 27%. The intercept declines dramatically from $\gamma = +7.6\%$ to $\gamma = -27\%$, and the slope rises from 0.6 to 1.9. But the $\chi^2(1)$ likelihood ratio statistic is 1,807, an even more spectacular rejection.

So far, the estimates raise slope coefficients a good deal in order to lower alphas. We might want to keep the estimate from following this path in order to examine the evidence against the core troubling estimate of high average arithmetic returns, rather than to excuse such returns by large poorly measured betas. In the ER = 15% rows of Tables 8 and 9, I impose an average arithmetic return of 15%, the same as the S&P500 in this sample, estimating the mean and variance of returns directly, i.e., restricting the “no δ ” estimate of Tables 3 and 4. The model might have kept the high standard deviation, and matched it with a -50% or so mean log return in order to reduce $\mu + \frac{1}{2}\sigma^2$. Instead, the dynamic evidence for a mean log return near zero is so strong that the estimate keeps it, with $E \ln R = -3.3\%$ in the IPO/acquisition sample and $E \ln R = -8.9\%$ in the round-to-round sample. The estimate reduces standard deviation accordingly, to 60% in the IPO/acquisition sample and 69% in the round-to-round sample. These variance reductions are just enough to produce the desired 15% mean arithmetic return via $e^{\mu + \frac{1}{2}\sigma^2}$. However, this reduction in variance does great damage to the model’s ability to fit the dynamic pattern of new financing. The measurement error probabilities rise to 28% in the IPO/acquisition sample, and to 9.9% in the round-to-round sample. The $\chi^2(1)$ likelihood ratio statistics are 2,523 for IPO/acquisition and 3,060 in the round-to-round samples, even more decisively rejecting the ER = 15% restriction.




Fig. 13. Simulated fates of venture capital investments, imposing $\alpha = 0$. In the “ $\alpha = 0$ ” case (triangles), I impose the condition that the arithmetic α is zero and maximize likelihood over the remaining parameters. In the “ $\alpha = 0$, others unchanged” case (no symbols), I change the intercept γ to produce $\alpha = 0$, leaving other parameters unchanged. Squares give the data. Dashed lines plot IPO/acquisitions; solid lines plot failures.

Where is the great violence to the data indicated by these likelihood ratio statistics? Fig. 13 compares the simulated fates in the $\alpha = 0$ restricted models to the simulated fates with the baseline estimates, and Table 11 characterizes the simulated distribution of observed returns with the various restricted models. Table 11 is meant to convey the same information as the return distribution in Figs. 6 and 8 in more compact form.

I start by lowering the intercept γ to produce $\alpha = 0$ with no change in the other parameters. The “ $\alpha = 0$, others unchanged” lines of Fig. 13 shows that this restriction produces far too many bankruptcies and too few IPO/acquisitions. (The dashed line with no symbols, representing the estimate’s prediction of IPO/acquisitions, is well below the dashed line with squares, representing the data; and the solid line with no symbols, representing the estimate’s prediction of out of business projects, is well above the solid line with squares, representing failures in the data.) As the back-of-the-envelope calculation suggested, a low mean log return implies that the distribution of values moves to the left over time, so we have an inadequate right tail of successes and too large a left tail of failures.

Table 11
Moments of simulated returns to new financing or acquisition under restricted parameter estimates

Horizon (years):	1. IPO/acquisition sample					2. Round-to-round sample				
	1/4	1	2	5	10	1/4	1	2	5	10
(a) E log return (%)										
Baseline estimate	21	78	128	165	168	30	70	69	57	55
$\alpha = 0$	11	42	72	101	103	16	39	34	14	10
ER = 15%	8	29	50	70	71	19	39	31	13	11
(b) σ log return (%)										
Baseline estimate	18	68	110	135	136	16	44	55	60	60
$\alpha = 0$	13	51	90	127	130	12	40	55	61	61
ER = 15%	9	35	62	91	94	11	30	38	44	44

The “ $\alpha = 0$ ” lines of Fig. 13 present the simulated histories when I impose $\alpha = 0$, but allow ML to search over the other parameters, in particular raising the slope coefficient β and measurement error, so as to give $\alpha = 0$ while keeping a large mean arithmetic return. Now the estimate can match the pattern of successes (the dashed lines with triangles and squares are close), but it still predicts far too many failures (the solid line with triangles is far above the solid line with squares). The 20% lower mean log and 30% lower mean arithmetic return in this estimate still leave a distribution that marches off to the left too much. The ER = 15% restriction and the round-to-round sample behave similarly.

