

The Quantitative Economics of Venture Capital *

Robert E. Hall

Hoover Institution and Department of Economics

Stanford University

National Bureau of Economic Research

rehall@stanford.edu stanford.edu/~rehall

Susan E. Woodward

Sand Hill Econometrics, Inc.

swoodward@sandhillecon.com sandhillecon.com

January 27, 2007

Abstract

Over the period from 1987 through 2003, outside investors in venture capital received a financial return around 7 percent per year higher than the risk-adjusted benchmark. We measure the excess return as the alpha of the Capital Asset Pricing Model. Founders earned an average of \$21 million from each company that succeeded in attracting venture funding. But founders are generally specialized in their own companies and bear the burden of the idiosyncratic risk. We show that the founders might be willing to sell their interests for only \$1 million at the outset rather than face that risk. Venture capitalists received an average of \$8 million in revenue from each company they backed. We measure all of these figures from a comprehensive body of data on venture-backed companies, covering the entire universe of such companies. To measure the risk-adjusted pure excess return, we develop a version of the Capital Asset Pricing Model that handles the correlation of excess returns with general stock-market returns. We show that venture is riskier than the stock market, with a beta of 1.3.

*We are grateful to Ravi Jagannathan and numerous seminar participants for comments and to Katherine Litvak for data on limited partner terms.

1 Introduction

Venture capital plays a key role in a modern economy. Venture capitalists—the general partners in venture-capital funds—provide funds for entrepreneurs at early stages, long before they gain access to public capital markets and for projects that banks will not finance. Venture provides funding secured only by the general partners’ beliefs that a new business promises to make real money in the future. Though venture is best known in recent years for incubating high-tech companies like Google, its reach is broader, including low-tech companies like Federal Express. The only important alternative to venture funding for a new company with significant capital requirements is partnering with an established corporation. Banks finance new businesses only to the extent they provide collateral such as accounts receivable and physical capital, usually a small part of a startup’s total capital. Most of the assets of a startup are intangibles, ineligible for debt financing.

This paper is about venture, not private equity more generally. We exclude companies that account for the bulk of the value of private equity. These include small businesses with little likelihood of eventual public ownership and companies in buyout funds. Most venture-backed companies are in high technology, including biotechnology.

Venture capitalists are actively involved in their portfolio companies. They hold board seats and advise on hiring and purchases of other key inputs. They sometimes manage the transition from management by founders—often scientists—to management by professional managers. They are actively involved in the transition out of the venture phase, to public trading through an IPO, to acquisition by a larger company, or to liquidation.

Venture capitalists deploy funds raised from limited partners, usually endowments, pension funds, or wealthy individuals, who invest based on the reputations of the general partners. They have no voice in the selection of companies or other activities of the general partners. The limited partners become shareholders of the portfolio companies.

Venture capitalists earn substantial returns on their expertise. They price their services in much the same way as do the managers of hedge funds—they impose an annual charge on the amount invested by the limited partners and take a fraction of the capital gains delivered by successful companies in the fund.

We examine the economics of venture capital from the perspectives of the three major types of players—limited partners who provide the money, founder-entrepreneurs who invest the money and produce the returns, and general partners who organize funds and choose

and monitor portfolio companies.

Our study of the value of the venture process uses the principles of modern financial economics. We ask whether the stream of later returns from an investment competes with alternative uses of the same funds invested at the same time. We measure alpha, the return from venture in excess of the risk-adjusted benchmark. We measure the risk of venture according to the standard Capital Asset Pricing Model, in terms of the covariance of the return to a limited partner's investment with the discount factor derived from the return to a broad portfolio of available investments, expressed as beta.

Our basic finding is that venture investments earned well above the risk-adjusted benchmark over the period we study, 1987 through 2003. We estimate that the limited partners earned about 7 percent above the CAPM risk-adjusted benchmark. The finding of a pure risk-adjusted excess return is statistically unambiguous. Recent returns in venture have not been so high—the success of venture in the late 1990s attracted a flood of money that may have driven returns down to the market level. We do not regard our finding of a positive pure risk-adjusted excess return as a financial anomaly—rather, it reflects the attention that venture is receiving today as a result of its exceptional and probably non-reproducible success over the past two decades.

We find that the beta of venture is about 1.3—venture investments are exposed to significant levels of systemic risk. Returns are high when the market is strong and low when it is weak. Hence the required return—the CAPM benchmark—is high at all times. Adjustment for risk is essential in understanding its returns.

In the analysis of returns to securities traded in thick markets, the measurement of alpha and beta is straightforward and well understood. They are the constant and slope coefficient of the regression of returns from one investment on the returns from a broad market portfolio. Transplanting this procedure to venture involves some special econometric issues. First, pure risk-adjusted excess returns to venture are correlated with return from the broad market, precluding standard CAPM regressions, where the correlation creates a missing-variable bias that overstates alpha. Second, important information is missing for about 60 percent of venture fundings. Companies disseminate information with more enthusiasm and frequency when the values are high than when they are low. Using a unique data source, we develop a direct solution to this sample selection problem.

Founder-entrepreneurs—the shareholders not associated with venture capital—are com-

compensated for their ideas and efforts at an average rate of \$22 million per venture-funded company. Because the founders are specialized in their own companies and often do not have much other wealth, they would be willing to trade a much smaller payment made with certainty for the highly uncertain payoff to a startup. We give an example where the founders would take \$1 million in exchange for the gamble on their company with an expected payoff of \$21 million.

Venture capitalists are the third group of participants whose earnings we study. They were compensated at a rate of \$8 million per company. Like all participants in venture, their compensation was huge for a small number of companies and small for the majority.

2 Earlier Measures of Venture's Financial Characteristics

With respect to the return and risk of venture investments, our work is most closely related to Cochrane (2005). Cochrane examines returns in relation to market returns in the broad equity market. He works in a maximum likelihood framework, which influences a number of choices that differentiate his approach from ours, notably his assumption of a log-normal distribution of returns. Because the time elapsed between valuations is highly endogenous and is correlated with returns, it is difficult to test the assumption of log normality. Cochrane's translation of log-returns back into arithmetic returns rests on the assumption of log-normality. He does not work with arithmetic returns directly in the main part of the paper.

Cochrane puts a great deal of effort into disentangling endogenous timing of valuation events from returns. As he observes, returns tend to be roughly the same whether the time elapsed between valuation events is long or short—they do not rise with elapsed time as they would if valuation events occurred according to an exogenous process independent of value. Annualized returns are extremely high for closely spaced valuation events and low for those distantly separated in time. In Cochrane's framework with log returns, disentangling is essential. By contrast, in our framework based on levels of values, not logs, endogenous timing of valuation events is not an obstacle to a straightforward approach to estimating the return and risk of venture investments.

Cochrane measures the gross returns on venture investments. These are not the returns received by the suppliers of venture capital, the limited partners. Our analysis distinguishes

the receipts of the three claimants on venture companies—the founders, the general partners of venture funds, and the limited partners. Financial valuation is only appropriate for the limited partners, who provide only financial capital. The founders and general partners provide small amounts of capital but mainly contribute their human capital. We study the actual net cash receipts considering all of the features of the limited partners' claims on venture companies, including the charges of the general partners and the preferences that the limited partners enjoy in the distribution of cash in cases where they do not convert their shares into common stock.

Finally, we make use of a much improved body of data on venture investments and outcomes in comparison to Cochrane's data. In addition to six more years of coverage, our data report unfavorable outcomes for a much larger fraction of the companies included in Cochrane's database.

Cochrane's derived alpha, the pure annual excess return over the risk-adjusted benchmark, is 32 percent per year, with a standard error of 9 percent. We believe that this figure is an overstatement of the excess return, which we find to be about 7 percent per year. The overstatement may arise from his assumption of log-normality. Our work makes no parametric assumption about underlying distributions of the unobserved per-period returns.

