

NBER WORKING PAPER SERIES

THE INCENTIVES TO START NEW COMPANIES:
EVIDENCE FROM VENTURE CAPITAL

Robert E. Hall
Susan E. Woodward

Working Paper 13056
<http://www.nber.org/papers/w13056>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2007

We are grateful to Ravi Jagannathan and numerous seminar participants for comments and to Katherine Litvak for data on limited partner terms. Hall's research is part of the Economic Fluctuations and Growth Program of the National Bureau of Economic Research. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Robert E. Hall and Susan E. Woodward. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Incentives to Start New Companies: Evidence from Venture Capital
Robert E. Hall and Susan E. Woodward
NBER Working Paper No. 13056
April 2007
JEL No. G12,G24,G32,L14

ABSTRACT

The standard venture-capital contract rewards entrepreneurs only for creating successful companies that go public or are acquired on favorable terms. As a result, entrepreneurs receive no help from venture capital in avoiding the huge idiosyncratic risk of the typical venture-backed startup. Entrepreneurs earned an average of \$9 million from each company that succeeded in attracting venture funding. But entrepreneurs are generally specialized in their own companies and bear the burden of the idiosyncratic risk. Entrepreneurs with a coefficient of relative risk aversion of two would be willing to sell their interests for less than \$1 million at the outset rather than face that risk. The standard financial contract provides entrepreneurs capital supplied by passive investors and rewards entrepreneurs for successful outcomes. We track the division of value for a sample of the great majority of U.S. venture-funded companies over the period from 1987 through 2005. Venture capitalists received an average of \$5 million in fee revenue from each company they backed. The outside investors in venture capital received a financial return substantially above that of publicly traded companies, but that the excess is mostly a reward for bearing risk. The pure excess return measured by the alpha of the Capital Asset Pricing Model is positive but may reflect only random variation.

Robert E. Hall
Hoover Institution
Stanford University
Stanford, CA 94305-6010
and NBER
rehall@gmail.com

Susan E. Woodward
115 Everett Ave.
Palo Alto, CA 94301
swoodward@sandhillecon.com

1 Introduction

Startup companies involve the collaboration of entrepreneurs, financial organizers, and suppliers of capital. We measure the incentives provided in the setting of venture capital to each of the three groups. In that setting, the entrepreneurs are scientists, engineers, and executives, the financial organizers are venture capitalists (general partners of venture funds) and the suppliers of capital are the limited partners of venture funds.

Our most important finding is that the reward to the entrepreneurs who provide the ideas and long hours of hard work in these startups is remarkably small, once risk is taken into consideration. The contract between venture investors and entrepreneurs imposes the burden of the idiosyncratic risk of a startup on its entrepreneurs. Far from shifting the risk toward the diversified investors, the contract exacerbates the risk the entrepreneurs face. Although the average ultimate cash reward to the entrepreneurs of a company that succeeds in landing venture capital is \$9 million for the group, most of this expected value comes from the small probability of a great success. With a coefficient of relative risk aversion of 2, the entrepreneurs would sell their holding as of the time they receive venture funding for about \$900,000 to avoid the undiversified risk of their claim on the company.

We also quantify the earnings of venture capitalists, the general partners in venture funds who make the decisions about startup investments and about further investments as companies develop. 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 arrange the transition from management by entrepreneurs—often scientists—to management by professional managers. They manage the transition out of the venture phase, to public markets through an IPO, to acquisition by a larger company, or to liquidation. 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 the fund. We find that this revenue is about \$5 million per venture-backed company. Although the revenue has similar idiosyncratic risk at the company level as does the entrepreneurial reward, venture capitalists are far more diversified.

Venture capitalists deploy funds raised from limited partners, usually pension funds, endowments, and wealthy individuals, who invest based on the reputations of the general

partners. The limited partners have no voice in the selection of companies or other activities of the general partners. The limited partners become shareholders of the portfolio companies. We quantify their earnings using the principles of modern financial economics. We ask whether the later returns from an investment compare 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 general stock market, expressed as beta. Our basic finding is that venture investments earned about the risk-adjusted benchmark over the period we study, 1987 through 2005, although our estimate of alpha—two percent per year—implies an important reward to venture investing, the confidence interval for the estimate of alpha includes zero comfortably.