In Table 11, the mean returns to new financing or acquisition under the restricted models are often less than half the mean returns under the unrestricted model and in the data. The standard deviations are often a poor match in some of the parameterizations. The restricted estimates also miss facts underlying the beta estimates, although I do not graph this phenomenon.

In sum, the data speak strongly against lowering the arithmetic alpha to zero, either by lowering mean arithmetic returns or by raising betas. To believe such a parameterization, we must believe that beta is much larger than estimated, we must believe that the data are measured with much more error, we must believe that the data substantially understate the frequency and timing of failure (as indicated by Fig. 13), and we must believe that the sample systematically overstates the returns to IPO, acquisition, and new financing by as much as a factor of two, as indicated by Table 11.

9. Robustness

I check that the anomalous IPO market at the end of the sample, measurement error, and the imputation of returns to out-of-business projects do not affect the results.

9.1. End of sample

We might suspect that the results depend crucially on the anomalous behavior of the IPO market during the late 1990s, and the unfortunate fact that the sample stops in June of 2000, just after the first Nasdaq crash. (This fact is not a coincidence—the data collection for this project was commissioned by a now defunct dot-com.)

To address this concern, the “pre-1997” subsample uses no information after January 1997. I ignore all rounds that are not started by January 1997, and I treat all rounds started before then that have not yet gone public, been acquired, had another round, or gone out of business by January 1997 as “still private”. The “Dead 2000” sample assumes that all firms still private as of June 2000 go out of business on that date. This experiment also provides a way to address the “living dead” bias: some firms that are reported as still alive are probably really inactive and worthless. Assuming *all* inactive firms are worthless by the end of the sample gives a bound on how severe that bias could be.

As Table 2 shows, about two-thirds of the venture capital financing rounds begin after January 1997, so the concern that the results are special to the subsample is not unfounded. However, the main difference in fates is that firms are much more likely to *fail* in the post-1997 sample. The fraction that go public, etc., is virtually identical. If we assume that all firms alive in June 2000 go out of business, we increase the out-of-business fraction dramatically, at the expense of the “still private” category.

In Tables 8 and 9, the mean and standard deviation of log returns are essentially the same in the pre-1997 sample as in the base case. The main difference is the split of the mean return between slope and intercept for the IPO/acquisition sample. The slope coefficient δ switches sign from +1.7 to -0.8 , and the intercept γ rises from -7.1% to $+11\%$. The association of IPOs with the stock price rise of the late 1990s is the major piece of information identifying the slope δ . Since volatility is unchanged and the mean log return is unchanged, mean arithmetic returns are essentially unchanged in the pre-1997 sample. Since the slopes decline, arithmetic alphas actually *increase* in the pre-1997 sample, to 48% in Table 8 and 40% in Table 9.

Assuming that all firms still private in June 2000 go out of business on that date plays havoc with the estimate. The failure cutoff k increases to 150% of its initial value in Table 8 and 108% in Table 9, naturally enough, as the chance of failure increases dramatically. The other parameters change a bit, as they must still account for the successes in the pre-2000 data despite much higher k . Both samples show a much larger mean log return, and the IPO/acquisition sample shows a somewhat smaller variance.

In the end, the mean arithmetic returns and alphas are the same or higher in the pre-1997 and Dead 2000 samples. As always, the idiosyncratic variance remains large and it is not paired with a huge negative mean. Thus, neither the late 1990s boom nor a “living dead” bias is behind the central results.