In his Table 7, Cochrane presents regressions that resemble standard CAPM regressions of asset returns on market returns. He finds an absurdly high alpha of 462 percent over the holding period and a beta of 2.0. Our approach, derived from the standard CAPM, deals with essentially the same concepts, but discounts returns back to the time of investment rather than working with the covariance of venture returns with market returns. We obtain the reasonable value for alpha of around 20 percent over the roughly three-year holding period. We believe that the problem with standard CAPM regressions arises from the endogeneity of the excess return—the standard regression suffers from a severe omitted variables bias by treating alpha as a constant. The method we have developed avoids this bias.

In our opinion, Cochrane's work is an important contribution to describing interesting features of the venture process, especially the endogenous timing of valuation events. We think, however, that we have found a better way to measure the excess return to venture investments.

Jones and Rhodes-Kropf (2004) study quarterly data on the returns to venture-capital funds, for the returns reported by general partners. They use the Fama-French three-factor

CAPM regression model. To deal with the substantial problem of stale valuations of portfolio companies between rounds and before IPOs or other exit events, they include four lagged values of each of the three factors. They report an annual alpha of about 5 percent, but with a standard error of about 4 percentage points. Apparently because of the large standard error, they consider this to be a small alpha, equivalent to zero, despite the high actual estimated value. Our estimate of alpha, more precisely estimated, is within the confidence interval of their estimate. In addition, we estimate an annual one-factor CAPM using the same data source and find an alpha of about 10 percent per year, also with low precision, but consistent with the findings of this paper. We do not believe that an alpha of 5 or 10 percent per year is small.

Kaplan and Schoar (2005) study the returns to venture capital using two metrics popular among venture capitalists, the internal rate of return and the public market equivalent. The latter is the ratio of the exit value of a venture investment to the value at the same time of the same amount of earlier investments in a public stock index, in their case the S&P 500. They frame the problem solved in this paper neatly in the remark, “...we do not attempt to make more complicated risk adjustments than benchmarking cash flows with the S&P 500 because of the lack of true market values for fund investments until the investments are exited.”

Moskowitz and Vissing-Jorgensen (2002) study the return to capital among privately held businesses in general. Their paper is sometimes cited in connection with venture returns. The universe of privately held companies in their study is vastly larger than venture. It includes many small businesses such as dry cleaners that have no connection with venture. Their finding of a substantially negative excess return to capital is not informative about the financial issue of this paper, the risk-adjusted return to venture investments.

With respect to our measures of the rewards to founders and general partners, we are not aware of any earlier research that quantifies the rewards on a per-company basis, the focus of our work.

Our method for measuring excess returns and risk considers the same problem as in the econometric literature on nonsynchronous trading—see Lo and MacKinlay (1990) and earlier papers cited there—but we make a different assumption about the information available to the econometrician. The earlier literature assumes the econometrician observes only reported returns that may be based on trades that occurred at an earlier but unknown time. We

develop an approach based on the observability of trading dates in which the time of trades is endogenous. Cochrane also assumes observability of endogenous trading dates, but uses rather a different econometric approach, as discussed above.

3 The Venture Process

Venture funds invest in developing companies in financing rounds. The standard convention is to designate the first round of venture financing as the A round, the second the B round, and so on. A syndicate of venture funds will provide a few million dollars in early rounds and substantially more in later rounds, for promising companies whose revenues do not cover their operating and development expenditures.

General partners organize venture funds. They recruit financing commitments from limited partners—usually endowments and pension funds—and choose the companies that will receive financing. The limited partners pay into the funds as required when the general partner provides funding. The limited partners receive most of the cash returned by venture investments, except that the general partner retains almost 3 percent per year of the amount invested in companies still in the fund plus 20 to 25 percent of the cash returned to the limited partners above their original investment when a company undergoes an event such as an IPO or acquisition.

Venture funds generally hold convertible preferred shares in their portfolio companies. The preference requires that the funds receive a specified amount of cash back before the common shareholders (usually only the founders) receive any return. In a successful outcome, the venture funds convert their shares to common stock. The primary purpose of the preferences is to prevent the insiders in a company from paying the venture funds out to themselves as dividends, not to improve the expected returns of the venture investors. It is relatively unusual—but not unknown—for the investors to benefit from the preferences. In addition, venture funds may hold non-convertible preferred shares, in which case they receive the preferences as cash even in the best outcomes.

Most of the return to venture investors comes from occasional large gains. Many venture-backed companies expire without returning any cash to investors. The custom of venture capital is to shut companies down before bankruptcy—venture capitalists have long-term relations with the suppliers who might be harmed by an insolvent portfolio company. The largest returns generally come from IPOs, but acquisitions sometimes provide high returns

as well. On the other hand, many acquisitions occur at low prices and are effectively liquidations. A fraction of venture-backed companies remain for many years as stand-alone operations, able to pay their employees out of revenue, but generate no returns for investors.

We will adopt the standard and convenient vocabulary for describing the evolution of the value of a venture-backed company. When a round of funding occurs, the venture syndicate negotiates a price per share with the founders or other management of the company. This price, multiplied by the number of shares outstanding before the new funding, is called the *pre-money value* of the company. The sum of the pre-money value and the amount newly invested is the *post-money value*. The two values together fully describe the financial evolution of the company, without reference to the share prices or the number of shares. The return ratio earned by shareholders is the ratio of the new pre-money value to the previous post-money value. The pre-money value is adjusted by GP fees and preferences in the case of an exit event.

4 Model and Estimation Framework

Our basic approach examines the discounted future value received by investors from a venture investment in relation to the amount invested. In a frictionless capital market—and with the proper approach to discounting—we should find the two to be equal, on the average. We employ the basic logic of the Capital Asset Pricing Model to determine the discounts, though our estimation does not use a CAPM regression. Cochrane (2001), Chapter 6, discusses the relation between the CAPM and stochastic discounting.

4.1 The obstacle to using the CAPM regression

By analogy with standard practice for securities with pricing at fixed intervals, one might consider measuring the alpha and beta of venture capital from a CAPM regression,

$$r_{v,t} = \alpha + \beta r_t + \epsilon_t. \tag{1}$$

Here $r_{v,t}$ is the return in excess of the risk-free rate from the portfolio of venture investments made in period t , and r_t is the return in excess of the risk-free rate from a broad portfolio of stocks, with the same amount invested as was the case for venture and liquidating along the same schedule. The quantity $(1 + r_{v,t})/(1 + r_t)$, without deduction of the risk-free rate, is called the *public market equivalent* and is a widely used metric of performance among venture

capitalists. It is the ratio of the proceeds from a venture portfolio as a ratio to the proceeds from a similarly timed investment in the broad market. It lacks any explicit adjustment for risk. The CAPM regression using the same data might provide such an adjustment— α is the risk-adjusted pure excess return.

An immediate objection to this approach is the endogeneity of the holding periods of the investments in each portfolio. Cochrane (2005) emphasizes this endogeneity and accounts for it explicitly. But simulations demonstrate that the endogeneity of holding periods by itself is not an obstacle. We have run CAPM regressions on experimental data with exit times occurring earlier when returns are higher. We generated the data with $\alpha = 0$. The regressions return alphas of zero and betas that properly reflect the diminution of risk that return-based exits bring about.

Nonetheless, CAPM regressions on holding-period returns give absurd estimates. We earlier noted that Cochrane obtained a gigantic α in this framework. We obtain similar results. The reason is simple and could arise in the standard setting of returns measured at fixed intervals—it is not a special feature of endogenous holding periods. It may be related to that endogeneity, however. Suppose that the pure excess return varies over time:

$$r_{v,t} = \alpha_t + \beta r_t + \epsilon_t. \quad (2)$$

A substantial empirical literature considers this setup, with an auxiliary model showing how α_t varies along with observed exogenous variables. Estimation by regression remains appropriate. We do not have a reliable auxiliary model, and, in any case, it would be driven by endogenous variables, so identification would be problematical.