We find that the beta of venture is about 1.7—venture investments are exposed to significant levels of systematic risk. Returns are high when the stock market is strong and low when it is weak. The required return—the CAPM benchmark—reflects the systematic risk. Adjustment for risk is essential in understanding the returns to venture-backed companies.

In addition to the three main claimants on venture-backed companies, our quantification includes the payouts to angel and other non-venture investors and to employees who receive stock options as part of their compensation.

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. We also exclude startups financed by the receipts of entrepreneurs in earlier successful startups—our data are drawn from companies who received venture funding.

Venture capital plays a key role in a modern economy. Venture is a unique source of 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 assets. Most of the assets of a startup are intangible, ineligible for debt financing.

With respect to our measures of the rewards to entrepreneurs 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. We discuss the extensive earlier research on the financial returns to venture investors in a later section.

2 The Startup Process

At the outset, startups are usually operated and financed by the entrepreneurs themselves. Others—often friends and family—may invest as founding common shareholders as well. These investors are called angels. Unless the founders are wealthy, they need outside financing. A main task early in a startup is to find investors. Two primary types of investors are corporate partners and venture investors. Our concern is with the companies that succeed in obtaining venture funding by convincing some venture capitalists that the new business has a prospect of being highly profitable.

Venture funds invest in developing companies in financing rounds. 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 pension funds, endowments, and wealthy individuals—and choose the companies that will receive financing. 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 preferred shares in their portfolio companies. The preference requires that the funds receive a specified amount of cash back before the common shareholders (the entrepreneurs, angels, and employees) receive any return. In the majority of cases, the preferred shares are convertible. In a successful outcome, the venture funds

convert these shares to common stock. In addition, venture funds may hold non-convertible preferred shares, in which case they receive the preferences as cash even in the best outcomes.

A huge literature portrays the standard venture financial contract as the constrained optimum of a challenging mechanism design problem. This research explains key features, including the assignment of a share of the ultimate value to the entrepreneurs, multiple stages of financing, and debt instruments (preferences) that convert to equity. Some of the more prominent contributions include Admati and Pfleiderer (1994), Schmidt (2003), Casamatta (2003), and Repullo and Suarez (2004). Wilmerding (2003) and Bagley and Dauchy (2003) explain the terms of venture contracts from the perspective of venture capitalists and their lawyers.

The dominant factor in this literature is moral hazard. Venture investors and their agents, the general partners of venture funds, are unable to monitor the efforts of entrepreneurs to commercialize their ideas. Consequently, the entrepreneurs are paid in proportion to the actual commercial success of their companies. This alignment of incentives comes at the cost of a substantial diminution in the strength of the incentive because of the idiosyncratic risk that entrepreneurs are unable to insure. Alternatives with less risk, such as paying entrepreneurs salaries in place of equity, apparently provide such weak incentives that the relationship based on equity incentives weakened by idiosyncratic risk is still optimal.

Most of the return to venture investors comes from occasional large gains. Most venture-backed companies expire without returning any cash to owners. 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. Some venture-backed companies remain for many years as stand-alone operations, able to pay their employees out of revenue, but generating no returns for shareholders.

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 entrepreneurs 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.

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

$$s_{i,i} = \frac{f_i}{f_i + v_i}. \quad (1)$$

Later rounds dilute the share according to the recursion,

$$s_{i,j} = \frac{s_{i,j-1}v_j}{f_j + v_j}. \quad (2)$$

The exit value of round i investors is $x_i = s_{i,N+1}x$. Appendix 1 describes the modifications in this recursion that apply when the value of a company declines from one round to the next, thanks to anti-dilution provisions benefiting investors who paid more than the current share price in an earlier round.

3 Measurement Issues and Data

3.1 Data on venture transactions

We use a database compiled by Sand Hill Econometrics on venture investments in startups and on the fates of venture-backed companies. The 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 exit 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 55,307 funding rounds and 19,434 exit events. Among the exit values used in the analysis, 1,936 are IPOs, 4,832 are acquisitions, and 3,095 are confirmed zero-value exits. We impute positive exit values at the end of the sample in 2006 for 9,571 companies, of which 8,787 have positive imputed values and 784 are taken as zero-value exits because the last funding event occurred before 1992. Many of the companies with positive imputed interim values will ultimately expire without shareholder value—our positive imputation is the working of the general principle that a forecast has less dispersion than the actual outcome.