9.2. Measurement error and outliers

How does the measurement error process affect the estimates? In the rows of Tables 8 and 9 marked “no π ”, I remove the measurement error process. This change raises the estimated standard deviation from 89% to 115% in Table 8 and from 84% to 104% in Table 9. Absent measurement error, we need a larger variance to accommodate tail returns. The mean log return is unaffected. Higher variance alone would drive more firms to failure, so the failure cutoff k drops from 25% to 11% in Table 8 and 21% to 11% in Table 9. All the other estimated parameters are basically unchanged. Raising the variance raises mean arithmetic returns to 85% and α to 67% in Table 8, and mean arithmetic returns to 65% and α to 60% in Table 9. Repeating the whole set of estimations without measurement error, the largest difference, in addition to the larger variance, is much less stability in slope coefficients δ across subsamples. A few large returns, very unlikely with a lognormal distribution, drive the δ estimates without measurement error. The estimates vary as the few influential data points jump in and out of subsamples.

The likelihood ratio statistic for $\pi = 0$ is 170 in the IPO/acquisition sample and 864 in the round-to-round sample. The 5% critical value for a $\chi^2_{(1)}$ is 3.84. Whether we interpret the measurement error process as such, or as a device to induce a fatter tail in the true return process,⁷ the model wants it.

The measurement error process does *not* just throw out large returns, which are plausibly the most interesting part of venture capital. It largely throws out reasonable returns that occur in a very short time period, leading to very large annualized returns. Even if not errors, these events are a separate phenomenon from what most of us think as the central features of venture capital. Venture capital is about the possibility of earning a very large return in a few years, not about the chance of “only” doubling your money in a month.

To document this interpretation, I examine “outliers,” the data points that contribute the greatest (negative) amount to the log likelihood in each estimate. Without measurement error, the biggest outliers are IPO/acquisitions that have moderately large positive—and negative—returns in a short time span, not large returns per se. With measurement error, the outliers are old (eight to ten years) projects that eventually go public or are acquired with very low returns—5% to 20% of initial value. This finding is sensible. Since low values exit, and the probability of going public is very low at a low value, it is hard to attain a very low value and go public without failing along the way.

To check further that the high mean returns and alphas are not driven by anomalous quick successes, implying huge annualized returns, I try replacing the actual age of returns in the first year with “1 year or less”, i.e. summing the

⁷We cannot interpret this exact specification of the measurement error process as a fat-tailed return distribution. The measurement error distribution is applied only once, and does not cumulate. A fat-tailed return distribution, or, equivalently, the addition of a jump process, is an interesting extension, but one I have not pursued to keep the number of parameters down and to preserve the ease of making transformations such as log to arithmetic based on lognormal formulas.

probability of an IPO over the first year rather than using the probability of achieving the IPO on the reported date. This variant has practically no effect at all on the estimated parameters.

In sum, the measurement error complication does not drive the large alphas. Quite the opposite, adding measurement error reduces the volatility-induced mean arithmetic returns and alphas by accounting for the occasional quick large returns.

9.3. Returns to out-of-business projects

So far, I have implicitly assumed that when a firm goes out of business, the investor receives whatever value is left. What if, instead, investors get nothing when the firm goes out of business? This change adds a lumpy left tail to the return distribution. Perhaps this lumpy left tail is enough to get rid of the troublesome alphas. Mean log returns become $-\infty$, and the standard deviation of log returns $+\infty$, but we can still characterize the mean and standard deviation of arithmetic returns.

To answer this question, I simulate the model at the baseline parameter estimates, and find the probability and value of all the various outcomes. I then calculate the annualized expected arithmetic return, assuming that investors get zero return for any project that goes out of business. (Since we are aggregating payoffs at different horizons, I calculate the arithmetic discount rate that sets the present value of the cash flows to one.) The average arithmetic return declines only from 58.72% to 58.38%. This modification has so little effect because the failure values k are quite low, around 10% of the initial investment, and only 9% of firms fail. Losing the last 10%, in the 9% of investments that are down to 10% of initial value, naturally has a small effect on average returns.

10. Comparison to traded securities

If we admit large arithmetic mean returns, standard deviations, and arithmetic alphas in venture capital, are these findings unique, or do similar traded securities behave the same way?