To see the problem most clearly, consider the rewriting of the CAPM regression as

$$r_{v,t} = \alpha + \tilde{\alpha}_t + \beta r_t + \epsilon_t, \quad (3)$$

where α is the average value of α_t and $\tilde{\alpha}_t = \alpha_t - \alpha$ is the deviation from the average. If the regression has a constant and r_t as right-hand variables, the estimates of α and β are biased because of omitting the variable $\tilde{\alpha}_t$. In particular, a negative correlation of $\tilde{\alpha}_t$ and r_t could explain low values of the estimate of β and correspondingly high values of α .

As we will show, there is a pronounced negative correlation of the pure excess return and the market return. We are able to measure the correlation by using a different approach within the CAPM framework, based on using the CAPM discount factor rather than the CAPM regression.

4.2 The CAPM in discount form

We will start with the one-period-ahead valuation problem. We seek to find the value this period, V , of a random payoff, x , next period. Standard consumer theory implies that the value of next period's income relative to this period's is the marginal rate of substitution, m . Thus we take the present value to be

$$V = E(mx). \tag{4}$$

This approach takes account of risk in the following way: A risky payoff is one that is low if times turn out to be bad and the realization of m is high because future marginal utility is high. Such a payoff receives a lower present value than one with a less perverse correlation with next period's consumption. Because the same stochastic m discounts all payoffs whatever their risks, it is called the *universal stochastic discounter*.

Another approach to forming present values of risky stochastic future payoffs is to apply a risk-adjusted non-stochastic discount factor to the expected payoff. This approach is completely consistent with the one based on the universal discounter. We write

$$V = E(mx) = E(m)E(x) + \text{Cov}(m, x). \tag{5}$$

Here we apply the non-stochastic discount factor $E(m)$ to the expected payoff, $E(x)$ and then make the risk adjustment $\text{Cov}(m, x)$. Because most payoffs are higher in good times than bad, the covariance is usually negative and the risk adjustment lowers the present value below $E(m)E(x)$.

The CAPM takes a stand on measuring the stochastic discounter by linking it to a broad measure of the state of the economy, an index of the entire stock market, and to the risk-free interest rate, r_f . We let r be the return in excess of that rate that investors earn on the broad market portfolio. A universal stochastic discounter associated with the CAPM is

$$m = \frac{a - br}{1 + r_f}. \tag{6}$$

Discounters come in families, so this is not the only one based on r , but it is the convenient form for the purposes of this paper.

To determine the constants a and b , we make the reasonable assumptions that the current prices of risk-free one-period debt (Treasury bills) and of the broad stock index are equal to the present values of their payoffs. For a bill that sells for \$1 today and pays $1 + r_f$ next

period, the pricing condition is

$$E[m(1 + r_f)] = 1. \quad (7)$$

For the stock market, we observe that one can borrow a dollar at the risk-free rate, put the dollar into the stock market, and earn the difference between their payoffs, r , next period. This move should discount to zero:

$$E(mr) = 0. \quad (8)$$

Dybvig and Ingersoll (1982) show that the value of m satisfying both conditions is

$$m = \frac{1}{1 + r_f} \left(1 + \frac{\bar{r}^2}{\sigma^2} - \frac{\bar{r}}{\sigma^2} r \right), \quad (9)$$

where \bar{r} and σ^2 are the mean and variance of r .

Plugging this into the present value equation, we get

$$V = \frac{1}{1 + r_f} \bar{x} - \frac{\bar{r}}{1 + r_f} \frac{\text{Cov}(x, r)}{\sigma^2}. \quad (10)$$

We let β_x be the covariance divided by the variance, which is the coefficient of the regression of x on r . Some algebra shows that

$$V = \frac{\bar{x} - \bar{r}\beta_x}{1 + r_f}. \quad (11)$$

A standard application of this approach deals with a particular stock with a return r_s , so $x = 1 + r_s$. We believe that the stock, though we can buy it for \$1, is actually worth

$$V = 1 + \frac{\alpha}{1 + r_f}. \quad (12)$$

Here α is the one-period excess return needed to explain our belief that it is worth more than \$1. We insert this into the equation above and do some more algebra to find

$$\bar{r}_s = \alpha + r_f + \beta_s \bar{r}. \quad (13)$$

This is the standard CAPM equation showing that the expected return to a stock is the sum of its excess return α , the risk-free rate, r_f , and its beta times the expected return of the broad market over the risk-free rate, \bar{r} . It is a short step from here to replace the two means by their values at t and treat the resulting relationship as a time-series regression.

For the reason discussed earlier, we will not take that step. We will operate directly with the universal discounter of equation (9) in equation (4) to form present values. The reason

for the discussion of the more standard beta-based approach is to establish that our approach has the same foundation and content as that approach—we employ a totally standard CAPM model.

The standard approach finds alpha and beta from the CAPM regression. Instead, we find α from

$$\alpha = (1 + r_f) \frac{E(mx)}{f} - 1, \quad (14)$$

where f is the amount that investors pay in the market to buy the right to receive the stochastic return x next period. Then we find beta from the ratio of the venture return to the return on the broad stock market. In the one-period case, our results would be identical to those from the CAPM regression approach.

We noted earlier that the covariance term in equation (5) is the adjustment for risk. Now we show that the ratio of the covariance for the return ratio of an investment to the covariance for the return ratio for the general market is the CAPM beta. The return ratios are

$$R_v = \frac{x}{f} \quad (15)$$

and

$$R = 1 + r_f + r. \quad (16)$$

We assume that the present value of R_v is $1 + \alpha/(1 + r_f)$ and the present value of R is, by construction, one. The ratio of the covariances is

$$\frac{1 - E(m)E(R_v) - \alpha/(1 + r_f)}{1 - E(m)E(R)}. \quad (17)$$

Recall that $E(m) = 1/(1 + r_f)$. Write $E(R_v) = 1 + \bar{r}_v$. Then the ratio simplifies to

$$\beta_v = \frac{\bar{r}_v - r_f - \alpha}{\bar{r}}, \quad (18)$$

which is the ratio of the excess of the venture return over $r_f + \alpha$ to \bar{r} , the excess of the market return over r_f . Finally, we substitute the CAPM equation, (13), to determine that the ratio is β , which we will call β_v , the beta of venture investments. One way to express the CAPM model is that the ratio of excess returns measures the risk in the sense of beta.

4.3 Application to venture investments

Venture investors make a series of investments, f_1 through f_N , in months t_1 through t_N . Immediately before a round, the pre-money value of the firm is v_i . At time τ , either the

company undergoes an initial public offering, is acquired, or ceases operations, with an exit cash payoff to the investors of x .

We let $s_{i,j}$ be the ownership share of the company attributable to the investment in round i as of round j . The initial ownership share is

$$s_{i,i} = \frac{f_i}{f_i + v_i}. \quad (19)$$

Later rounds dilute the share according to the recursion,

$$s_{i,j} = \frac{s_{i,j-1}v_j}{f_j + v_j}. \quad (20)$$

We study the return from investing v_i in round i and receiving cash $x_i = s_{i,N+1}x$ when the company exits the portfolio. This setup is the second of Cochrane's two approaches. Because investors almost never have a chance to cash out in a later venture round, we measure the present value of the actual cash they receive at the end.

For notational simplicity, we will number all of the investments in the database sequentially as k , running over rounds and companies, and let τ_k be the exit month for investment f_k . All of the rounds for a given company have the same τ_k .