	<i>Number</i>
Exits	19,434
IPO	1,936
Acquisition	4,832
Confirmed to have ceased operations with no value	3,095
Imputed because of lack of observed exit	9,571
With positive value	8,787
With zero value (more that 15 years since last funding)	784
Funding rounds	55,307
With ownership share from new investment revealed	14,396
With imputation of ownership share	40,911
Second look rounds (all with value revealed)	1,292
With value not revealed in main data	762

Table 1. Counts in Database

Of the 55,307 funding rounds included in the analysis, we can infer the venture share directly from the reported value in 14,396 of the rounds. In the remaining 40,911 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. That sub-sample contains 762 rounds.

3.2 Data for the CAPM

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 at a monthly frequency, we use the monthly average of the LIBOR rate, and, at a quarterly frequency, the return on 3-month

Treasury bills.

3.3 Accounting for the general partners' charges

The general partner in a venture fund charges a management fee that 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 to the limited partners 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 rather than invested 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.

The carry is imposed on funds, not separately on companies in funds. The general partners do not receive the carry from a successful company unless the fund reaches the point of repaying the limited partners' original capital. We do not have information about fund ownership of the companies in our data. The practice of syndication of venture investments would make analysis of funds at the company level challenging, as a number of funds typically invest in each round of each company.

To account for the fact that carry does not apply to successful companies in unsuccessful funds, we carry out a simulation study. In each year, we form 1000 synthetic funds, each of which invests in around 100 rounds drawn at random for that year from our data. Specifically, the number of rounds in a fund is drawn from a uniform distribution running from 10 to 190 and the fraction of a round held by the fund is drawn from a uniform distribution running

from 0.01 to 0.27. Thus the average ownership fraction is 0.14, corresponding to 7 funds investing in a round.

For each year, we calculate the carry at the fund level summed over all funds and the carry at the round level for the rounds owned by those funds. The ratio of the two is the adjustment factor we apply to calculate the effective carry from the notional round-level carry. Table 2 shows the results of the calculation. The adjustment factor is close to one in years of high returns, but in the years of lowest returns—investments made in 2001 being the worst—the adjustment factor is small and the effective carry is low.

<i>Year</i>	<i>Fraction of gains occurring in funds with fund-level gains</i>
1987	0.72
1988	0.70
1989	0.73
1990	0.70
1991	0.79
1992	0.80
1993	0.81
1994	0.84
1995	0.83
1996	0.83
1997	0.85
1998	0.83
1999	0.56
2000	0.21
2001	0.08
2002	0.56
2003	0.79
2004	0.76
2005	0.78
2006	0.93

Table 2. Factor for Adjusting the Carry to Account for Companies in Unsuccessful Funds, Based on Simulation of Synthetic Funds

3.4 Preferences

The standard financial contract gives venture convertible preferred equity in a company—see Kaplan and Stromberg (2003), Table 1. In about half of these cases, 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 that altered the preferences may have occurred among the disappointed claimants on a company. We understand that venture investors sometimes bargain away their preferences and become common shareholders in order to induce the entrepreneurs to agree to a disappointing exit plan. Wilmerding (2003) writes, in connection with a low potential exit value, “. . .at those levels, where management will not receive much from a sale, the preferred shareholders will likely be forced to give up some of their return to make the deal work” (p. 52). 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 entrepreneurs, holding only common stock, with option positions that are far out of the money because of the preferences. In this situation, the entrepreneurs 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.

In some cases, one round of preferred stock has priority over another, but we lack information on priorities, so we assume that all rounds of preferred shares have equal priority and divide the available cash in proportion to the amount invested. Kaplan and Stromberg (2003), Table 2, report that 71 percent of venture contracts grant preferences in excess of the amount invested, often in the form of a cumulative dividend, which averages 8 percent per year. We take the preference to be 125 percent of the amount invested, corresponding to about 3 years of the dividend.

For non-participating preferred shares, the cash payout is

$$\max(\min(\bar{s}X, \bar{s}P), sX), \tag{3}$$

where \bar{s} is the share of a given venture round among all venture (preferred) shareholders and s is the share of the round among all shareholders, X is the exit value of the company, and P is the preference amount for the round. For participating shares, the payout is

$$\min(\bar{s}X, \bar{s}P + s(X - P)). \tag{4}$$

We find that total preference receipts under these assumptions were 9.3 percent of cash paid out to limited partners. This adds 13.5 percentage points to our estimate of the holding-period return to venture investments.