Table 12 presents means, standard deviations, and market model regressions for individual small Nasdaq stocks. To form the subsamples, I take all stocks that have market value below the indicated cutoffs in month t , and I examine their returns from month $t + 1$ to month $t + 2$. I lag by two months to ensure that erroneously low prices at t do not lead to spuriously high returns from t to $t + 1$, though results with no lag (selection in t , return from t to $t + 1$) are in fact quite similar. I examine market value cutoffs of two million, five million, 10 million and 50 million dollars. The average venture capital financing round in my sample raises \$6.7 million. The first five deciles of all Nasdaq market value observations in this period occur at 5, 10, 17, 27, and 52 million dollars, so my cutoffs are approximately the $\frac{1}{20}$, $\frac{1}{10}$, $\frac{2}{10}$, and $\frac{1}{2}$ quantiles of market value. Small Nasdaq stocks have a large number of missing return observations in CRSP data, most due to no trading. I ignore missing returns

Table 12
 Characteristics of monthly returns for individual Nasdaq stocks

	N	$E(R)$	$\sigma(R)$	$E(R^{vw})$	α	β	R^2 (%)
ME < \$2M, arithmetic	22,289	62	175	54	53 (4.0)	0.49	0.18
ME < \$2M, log		-1.9	113		-15 (2.6)	0.40	0.30
ME < \$5M, arithmetic	72,496	37	139	29	27 (1.8)	0.60	0.44
ME < \$5M, log		-5.1	103		-26 (1.3)	0.57	0.77
ME < \$10M, arithmetic	145,077	24	118	16	12 (1.1)	0.76	0.99
ME < \$10M, log		-5.8	93		-31 (0.9)	0.66	1.3
All Nasdaq, arithmetic	776,290	14	81		3.1 (0.3)	0.91	3.1
All Nasdaq, log		-3.4	72.2		-22 (0.3)	0.97	4.6

Note: Mean returns, alphas, and standard deviations are annualized percentages: means and alphas are multiplied by 1,200, and standard deviations are multiplied by $100\sqrt{12}$. $E(R^{vw})$ denotes the value-weighted mean return. α , β , and R^2 are from market model regressions, $R_t^i - R_t^{ib} = \alpha + \beta(R_t^m - R_t^{ib}) + \varepsilon_t^i$ for arithmetic returns and $\ln R_t^i - \ln R_t^{ib} = \alpha + \beta(\ln R_t^m - \ln R_t^{ib}) + \varepsilon_t^i$ for log returns, where R^m is the S&P500 return and R^{ib} is the three-month T-bill return. The sample consists of all Nasdaq stocks on CRSP, January 1987–December 2001, including delisting returns. Each set of rows considers a different upper limit for market value (ME) in month t . Returns are then calculated from month $t + 1$ to month $t + 2$. Missing return data (both regular and delisting) are ignored if the security is still listed the following period, or if the previous period included a valid delisting return. Other missing returns are assumed to be -100%. Log returns are computed ignoring observations with -100% returns. Parentheses present simple pooled OLS standard errors ignoring serial or cross correlation.

when the company remains listed at $t + 3$, or if month $t + 1$ is a delisting return. I treat all remaining missing returns as -100%, to give the most conservative estimate possible. This sample is different than the CRSP deciles in several respects. First, the cutoff for inclusion is a fixed dollar value rather than a decile in the selection month; as a result the numbers and fraction of the Nasdaq in each category fluctuate over time. Second, I rebalance each month rather than once per year. I make both changes in order to better control the characteristics of the sample. With 100% standard deviations, stocks do not keep their capitalizations for long.

The estimates in Table 12 for the smallest Nasdaq stocks are surprisingly similar to the venture capital estimates. First, the mean arithmetic return of the smallest Nasdaq stocks is 62%, comparable to the 59% mean return in the baseline venture capital estimates of Tables 3 and 4. As we increase the size cutoff, mean returns gradually decline. The mean return of all Nasdaq stocks is only 14.2%, similar to the S&P500 return in this time period. Value-weighted averages are lower, but still quite high—the basic result is not a feature of only the smallest stocks in each category.