We define the realized discounted value as of time t of a dollar received in month τ , as

$$M_{t,\tau} = m_t \cdots m_{\tau-1}. \quad (21)$$

Because the returns to venture investments made in the same period are correlated, we adopt the approach often used in finance to investigate excess returns, by forming a portfolio. We group together all investments made in a period and treat them as a single investment with payouts at various times in the future. Let K_t be the set of investments made in t . The valuation condition is

$$\frac{E \left(\sum_{k \in K_t} M_{t,\tau_k} x_k \right)}{\sum_{k \in K_t} f_k} = 1 + \alpha_t. \quad (22)$$

Here we have adopted a slightly different way to characterize the excess return than is standard in the one-period CAPM. Our α_t is the proportion of extra present value as of the time of the investment made in period t . In the standard CAPM, α is discounted at the risk-free rate—see equation (12). The difference is trivial in the one-period case, but the standard definition of α gives anomalous results in the venture setting.

We remove the expectation operator and thus incorporate the expectation error in α_t :

$$\alpha_t = \frac{\sum_{k \in K_t} M_{t,\tau_k} x_k}{\sum_{k \in K_t} f_k} - 1. \quad (23)$$

This expression is the observed return for investments made in period t —it is the discounted proceeds divided by the sum of the investments, less one.

If venture had no common stochastic return and enjoyed no excess return, α_t would be zero in each period, t . The time series for α_t shows the common venture-specific stochastic element of return. The standard error of the average of α_t measures the statistical reliability of the estimate of the excess return, α .

Our overall measure of alpha is

$$\alpha = \frac{\sum_k M_{t_k, \tau_k} x_k}{\sum_k f_k} - 1. \quad (24)$$

We use this in the calculation of beta, to be described shortly.

It would be tempting to characterize the pure excess return as $\alpha(\tau_k - t_k)$, so that α would be a monthly rate. But because the time elapsed between funding rounds is highly endogenous and not in the information set of investors at the beginning of the period, this formulation is prohibited. Instead we restate the estimated α at annual rates after estimation by raising one plus the estimated α to the power of the reciprocal of the typical holding period, which is just under three years, and subtracting one.

4.4 Risk

Our objective here is to develop a risk measure analogous to the beta of the standard CAPM.

With return ratio

$$R_{v,k} = \frac{x_k}{f_k}, \quad (25)$$

we write the valuation condition for venture round k as

$$E(M_k R_{v,k}) = 1 + \alpha. \quad (26)$$

Here $M_k = M_{t_k, \tau_k}$. We proceed as before by breaking the expectation into covariance and expectation components:

$$\text{Cov}(M_k, R_{v,k}) + E(M_k)E(R_{v,k}) = 1 + \alpha. \quad (27)$$

Let R_k be cumulative return ratio for the market over the period:

$$R_k = (1 + r_{t_k} + r_{f,t_k}) \cdots (1 + r_{\tau_{k-1}} + r_{f,\tau_{k-1}}). \quad (28)$$

We posit that the venture beta, β_v , describes the ratios of the covariance of the return ratio for venture and the discounter to the covariance of the market return ratio and the discounter over the same holding period:

$$\text{Cov}(M_k, R_{v,k}) = \beta_v \text{Cov}(M_k, R_k). \quad (29)$$

We also have, from the pricing condition for the market return over the round k holding period,

$$\text{Cov}(M_k, R_k) + E(M_k)E(R_k) = 1. \quad (30)$$

Putting these together and solving for beta, we find

$$\beta_v = \frac{E(M_k)E(R_{v,k}) - 1 - \alpha}{E(M_k)E(R_k) - 1}. \quad (31)$$

As in the standard CAPM, beta is the ratio of the return to venture in excess of the risk-free rate to the return to the market in excess of the risk-free rate, measured as $E(M_k)$.

To apply this formula, we form sample counterparts of the expectations using averages of all observations weighted by round funding, f_k . We calculate

$$\bar{M} = \frac{\sum_k M_k f_k}{\sum_k f_k}, \quad (32)$$

$$\bar{R} = \frac{\sum_k R_k f_k}{\sum_k f_k}, \quad (33)$$

$$\bar{R}_v = \frac{\sum_k x_k}{\sum_k f_k}, \quad (34)$$

and then calculate beta as

$$\beta_v = \frac{\bar{M}\bar{R}_v - 1 - \alpha}{\bar{M}\bar{R} - 1}. \quad (35)$$

4.5 Option value

Many discussions of venture capital emphasize options. At each stage, investors have the option to provide more cash to keep a company going or of letting the company expire for lack of cash. A model of the value of a particular company's prospects as of the time of inception would need to deal with the sequence of real options that define those prospects. Our approach is fully consistent with the option character of a venture company. The universal stochastic discounter is, in principle, truly universal. The expected value of the product of the discounter and the payoff from an option provides the value of the option. Specific

option-value formulas, such as Black-Scholes, use the additional structure of a particular type of option to generate useful formulas. We do not pursue the development of any of these formulas. Rather, we take advantage of the universality of the discount rate. Averaging over many companies, each governed by sequences of real options, gives us the present value without delving into the details of the options. Our approach gives the value achieved, on the average, from venture investments even if the decision makers do not make optimal decisions. Our α describes the risk-adjusted return in the real world, including any agency or limited-information obstacles to the optimal exercise of options.

5 Measurement Issues and Data

5.1 Data on venture transactions

Our data are drawn from a variety of sources, including several commercial data vendors. The vendors concentrate on reporting funding and valuations for venture investments and are less likely to report termination events, especially shutdowns and acquisitions at low values. Sand Hill Econometrics has used a wide range of sources to augment coverage of these adverse termination events.

One important source of valuation data is S-1 statements filed by venture-backed companies when they go public. These statements often give a funding history for the company. Because an IPO is a favorable event, the back-filling of round values from S-1s is a source of return-based selection in the data.

Table 1 describes the data. Our general database reports 54,699 funding rounds and 13,049 exit events for 19,434 companies. We omit 6,385 companies with 17,241 funding rounds because the funding occurred in 2001 or later and the companies have not exited. Among the exit values used in the analysis, 1,936 are IPOs, 4,832 are acquisitions, and 6,281 are zero-value exits, of which 3,186 are inferred because the last funding event occurred before 2001.

Of the 37,458 funding rounds included in the analysis, we can infer the venture share directly from the reported value in 14,368 of the rounds. In the remaining 23,090 rounds, we impute the venture share of ownership as described below. For this purpose, we use the second-look database of 1,292 funding rounds where the values (and thus venture shares) are reported for companies with missing valuations in the general database.

	<i>Number</i>
Company basis	
Companies	19,434
Companies omitted because of no exit event	6,385
Exits used in analysis	13,049
IPO	1,936
Acquisition	4,832
Ceased operations with no value	6,281
Companies assigned zero exit value	3,186
Investment round basis	
Funding rounds	54,699
Rounds omitted because of no exit event	17,241
Rounds used in analysis	37,458
Funding rounds with value revealed	14,368
Funding rounds with imputation	23,090
Second look	
Funding rounds	1,292

Table 1. Counts in Database

5.2 Data for the universal discounter

To measure the return from the general stock market, we use the Dow Jones Wilshire 5000 Index (www.wilshire.com). To measure the risk-free rate, we use the one-year Treasury bill rate stated as a monthly rate.

5.3 Accounting for the general partners' charges

The general partner in a venture fund charges a management fee that is not contingent on the outcome of a fund's companies and is usually a percentage of committed or invested funds. In addition, the general partner keeps the carry, a contractual fraction of the positive part of the gain, $x_i - f_i$. We treat the management fee charge as an addition to the amount invested and deduct the carry from the cash proceeds, in forming the return R_i and in calculating the cost of the carry.

Litvak (2004) reports that the present value at a 10-percent discount rate of the management fees in her sample is 13.85 percent of committed capital. She reports (p. 60) the relation between committed and invested capital over the 11-year typical life of a fund. The

management fee as a percent of invested capital implicit in her present value is 2.9 percent. This figure is slightly higher than the standard one of 2.5 percent because some funds impose the fee on committed capital, at least in the early years of the fund's life.