3.5 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 entrepreneurs 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 entrepreneurs 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 a more moderate provision called weighted-average adjustment.

Appendix 1 describes the two types of anti-dilution adjustments. Both are modifications of equation (2) to shift ownership shares toward venture investors who paid more than the current price for their shares, where the price paid is measured on a post-conversion basis. We calculate updated ownership shares for down rounds using both types of adjustment and take the weighted average, using the figures from Kaplan and Stromberg (2003).

Down-round anti-dilution provisions shift venture ownership upward and non-venture (entrepreneurs, angels, and employees) downward by an average of 4.8 percentage points. These provisions have an important role in our central finding of high risk to entrepreneurs, because they tend to reduce the reward to entrepreneurs disproportionately in the less favorable outcomes.

3.6 Additional ownership for entrepreneurs and employees

The contract between venture investors and entrepreneurs often includes provisions for additional ownership based on company performance. The entrepreneurs vest in the shares upon reaching milestones in the contract. In addition, non-founder employees vest in stock options based on longevity and other factors. Kaplan and Stromberg (2003), Table 2, report on both elements of vesting in terms of ownership shares if no vesting occurs and if all vesting occurs. The average founders' share rises from 24.3 percent to 31.1 percent upon full vesting

and the average share of non-founder and employee ownership rises from 20.2 percent to 22.2 percent.

We interpret the initial 20.2 percent as the share of non-entrepreneur, non-venture investors—angels and friends of the family. We interpret the initial non-venture share observed in our data as the sum of the entrepreneurs’ share and the 20.2 percent. We interpret the incremental 2.0 percent as the ownership of non-founder employees upon a successful exit such as an IPO. We interpret the incremental 6.8 percent of entrepreneur ownership in the same way.

We construct an index of success, with respect to vesting, as

$$z = \min \left(1, \frac{X}{\psi F} \right), \quad (5)$$

where X is the gross value to all existing shareholders at exit and F is the sum of all rounds of venture investing. The parameter ψ is interpreted as the ratio where all possible vesting occurs. We take $\psi = 8$. We boost the entrepreneur and non-venture investor share by $0.068z$ and impute a share to non-founder employees of $0.02z$.

4 Imputing Missing Data

4.1 Venture ownership shares

In our data, the amounts invested by venture are reported quite fully. As Table 1 shows, however, the valuation of the company at the time of the investment is often not reported. We impute the missing data on the form of the ownership share acquired through the new investment. The share implies pre- and post-money values of the company at the time of the investment.

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,

$$s_{i,i} = \frac{f_i}{f_i + v_i} \quad (6)$$

because the recursion in equation (2) can be written:

$$s_{i,j} = s_{i,j-1}(1 - s_{i,i}). \quad (7)$$

We use the nonlinear logit regression specification

$$s_{i,i} = \frac{1}{1 - \exp(-X_i\delta)}, \quad (8)$$

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
- Amount raised in this round
- Cumulative increase in the Wilshire index over the 2 years preceding this round

Our specification has a complete set of interactions by round number, except that we fit the same coefficients for rounds 5 and higher.

Table 3 shows the results of this regression.

4.2 Imputing values for lingering companies

For the more recent years, our strategy encounters the problem that some companies linger for substantial periods without an exit event—they do not go public, are not acquired, and do not fail, as far as our data collection efforts can determine. In many cases, the companies are continuing in operation despite the lack of further venture funding, while in others, the companies are effectively defunct but their shutdowns have not been recorded. Because the lingering companies have low average payoffs to entrepreneurs and investors, we cannot

<i>Round</i>	<i>Constant</i>	<i>Amount raised</i>	<i>Stock market</i>
1	0.141 (0.306)	0.004 (0.004)	-0.450 (0.228)
2	-0.094 (0.247)	0.003 (0.003)	-0.656 (0.201)
3	0.017 (0.232)	0.003 (0.002)	-0.948 (0.210)
4	0.046 (0.297)	0.010 (0.004)	-1.192 (0.250)
5 or higher	0.057 (0.219)	0.014 (0.003)	-1.276 (0.198)

Standard error of the regression: 0.160

Table 3. Coefficients for Logit Specification for Venture Share of Ownership in a Funding Round

simply omit them from our analysis without creating an upward selection bias. To measure excess returns for the more recent years, we impute values to the lingering companies.