Second, these individual stock returns are very volatile. The smallest Nasdaq stocks have a 175% annualized arithmetic return standard deviation, even larger than the 107% from the baseline venture capital estimate of Table 3. The \$5 million and \$10 million cutoffs produce 139% and 118% standard deviations, quite similar to the 107% of Table 3. Even stocks in the full Nasdaq sample have a quite high 80.7% standard deviation.

Third, as in the venture capital estimates, large arithmetic mean returns come from the large volatility of log returns, not a large mean log return. The mean and

standard deviation of the small stock log returns are -1.9% and 113% , comparable to the 15% and 89% venture capital estimates in Table 3.

Last, and most important, the simple market model regression for the smallest Nasdaq stocks leaves a 53% annualized arithmetic alpha, even larger than the 32% venture capital alpha of Table 3. This is a feature of only the very smallest stocks. As we move the cutoff to $\$5$ and $\$10$ million, the alpha declines rapidly to 27% —comparable to the 32% of Table 3—and 12% , and it has disappeared to 3% in the overall Nasdaq. The slope coefficients β are not large at 0.49 , gradually rising with size to 0.91 . These values are smaller than the 1.9 β of Table 3. As in Table 3, the log market models leave substantial negative intercepts, here -15% declining with size to -22% , similar to the -7% of Table 3. As in Table 3, the arithmetic alpha is induced by volatility, not by a large intercept in the log market model. The R^2 are of course tiny with such large standard deviations, though the higher R^2 for the log models suggest that they are a better statistical fit than the arithmetic market models.

The finding of high arithmetic returns and alphas in very small Nasdaq stocks is unusual enough to merit a closer look. In Table 13, I form portfolios of the stocks in the samples of Table 12. By the use of portfolio average returns, the standard errors control for cross correlation that the pooled statistics and regressions of Table 12 ignore. In the first panel of Table 13 we still see the very high average returns in the small portfolios, declining quickly with size. Despite the large standard deviations, the mean returns appear statistically significant.

The large average returns (and alphas) are not seen in the CRSP small decile. The CRSP small Nasdaq decile is comparable to the $\$10$ million cutoff. To see the high returns, you have to look at smaller stocks and control the characteristics more tightly than annual rebalancing in CRSP portfolios allows.

In the second panel of Table 13, I run simple market model regressions for these portfolios on the S&P500 return. The market model regressions of the *individual* stocks in Table 12 are borne out in *portfolios* here. The smallest stocks show a 61.6% arithmetic alpha, and the second portfolio still shows a 31.6% arithmetic α . Though the standard errors are greater than those of Table 12, the alphas are statistically significant. By contrast, the relatively much smaller 12% α of the CRSP first (small) decile is not statistically significant. Again, the behavior of my smallest group is not reflected in the CRSP small decile.

Is the behavior of very small stocks just an extreme size effect, explainable by a large beta on a small stock portfolio? In the third panel of Table 13, I run regressions with the CRSP decile 1 on the right-hand side. Alas, this hope is not borne out. Though regressions of the small stock portfolios on the CRSP decile 1 return give substantial betas, up to 1.4 , they also leave a substantial alpha. The alpha is 43% in the $\$2$ million or less portfolio, declining quickly to 18% in the $\$5M$ portfolio and vanishing for larger portfolios. The small Nasdaq portfolio returns lose the correlation with the CRSP small Nasdaq decile, also suggesting something more than an extreme size effect.

Perhaps these very small stocks are “value” stocks as well as “small” stocks. The final panel of Table 13 runs a regression of the small Nasdaq portfolios on the three Fama–French factors. Though the SMB loading is quite large, up to 1.9 , nonetheless