Gompers and Lerner (1999), Table 2, report that the average carry in their sample of venture capital funds was 20.7 percent. The typical fund in their sample was launched in 1986. Litvak (2004) finds an average carry of 24.3 percent for a later sample centered around 1998. We interpolate linearly between these figures. Litvak finds that the carry percentage is applied to exit value less investment and less management fees in all of the fund agreements in her sample.

Litvak also comments on the use of distribution provisions that are in effect additional charges by GPs, but we have not so far incorporated them in our calculations.

5.4 Preferences

The standard financial contract gives venture convertible preferred equity in a company—see Kaplan and Stromberg (2003), Table 1. Venture investors almost always have a claim on liquidation value equal to the amount they invested. In almost three quarters of deals, venture has a greater claim in liquidation (Kaplan and Stromberg (2003), Table 2). In about half of these cases, the preferences include a cumulative dividend and in the other half, venture receives both its original investment and its common-stock value, sometimes with an upper limit on the common-stock value for payment of the original investment. The second form of claim is called participating preferred stock.

We do not observe the preference terms for each of the rounds of investment in our database. Further, we observe unfavorable outcomes—those in the range where the preferences would matter—*after* negotiations may have occurred among the disappointed claimants on a company that altered the preferences. We understand—without better than anecdotal evidence—that venture investors sometimes bargain away their preferences and become common shareholders in order to induce the founders to agree to a disappointing liquidation plan. Our data do not reveal if the cash from a low-value exit is distributed according to the original contracts or whether the parties have bargained to a jointly superior outcome once the bad news arrived.

The reason that a jointly superior bargain is available is that adverse events leave the founders, holding only common stock, with options that are far out of the money because of

the preferences. In this situation, the founders have little incentive to perform the functions needed to recover limited value from a disappointing outcome. They will prefer to continue rolling the dice unless a new deal can be struck that better aligns incentives.

We have carried out a sensitivity analysis that shows that the issue of preferences is not terribly important for our measure of the return to venture investments. We include in the limited partners' returns the extra value they would have received under the simple preference that gives the investor back the amount originally invested if that amount is greater than the value of the investor's common stock. The return from the preference is further limited by the amount of cash available. The claims of each round of investors are subordinated to claims of earlier rounds. On the one hand, we understate returns by excluding any cumulative dividend the investors may receive. On the other hand, we overstate returns to the extent that we impute preference receipts that the investors actually bargained away in order to receive what cash all common shareholders were able to salvage.

We apply the formula

$$[\text{Preference receipts for a round}] = \max(\min([\text{Liquidation value not paid to earlier rounds}] - [\text{this round's share of total liquidation value}], [\text{amount this round invested}] - [\text{this round's share of total liquidation value}]), 0)$$

We find that total preference receipts under this assumption, from exit values that were neither zero nor high enough to make preferences irrelevant, were 4.7 percent of cash paid out to limited partners. This adds about 1.6 percentage points to our estimate of the annual return to venture investments.

5.5 Down rounds and anti-dilution provisions in the finance contract

A down round occurs when the share price or pre-money value in one round is below the previous round. Almost all contracts between venture funds and founders shift value toward venture in the case of a down round, through anti-dilution provisions. Kaplan and Stromberg (2003) find that about a quarter of the contracts have full-ratchet language, meaning that the founders absorb enough of the decline in value to leave the value of venture's ownership at the same level as in the previous round. The other three quarters of contracts have more moderate provisions.

We have not so far found data on the magnitude of the shift in value in a down round to

make any adjustment for this factor. We understate the return to the limited partners and overstate the return to the founders as a result.

6 Imputing Missing Venture Ownership Shares

We do not measure returns from the interim valuations that are sometimes available when investments occur—our returns are measured from investment amounts to exit amounts, both of which are well reported (we discuss this topic in a later section). But the interim valuations are important for our work because they reveal the fraction of ultimate cash that the investors in each round receive.

We impute missing data for the venture shares by combining a standard missing-data approach with a unique body of data that provides a full solution to the problem of selection bias that plagues the imputation of missing data in most applications.

Our second-look database gives full information about valuations obtained from another source for more than a thousand of the financing rounds in the general data. We make the imputations of venture share from the second-look database. Here we know about missing data, from the perspective of the large body of data, and about true shares, from the fully reported data. Thus we can make a direct attack on the selection problem described above. We fit an equation to the actual shares in the second look data for the companies with missing valuations in the general database.

The venture share, s , needs to be tracked through the various rounds of financing as later rounds dilute the ownership of earlier rounds. The calculations require the most recent share,

$$s_{i,i} = \frac{f_i}{f_i + v_i} \tag{36}$$

because the recursion in equation (20) can be written:

$$s_{i,j} = s_{i,j-1}(1 - s_{i,i}). \tag{37}$$

We use the nonlinear logit regression specification

$$s_{i,i} = \frac{1}{1 - \exp(-X_i\delta)}, \tag{38}$$

where X_i is a vector of variables observed even when v_i is missing and δ is the corresponding vector of parameters.

As predictive variables, we use:

- Number of this round
- Cumulative increase in the Wilshire index over the 2 years preceding this round
- Amount raised in this round
- Number of months since the previous round, or zero for first round
- Indicator variable, 1 for first round, 0 for subsequent rounds

Table 2 shows the results of this regression.

<i>Variable</i>	<i>Coefficient</i>	<i>Standard error</i>
Constant	0.0216	(0.1494)
Funding number	-0.0493	(0.0174)
Stock market	-1.0051	(0.1031)
Amount raised	0.0065	(0.0013)
Months after earlier round	0.0078	(0.0036)
First round	0.6511	(0.1007)
Standard error of the regression:		0.154

Table 2. Coefficients for Logit Specification for Venture Share of Ownership in a Funding Round

7 Results

Figure 1 shows the basic data we use in studying the risk and return of venture. It shows annual percent returns for venture and the broad stock market by calendar quarter. The venture return is the ratio of cash ultimately received by limited partners to the amount of cash they invested in all venture companies in that quarter. Thus the returns are forward looking. They are stated at annual rates, but the cash is received, for the average dollar invested, 32 months after the investment. We measure the broad stock market from the Wilshire 5000 index. To put the broad market on the same forward-looking timing, we consider the returns that investors would have made if they had invested the same amount

in the Wilshire as was invested in venture, and then cashed out at the same time that venture paid off.

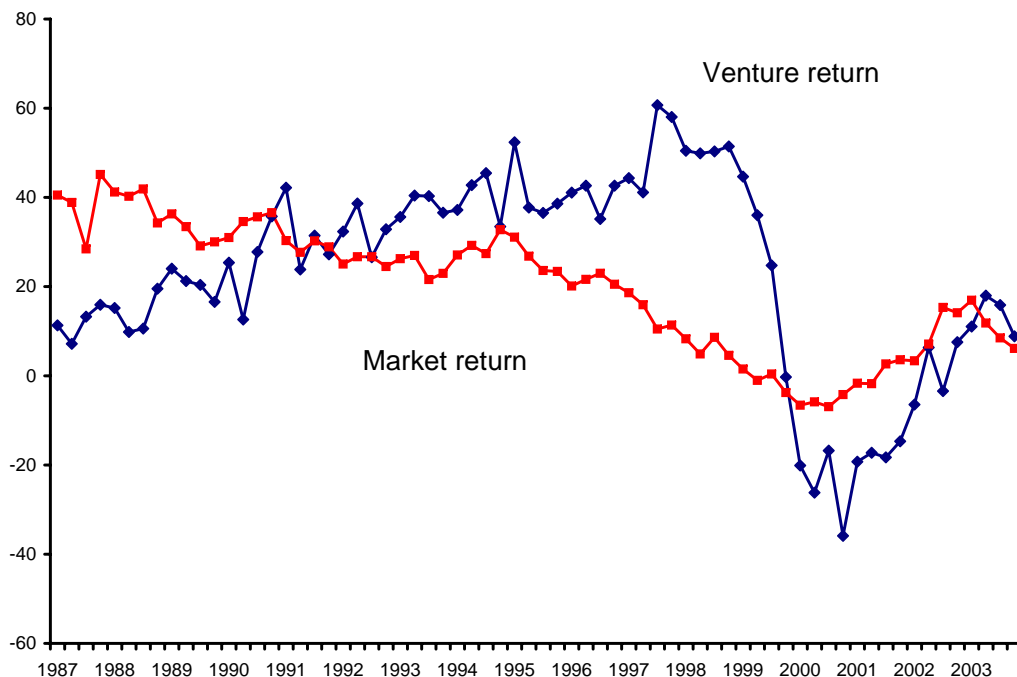


Figure 1. Annual Returns for Venture and for the Broad Stock Market, by Investment Quarter