Our approach to this imputation relies on our essentially complete data for earlier years about the outcomes of venture-financed companies. We use these data to measure the expected ratios of discounted exit value to funds invested conditional on the amount of time since the company’s last funding event, which we call the lingering duration. Earlier research has made it clear that returns decline with the duration of lingering—few companies that have neither funded nor exited for five years or more yield substantial returns to investors.

We study the expected gross return ratio, $\frac{x_k}{f_k}$, as a function of the lingering duration. We are interested in learning the expectation of the ratio conditional on a lingering duration of at least τ months. We use the expected ratio to assign an imputed return to the lingerers in the sample. For each value of τ , we find the sum of investments and of the gross realized value for all instances in our sample where a round of investment lingered for τ or more months. Our estimate of the conditional expectation of the gross value ratio is the ratio of the sum of gross realized value to the sum of investments.

Two examples may help clarify this calculation. First, consider a round of investment of \$1 million that yields investors a gross value of \$3 million. Eleven months separate that

round from the next round. In each of those months, we include both figures in the ratio we calculate for lingering durations up to 11 months. The next round is separated from its predecessor by 8 months. We include both figures again in the ratios for lingering durations up to 8 months. Finally the company does an IPO 15 months after the last funding. We include both figures again in the ratios for durations up to 15 months. The company makes no contribution to our calculations for lingering durations over 15 months. Second, consider a round of investment of \$1 million that yields the investors nothing after the firm discontinues operations after 127 months. We include the investment in the denominator of the ratio of all calculations up to 127 months.

Figure 1 shows the results of these calculations. As in certain other of our results, Google’s unusual history and success has a marked effect on the calculations. Google lingered for 47 months between its last round of funding in 1999 and its IPO in 2004. The company accounted for 61 percent of all the value generated in our entire sample from companies that lingered for 47 or more months. Google is not a large fraction of our entire sample, but it is unique in delivering a huge value ratio after lingering so long. We treat Google as a once-in-a-lifetime anomaly and omit it from our imputation procedure.




Figure 1. Expected Value Ratios Conditional on Lingering Duration

The line in Figure 1 without Google behaves reasonably smoothly and in the expected way. It starts at the unconditional mean of 1.78, because all the rounds are included for zero lingering duration. Thus we impute a round value for a firm that has just done a financing round at the end of our sample by multiplying the amount of the new investment by the unconditional mean. The ratio declines to a minimum of around one at two years. For example, we impute a value of 102 percent of investment if the firm has lingered for 24 months without more funding or an exit event. The expected value ratio rises after two years—firms that avoid shutdowns for that long without receiving additional funding tend to give their investors modest positive return ratios, though these correspond to low annual rates of return.

Most venture-backed companies do not linger long—they either deliver high value, cease operations, or get more funding. Half of all value ultimately returned by companies occurs among companies that linger 9 months or less and 90 percent among those that linger 23 months or less. The same proposition holds when measured by invested dollars—half of the investments are accounted for by firms lingering 10 months or less and 90 percent by those lingering 30 months or less.

5 Incentives for Entrepreneurs to Start Companies

Venture-backed companies typically have a single scientist or similar expert, or a small group, who supplies the original concept, contributes a small amount of capital, and finds investors to supply the bulk of the capital. These entrepreneurs, together with any angels, own all of the shares in the company prior to the first round of venture funding.

We do not apply a financial valuation model to the entrepreneurs' interest, for two basic reasons. First, the entrepreneurs 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 entrepreneurs contribute much more than money to the company—their investment is primarily human capital. They 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.

Entrepreneurs 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 entrepreneurs receive when a company exits the venture phase. If the exit is by way of an IPO, the entrepreneurs 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 entrepreneurs 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 2 shows the average exit value received by entrepreneurs, in millions of 2006 dollars, by the date of the exit. Prior to 1999, the figure was quite stable at \$10 million. It jumped to the \$50 million level in 1999 and 2000, then plunged to around \$3 million until 2003. Google's exit in August 2004 resulted in a spike. Most recently, the entrepreneurs' take has risen back to its normal level of around \$9 million. The average over the period was \$9.2 million in 2006 dollars. This figure is below the average over the quarters shown in the figure, because the estimated values of the 9,571 lingerers without reported exits are all included in the last quarter—they deliver only about \$2 million per exit and thus lower the average substantially.