Table 13
 Characteristics of portfolios of very small Nasdaq stocks

	CRSP Dec1	Equally weighted, ME <				Value weighted, ME <			
		\$2M	\$5M	\$10M	\$50M	\$2M	\$5M	\$10M	\$50M
$E(R)$	22	71	41	25	15	70	22	18	10
s.e.	8.2	14	9.4	8.0	6.2	14	9.1	7.5	5.8
$\sigma(R)$	32	54	36	31	24	54	35	29	22
$R_t - R_t^{tb} = \alpha + \beta \times (R_t^{S\&P500} - R_t^{tb}) + \varepsilon_t$									
α	12	62	32	16	5.4	60	24	8.5	0.6
$\sigma(\alpha)$	7.7	14	9.0	7.6	5.5	14	8.6	7.0	4.8
β	0.73	0.65	0.69	0.67	0.75	0.73	0.71	0.69	0.81
$R_t - R_t^{tb} = \alpha + \beta \times (\text{Dec1}_t - R_t^{tb}) + \varepsilon_t$									
ρ	1.0	0.79	0.92	0.96	0.96	0.78	0.92	0.96	0.91
α	0	43	18	4.7	-2.7	43	11	-2.3	-5.7
$\sigma(\alpha)$		8.4	3.6	2.1	1.9	8.9	3.5	2.0	2.5
β	1	1.4	1.1	0.9	0.7	1.3	1.0	0.9	0.7
$R_t - R_t^{tb} = \alpha + b \times \text{RMRF}_t + s \times \text{SMB}_t + h \times \text{HML}_t + \varepsilon_t$									
α	5.1	57	26	10	-1.9	55	18	1.9	-7.0
$\sigma(\alpha)$	5.5	12	7.6	5.8	3.5	12	7.3	5.2	2.7
b	0.8	0.6	0.7	0.7	0.8	0.7	0.7	0.7	0.9
s	1.7	1.9	1.6	1.5	1.4	1.8	1.5	1.5	1.3
h	0.5	0.2	0.3	0.4	0.4	0.1	0.3	0.4	0.4

Note: Means, standard deviations, and alphas are annualized percentages. Portfolios are re-formed monthly. Market equity at date t is used to form a portfolio for returns from $t + 1$ to $t + 2$. CRSP Dec1 is the CRSP smallest decile Nasdaq stock portfolio. R^{tb} is the three-month T-bill return. The sample is for the period January 1987 to December 2001.

the smallest portfolio still leaves a 57% alpha, and even the \$5 million portfolio still has a 25% alpha. As expected, the alphas disappear in the larger, and especially value-weighted, portfolios.

In conclusion, it seems that something unusual happens to the very smallest of small Nasdaq stocks in this period. It could well be a period-specific event rather than a true ex ante premium. It could represent exposure to an unusual risk factor. These Nasdaq stocks are small, thinly traded, and illiquid; the CRSP data show months of no trading for some of them. For the purposes of this paper, however, the main point is that alphas of 30% or more *are* observed in *traded* securities with similar characteristics as the venture capital investments.

The small-stock portfolios are natural candidates for a performance attribution of venture capital investments. It would be gratifying if the venture capital investments showed a beta near one on the small-stock portfolios. Then the 50% or more average arithmetic returns in venture capital would be explained, in a performance

attribution sense, by the 50% or more average arithmetic returns in the small-stock portfolios during this time period. Alas, Tables 3 and 4 only found β of 0.5 and 0.2 on the small stock portfolios, and leave substantial alphas. Venture capital betas are poorly measured, so perhaps the β really is larger. However, the only conclusion from the *evidence* is that venture capital shows an anomalously large arithmetic alpha in this period, and the very smallest Nasdaq stocks also show an anomalously large arithmetic alpha. The two events are similar, but they are not the same event.

11. Extensions

There are many ways that this work can be extended, though each involves a substantial investment in programming and computer time, and could strain the stylized facts that credibly identify the model.

My selection function is crude. I assume that IPO, acquisition, and failure are only a function of the firm's value. One might desire separate selection functions for IPO, acquisition, and new rounds at the (not insubstantial) cost of four more parameters. The decision to go public could well depend on the market, as well as on the value of the particular firm, and on firm age or other characteristics. (Lerner (1994) finds that firms are more likely to go public at high market valuations, and more likely to employ private financing when the market is low.) I allow for missing data, but I assume that data errors are independent of value. It might be useful to estimate additional selection functions for missing data, i.e., firms that subsequently go public are more likely to have good round data in the VentureOne dataset, or that firms which go public at large valuations are more likely to have good final valuation data. I do no modeling of the decision to *start* venture capital projects, yet it is clear in the data that this is an endogenous variable.