In most years, venture returned more than the stock market. The gap reached its maximum in the late 1990s—venture was concentrated in the tech firms that enjoyed huge payoffs for investments made until about the end of 1998, which had IPOs or favorable acquisitions before the crash in 2000 and 2001. Venture investments made in the late 1980s and in the years 2000 through 2002 returned less than the stock market.

Figure 1 reveals that venture investment—even the investment diversified across all active venture-stage companies shown in the figure—is risky. When the broad stock market falls, venture falls by a greater proportion. The figure suggests that venture has a beta greater than one.

The method developed earlier in the paper answers the question of whether the extra return venture enjoys adequately compensates investors for the extra risk. The answer is unambiguously yes, for the years taken together. We cannot answer the question of whether the pure excess return to venture that occurred historically will govern the future. Our results show that venture has not had a risk-adjusted advantage over the stock market for almost a decade.

7.1 Return

Figure 2 shows the pure excess return α_t for investments made between the first quarter of 1987 and the last quarter of 2003. We omit later investments because their returns are not yet known. The horizontal line marks the benchmark of returns in financial equilibrium, where the discounted cash received has the same value as the cash paid in. The figure makes it clear that the investors in venture deals—the limited partners—generally beat the financial benchmark. Investments made in most quarters from 1989 through 1999 repaid more than amount they would merit in competition with other potential investments. Venture returns were spectacular during the boom of the late 1990s. Venture investments made in the late 1980s and early in the present decade fell short of the benchmark, however.

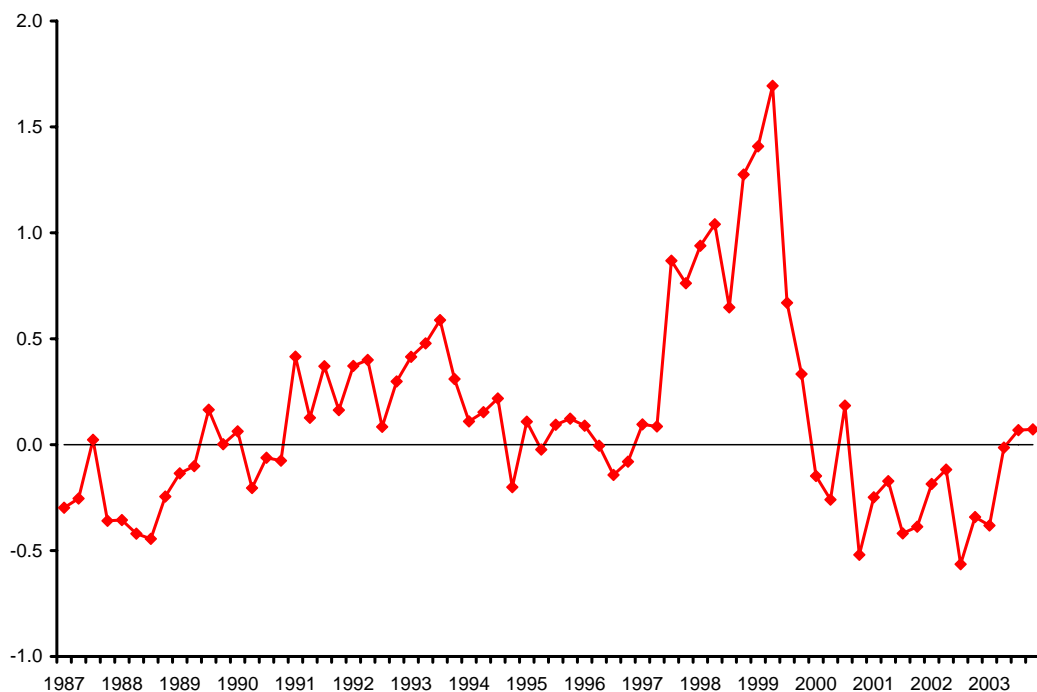


Figure 2. Realized Excess Present Values of Venture Investments, α_t , as Ratios to Amount Invested, by Investment Quarter

The average value of α_t , the pure excess return above the risk-adjusted cost of capital, is 12.0 percent. The standard error of the ratio is 5.5 percentage points—the finding of a positive α is statistically unambiguous. The average time between investment and receipt of cash is 32 months, so the annual risk-adjusted excess return is about 3 percent, with a standard error of 2.1 percentage points. The overall weighted average α is 19 percent.

The variations over time shown in Figure 2 combine the idiosyncratic random performance

of venture with systematic changes in excess returns. There is no doubt that the numbers are generally positive. The long sequence of negative figures in this decade may indicate a departure from the earlier superior performance of venture investments to one where the results, risk adjusted, were inferior to the broad stock market. Since 2000, venture funds have attracted huge amounts of investment, which may have driven down returns.

7.2 Risk

Table 3 calculates venture's beta using the approach described earlier.

<i>Measure</i>	<i>Description</i>	<i>Value</i>
$\alpha = E(MR_v) - 1$	Pure excess return to venture, over holding period	0.19
$E(M)$	Discount averaged over dollars invested	1.26
$E(R)$	Return ratio for the stock market for the timing and amount invested in venture	1.22
$E(R_v)$	Total dollars paid out by venture divided by total dollars invested	1.48
β_v	Venture's beta	1.26

Table 3. Calculation of Venture's Beta

8 Alternative Approaches to Estimation

8.1 The Sand Hill Index

The Sand Hill Index approximates quite closely the cumulative value of a continuously reinvested, value-weighted index of all venture-backed companies. It applies the philosophy of indexes like the S&P 500 and the Wilshire 5000 to the venture universe. Sand Hill constructs the index by interpolating between observed valuations and by using imputations for missing valuations similar to the ones used in this paper for missing venture ownership shares. The firm assigns a value to every known venture-backed company at the end of each month. The percentage change in the value during the month for the companies included at the end of the previous month is taken as the percentage change of the index during that month. The interpolation procedure uses information about publicly traded values during the period between observed valuations of a company. In addition, an extrapolation procedure

assigns values from the last observed valuation to the present, taking account of the declining value that generally occurs among companies after a few years without a valuation event. Although the index uses mainly estimated values, its basic properties are controlled by the same well-observed data used in this paper—the amounts paid in by venture and the cash received from favorable exit events.

Absent some special properties relating to the interpolation in constructing the index, we could carry out a standard monthly CAPM regression to find the alpha and beta of venture. None of the special treatment of endogenous holding periods developed earlier in this paper would be needed. But the index results in a slightly blurred view of venture performance because of the episodic valuations it rests upon. We deal with this issue by aggregating to annual time periods.

8.2 Venture Indexes from Venture Economics and Cambridge Associates

Venture Economics and Cambridge Associates produce indexes based on returns from funds. They rely on the values assigned by GPs to the companies remaining in the funds. Generally these values are kept at the last round value until a new round occurs or the company exits. Thus the valuations tend to be somewhat stale. Neither firm provides information about weighting. The returns are net of the charges of the GPs. Both indexes are published quarterly. Coverage of funds is probably not complete and may be biased toward funds enjoying higher returns.