Table 4 shows the distribution of entrepreneurs' exit value in millions of 2006 dollars. About 68 percent of companies yield no meaningful exit value to entrepreneurs. A large fraction of total value to entrepreneurs arises from the tiny fraction of startups that deliver hundreds of millions of dollars of exit value to the entrepreneurs. The table understates the dispersion of actual exit values because it includes the forecasted exit values of the non-exited companies at the end of the period. The contract between venture capital and entrepreneurs does essentially nothing to alleviate their financial extreme specialization in their own companies. Given the nature of the gamble revealed in Table 4, entrepreneurs would benefit by selling some of the value that they would receive in the best outcome in the bottom line, when they would be seriously rich, in exchange for more wealth in the more likely outcomes in the top lines. It would be hard to find a more serious violation of the Borch-Arrow optimality condition—equality of marginal utility in all states of the world—than in the case of venture entrepreneurs.

A diversified investor would be happy to trade this off at a reasonable price, given that




Figure 2. Average Exit Value to Entrepreneurs per Venture-Funded Company, Millions of 2006 Dollars, by Quarter of Exit

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 entrepreneurs pay themselves generous salaries. They use the exit value as an incentive for the entrepreneurs to perform their jobs. Moral hazard and adverse selection bar the provision of any type of insurance to entrepreneurs—they must bear the huge risk shown in Table 4.

The specialization forced upon the entrepreneurs by information limitations substantially lowers the incentive to be an entrepreneur. We can solve the equation,

$$u(\bar{W} + \tilde{W}) = \frac{1}{N} \sum_k u(\bar{W} + W_k), \quad (9)$$

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

$$u(W) = \frac{W^{1-\gamma}}{1-\gamma}. \quad (10)$$

Table 5 shows the non-stochastic equivalent wealth increment for a representative range of

<i>Entrepreneurs' exit value, millions of 2006 dollars</i>	<i>Percent of companies</i>	<i>Percent of total entrepreneurs' exit value</i>
0 to 1	67.4	0.7
1 to 10	19.7	7.7
10 to 50	8.6	22.3
50 to 100	2.4	18.4
100 to 200	1.3	19.0
200 to 500	0.5	15.9
500 to 1000	0.1	9.1
1000 and higher	0.03	6.9

Table 4. Distribution of Exit Value to Entrepreneurs per Venture-Funded Company, Millions of 2006 Dollars

values of the coefficient of relative risk aversion, γ . We take the pre-existing level of wealth to be \$3 million, a figure we take as illustrative of entrepreneurs without substantial wealth. With no risk aversion, the equivalent is the mean of the distribution of entrepreneurs' exit value, \$8.9 million. With a CRRA of 1, at the low end of the reasonable range, the equivalent is only \$1.9 million. At the reasonable CRRA of 2, it is just \$900,000 and at 3, \$600,000.

<i>Coefficient of relative risk aversion</i>	<i>Non-stochastic equivalent of venture payoff, millions of dollars</i>
0	9.2
1	2.0
2	0.9
3	0.6

Table 5. Non-Stochastic Wealth Increment Equivalent to Entrepreneur's Venture Opportunity, for a Range of Coefficients of Relative Risk Aversion

5.1 Implications

The median time from first venture funding to exit in our data is 49 months. We do not have data on the typical number of entrepreneur-founding shareholders, but we believe that three is probably representative. The return based on the \$900,000 figure is a little over \$6,000 per entrepreneur-month. The extreme idiosyncratic risk of venture-backed entrepreneurship and the inability of the venture contract to insure entrepreneurs against the risk result in a tiny incentive facing a prospective entrepreneur.

One interpretation is rational expectations—venture entrepreneurship attracts a flow of participants with low value of time and low risk aversion. Despite the low payoff to entrepreneurs, the flow is large enough to keep the overall payoff to venture startups at the observed level. The selection by value of time and risk aversion may preclude participation by many prospective entrepreneurs with good ideas but higher value of time and more risk aversion.

Another interpretation is biased expectations or unfounded optimism—entrepreneurs overstate the probability of success of their own projects.

Although our data refer only to venture-backed entrepreneurship, the same information limitations inhibit other ways that startups receive financial backing. Investment from a corporate partner is a common alternative to venture capital. And a great deal of development of new technologies and products occurs within established companies. In both of these situations, the scientists, engineers, and others who would be the entrepreneurs if they used venture backing face even less powerful personal incentives—they are basically paid by the hour, not by the value of what they create.