My return process is simple. The risks (betas, standard deviation) of the firm are likely to change as its value increases, as the breakout by financing round suggests. Multiple risk factors, or evaluation with reference to carefully tailored portfolios of traded securities, are obvious generalizations. I do not attempt to capture cross-correlation of venture capital returns, other than through identifiable common factors. Such residual cross-correlation is of course central to the portfolio question.

More and better data will certainly help. Establishing the dates at which firms actually go out of business is important to this estimation procedure. Many more projects were started in the late 1990s with very high market valuations than before or since. When we learn what happened to the post-crash generation of venture capital investments, the picture might change.

References

- Admati, A., Pfleiderer, P., 1994. Robust financial contracting and the role of venture capitalists. *Journal of Finance* 49, 371–402.
- Berk, J., Green, R., Naik, V., 2004. Valuation and return dynamics of new ventures. *Review of Financial Studies* 17, 1–35.

- Bygrave, W., Timmons, J., 1992. *Venture Capital at the Crossroads*. Harvard Business School Press, Boston.
- Chen, P., Baierl, G., Kaplan, P., 2002. Venture capital and its role in strategic asset allocation. *Journal of Portfolio Management* 28, 83–90.
- Gilson, R., Schizer, D., 2003. Understanding venture capital structure: a tax explanation for convertible preferred stock. *Harvard Law Review* 116, 874–916.
- Gompers, P., Lerner, J., 1997. Risk and reward in private equity investments: the challenge of performance assessment. *Journal of Private Equity*, 5–12.
- Gompers, P., Lerner, J., 2000. Money chasing deals? The impact of fund inflows on private equity valuations. *Journal of Financial Economics* 55, 281–325.
- Halloran, M., 1997. *Venture Capital and Public Offering Negotiation*, third ed. Aspen, Gaithersburg.
- Kaplan, S., Schoar, A., 2003. Private equity performance: returns, persistence and capital flows. *Journal of Finance*, forthcoming.
- Kaplan, S., Strömberg, P., 2003. Financial contracting theory meets the real world: an empirical analysis of venture capital contracts. *Review of Financial Studies*, forthcoming.
- Kaplan, S., Sensoy, B., Strömberg, P., 2002. How well do venture capital databases reflect actual investments? Unpublished working paper, University of Chicago.
- Lee, P., Wahal, S., 2002. Grandstanding, certification, and the underpricing of venture capital backed IPOs. Unpublished working paper, Emory University.
- Lerner, J., 1994. Venture capitalists and the decision to go public. *Journal of Financial Economics* 35, 293–316.
- Ljungqvist, A., Richardson, M., 2003. The cash flow, return and risk characteristics of private equity. NBER Working Paper 9454.
- Ljungqvist, A., Wilhelm, W., 2003. IPO pricing in the dot-com bubble. *Journal of Finance* 58-2, 577–608.
- Long, A., 1999. Inferring period variability of private market returns as measured by σ from the range of value (wealth) outcomes over time. *Journal of Private Equity* 5, 63–96.
- MacIntosh, J., 1997. Venture capital exits in Canada and the United States. In: Halpern, P. (Ed.), *Financing Growth in Canada*. University of Calgary Press, Calgary.
- Moskowitz, T., Vissing-Jorgenson, A., 2002. The returns to entrepreneurial investment: a private equity premium puzzle? *American Economic Review* 92, 745–778.
- Peng, L., 2001. Building a venture capital index. Unpublished working paper, University of Cincinnati.
- Reyes, J., 1990. Industry struggling to forge tools for measuring risk. *Venture Capital Journal*, Venture Economics, Investment Benchmarks: Venture Capital.
- Smith, J., Smith, R., 2000. *Entrepreneurial Finance*. Wiley, New York.
- Venture Economics, 2000. Press release, May 1, 2000 at www.ventureeconomics.com.