8.3 Results obtained from indexes

Figure 3 shows the three available indexes of annual returns to venture along with the annual return to the Wilshire 5000 stock-market index. We aggregate to annual returns by taking the percentage change from year-end to year-end to ameliorate the temporal blurring that occurs in the Sand Hill Index and the use of stale valuations in the other two indexes. All three venture indexes show similar movements over the period for which they overlap.

Table 4 shows annual CAPM regressions for the three venture indexes. They all agree with the earlier findings in this paper, though with relatively low statistical precision. The annual values of α , the risk-adjusted pure excess return to venture, are completely consistent with our earlier estimate of 9 percent, though with larger standard errors. The estimates of

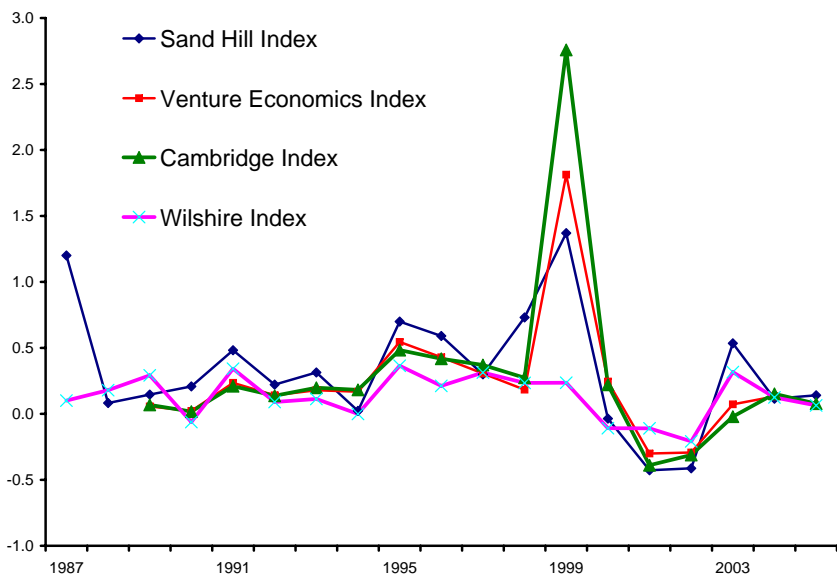


Figure 3. Annual Returns to Venture According to Three Indexes

<i>Index</i>	<i>alpha</i> (percent)	<i>beta</i>	<i>sigma</i> (percent)
Sand Hill	14.6 (9.7)	1.64 (0.54)	38
Venture Economics	10.4 (12.3)	1.08 (0.63)	44
Cambridge Associates	12.8 (17.7)	1.31 (0.93)	65

Table 4. Estimates of Alpha and Beta for Three Indexes of Venture Returns

β are also consistent with our earlier figure of 1.6.

The lower precision of index-based estimates arises from the need to use time-aggregation to overcome their inability to deal directly with timing issues originating in the episodic and endogenous timing of valuations. The method presented earlier in this paper deals directly with that issue.

9 Returns to Founders

Venture-backed companies typically have a single scientist or similar expert, or a small group, who supply the original concept, contribute a small amount of capital, and find venture funds

to supply the bulk of the capital. The founders own all of the shares in the company prior to the first round of outside funding.

We do not apply the CAPM valuation model to the founders' interest, for two basic reasons. First, the founders are specialized in ownership of the venture-stage firm. The CAPM, with its exclusive attention to the non-diversifiable risk arising from the covariance of one company's return with the market, does not apply to the investor who is heavily exposed to the idiosyncratic volatility of a single company. Second, the founders achieve their returns in large part from the human capital they provide to the company. The founders often have built up capital earlier in their careers in developing the idea or technology of the new company. They provide an intensive flow of new thinking and problem-solving as the company develops.

Founders generally receive modest salaries during the venture phase of the development of their companies. We do not have data on salaries. We do measure the value the founders receive when a company exits the venture phase. If the exit is by way of IPO, the founders receive public shares. We measure the value of those shares from the IPO price. In cases where the IPO price understates the actual market value, we understate the gains of founders correspondingly, because they almost never sell shares in the IPO, but rather are locked up by contract for a period of generally six months. The understatement is material only for a brief period in 1999 and 2000. Historically and today, IPOs are not generally underpriced.

Figure 4 shows the average exit value received by founders, in millions of 2006 dollars, by the date of the exit. Prior to 1999, the figure was quite stable at \$20 million. It jumped to the \$70 million level in 1999 and 2000, then plunged to around \$5 million until 2003. Google's exit in August 2004 resulted in a spike. Most recently, the founders' take has risen back to its normal level of around \$20 million. The average over the period was \$22 million in 2006 dollars.

Table 5 shows the distribution of founders' exit value in millions of 2006 dollars. About 68 percent of companies yield no exit value to founders. A large fraction of total value to founders arises from the tiny fraction of startups that deliver hundreds of millions of dollars of exit value to founders. The contract between venture capital and venture founders does essentially nothing to alleviate the financial extreme specialization of founders in their own companies. Given the nature of the gamble revealed in Table 5, founders would benefit by selling some of the value that they would receive in the best outcome in the bottom line,

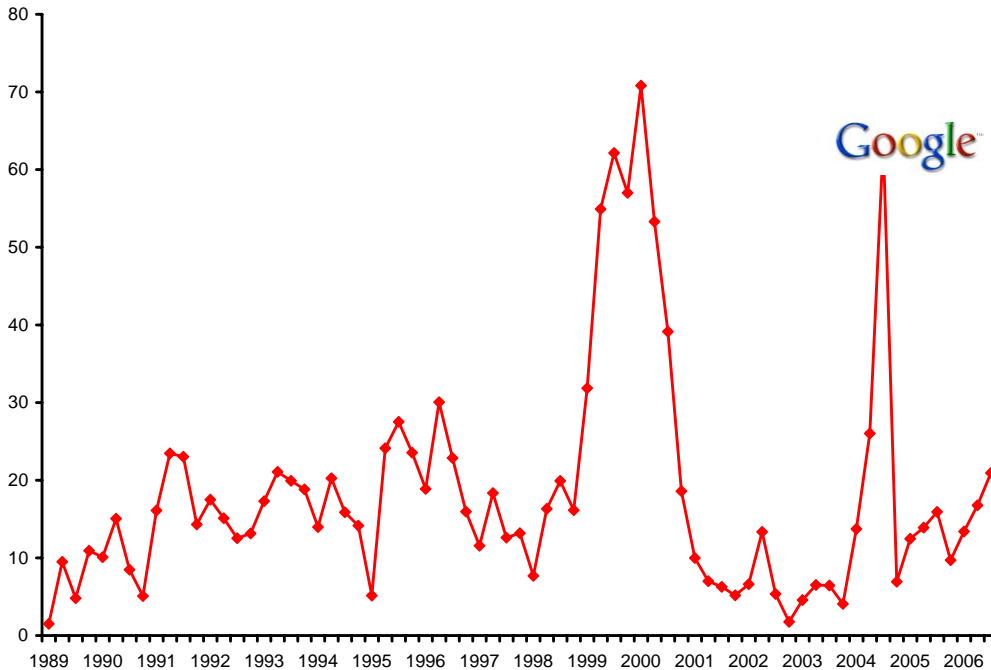


Figure 4. Average Exit Value to Founders per Venture-Funded Company, Millions of 2006 Dollars, by Quarter of Exit

when they would be seriously rich, in exchange for more wealth in the more likely outcomes in the top lines. A diversified investor would be happy to trade this off at a reasonable price, given that most of the risk is idiosyncratic and diversifiable. But venture capitalists will not do this—they don't buy out startups at the early stages and they don't let founders pay themselves generous salaries. They use the exit value as an incentive for the founders to perform their jobs. Moral hazard and adverse selection bar the provision of any type of insurance to founders—they must bear the huge risk shown in Table 5.

The specialization forced upon the founders by information limitations substantially lowers the incentive to be a founder. We can solve the equation,

$$u(\bar{W} + \tilde{W}) = \frac{1}{N} \sum_k u(\bar{W} + W_k), \quad (39)$$

to find the non-stochastic addition to wealth, \tilde{W} , that has the same value to the founders as the actual distribution in our data, W_k . \bar{W} is the founders' pre-existing level of wealth, including human capital. We assume that the founders have constant relative risk aversion preferences,

$$u(W) = \frac{W^{1-\gamma}}{1-\gamma}, \quad (40)$$

<i>Founders' exit value, millions of 2006 dollars</i>	<i>Percent of companies</i>	<i>Percent of total founders' exit value</i>
0 to 1	68.3	0.0
1 to 10	8.3	2.0
10 to 50	12.2	15.2
50 to 100	5.1	18.0
100 to 200	3.5	24.1
200 to 500	2.1	29.8
500 to 1000	0.4	11.9
1000 and hi	0.2	14.3

Table 5. Distribution of Exit Value to Founders per Venture-Funded Company, Millions of 2006 Dollars

with a coefficient of relative risk aversion, γ , of 2, a figure consistent with the evidence about this aspect of preferences. We take the pre-existing level of wealth to be \$3 million. The resulting value of \tilde{W} is \$0.99 million. It takes the prospect of \$21 million of expected wealth to generate only \$1 million of effective wealth, given the curvature of preferences. The typical group of founders would be willing to sell their company for \$1 million at the moment when the first venture capitalist agrees to invest in the company.

10 Earnings of General Partners

Gompers and Lerner (2004) discuss the services that general partners provide to the venture-backed companies in their portfolios and to their limited partners. These include raising the funds from limited partners, screening proposals from founders, and supervising the development of companies in their portfolios. Because most venture funding is syndicated, a venture-backed company receives these services in principle from a number of venture-capital firms, though in practice a lead firm provides most of them.

The general partners receive compensation for their services through their expense charges

of just under three percent per year of the funds invested in their portfolios plus the carry, which generates cash receipts of somewhat over 20 percent of the capital gain from each exiting company. Figure 5 shows the average earnings from these sources as of the exit date, when the GP receives the carry (the expenses are earned earlier, during the period averaging about three years between investment and exit).

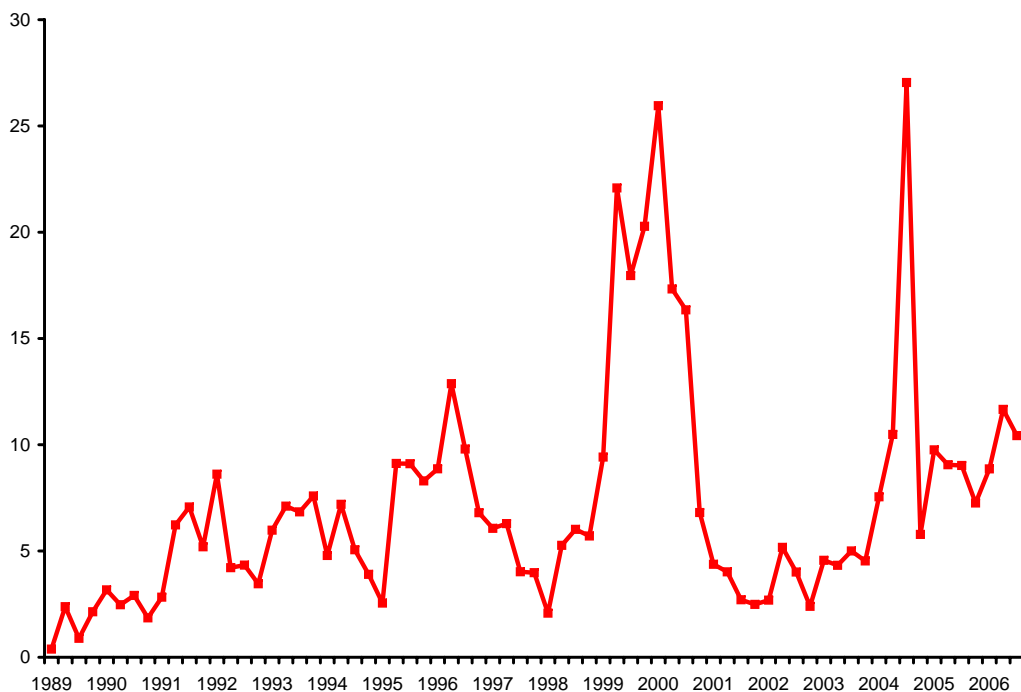


Figure 5. Average Exit Value to General Partners, per Venture-Funded Company, Millions of 2006 Dollars, by Quarter of Exit

GPs' earnings per company were stable at about \$8 million until the boom of 1999, fell to about \$3 billion in 2001 through 2003, spiked with Google, and then rose back to the \$10 million level. The average over the period was \$8.3 million.

11 Concluding Remarks

The venture capital institutions of the United States convert ideas into functioning businesses. The founders, equipped with ideas and willing to work hard for little current return, take their plans to venture capitalists. A fraction of them receive funding, at which point they enter our data. Once funded, the founders will receive \$21 million on the average from the ultimate sale or public listing of their company. Venture capitalists, who screen many

proposals for each one they fund, then help guide the company as it develops, receive \$8 million per company for their services.

Large pools of funds are potentially available to the companies that qualify for venture financing. Our results are compatible with the view that the managers of these funds—from endowments, pension funds, other financial institutions, and wealthy individuals—provide them to venture capitalists perfectly elastically at the risk-adjusted rate of return. We find that the average return over the past 20 years has exceeded the risk-adjusted benchmark by the considerable margin of 7 percent per year, but recent results have been closer to the benchmark.

References

- Cochrane, John H., *Asset Pricing*, Princeton University Press, 2001.
- , “The Risk and Return of Venture Capital,” *Journal of Financial Economics*, 2005, 75, 3–52.
- Dybvig, Philip H. and Jonathan E. Ingersoll, “Mean-Variance Theory in Complete Markets,” *Journal of Business*, 1982, 55 (2), 233–251.
- Gompers, Paul and Josh Lerner, “An Analysis of Compensation in the U.S. Venture Capital Partnership,” *Journal of Financial Economics*, 1999, 51, 3–44.
- and —, *The Venture Capital Cycle*, second ed., Cambridge, Massachusetts: MIT Press, 2004.
- Jones, Charles M. and Matthew Rhodes-Kropf, “The Price of Diversifiable Risk in Venture Capital and Private Equity,” July 2004. Graduate School of Business, Columbia University.
- Kaplan, Steven N. and Antoinette Schoar, “Private Equity Performance: Returns, Persistence, and Capital Flows,” *Journal of Finance*, August 2005, 60 (4), 1791–1823.
- and Per Stromberg, “Financial Contracting Theory Meets the Real World: An Empirical Analysis of Venture Capital Contracts,” *Review of Economic Studies*, 2003, 70, 281–315.
- Litvak, Kate, “Venture Capital Limited Partnership Agreements: Understanding Compensation Arrangements,” May 2004. University of Texas Law School.
- Lo, Andrew W. and A. Craig MacKinlay, “An Econometric Analysis of Nonsynchronous Trading,” *Journal of Econometrics*, 1990, 45, 181–211.
- Moskowitz, Tobias J. and Annette Vissing-Jorgensen, “The Return to Entrepreneurial Investment: A Private Equity Premium?,” *American Economic Review*, September 2002, 92 (4), 745–778.