6 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 entrepreneurs, 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, adjusted as discussed earlier for the fraction of companies that generate gains within funds that suffer losses. Figure 3 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 3. Average Exit Value to General Partners, per Venture-Funded Company, Millions of 2006 Dollars, by Quarter of Exit

GPs' earnings per company rose slowly to about \$10 million in the mid-1990s, rose sharply in the boom in 1999, fell to about \$3 billion in 2001 through 2003, spiked with Google, and then rose back to the \$5 million level. The average over the period was \$5.5 million. This figure includes what GPs would have earned if the valuation we assigned to non-exited companies at the end of 2006 was actually the exit value. As in the earlier case of the cash taken out by entrepreneurs, the dollar-weighted average is substantially lower than the time average, because the lingering companies tend to have low exit values. The earnings of GPs are 26 percent of funds invested, so they reduce the limited partners' return by 26 percentage points. Of this, 3 percentage points are from the fee and the remaining 23 percentage points from the carry.

Because venture funds hold dozens of companies, the earnings of the GPs have far less idiosyncratic risk than the earnings of the entrepreneurs. To illustrate this point, we repeat the calculations we did earlier for entrepreneurs, substituting a hypothetical GP who receives all of the earnings from a single fund. Again we calculate the non-stochastic increment to wealth that would replace the distribution of outcomes and assume as before that the GP's existing wealth is \$3 million. We use the data from our simulation of venture funds. The average GP earnings per fund is \$22 million, which is also the non-stochastic equivalent for a risk-neutral GP. With log preferences ($CRRA = 1$), the equivalent falls only to \$16 million, to \$12 million at a $CRRA$ of 2 and \$9 million at a $CRRA$ of 3. These figures overstate the burden of the idiosyncratic risk because most GPs are far wealthier than assumed here and also receive income from more than one fund.

Syndication of venture-fund investments is an important mechanism for diversifying the risks that GPs would otherwise face. Because the funds for any given venture-backed company come from a number of unaffiliated venture funds, each GP has earnings claims on small fractions of large numbers of companies. Diversification of GP earnings may be the explanation for the practice of syndication.

7 Incentives for Outside Investors

The Capital Asset Pricing Model provides the natural framework for studying the incentives facing outside investors. We assume, realistically, that these investors hold fully diversified portfolios and are immune to the idiosyncratic risk of the venture companies they hold. They will provide capital to venture-backed companies if the return compensates the systematic risk of venture investments. The CAPM identifies the general stock market—measured by the Wilshire 5000 index in our work—as the benchmark alternative investment and provides a way to measure the systematic risk of venture.

As we will show, our data are not the only way to measure the risk-adjusted returns received by limited partners in venture funds. Instead of inferring the earnings from the performance of the companies themselves by tracking the financial flow from exit events to the limited partners, one can measure the returns directly from the net receipts of the limited partners from the funds. We show that our calculations match the alternative approach closely and thereby validate our calculations.

Figure 4 shows the basic data we use in studying the risk and return of venture. It shows percent returns over holding periods for venture and the broad stock market by calendar quarter. The venture return is the ratio of cash ultimately received by limited partners from their investments in a quarter to the amount of cash they invested in that quarter. Thus the returns are forward looking. The average holding period weighted by dollars invested is 29 months. 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 4. Holding-Period Returns for Venture and for the Broad Stock Market, by Investment Quarter

In most years, venture returned at least as much as 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 around 1990 and in the years 2000 through 2002 returned less than the stock market.

To compare averages, we use the investments made from 1987 through 2003 and omit those from later years, where the returns for an important fraction of the investments are

imputed in our data because the companies have not yet exited. Weighted by dollars invested in venture, the holding period return for venture averaged 46 percent and the general stock market 19 percent. Venture earned a premium of 27 percentage points. The time average of the quarterly returns shown in Figure 4 for venture was 106 percent, compared to 69 percent for the general stock market, for a premium of 37 percentage points.

The venture premium is the sum of the reward to venture investors for bearing more risk than they would in the stock market and the pure excess return delivered by venture capitalists, alpha. Some commentators also attribute part of the venture premium to illiquidity, but we believe this effect is tiny because venture investors—pension funds, endowments, and wealthy individuals—have low random liquidity needs. They are naturally buy-and-hold investors and thus do not discount investments that require buying and holding.

Figure 4 suggests 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.

In the analysis of returns to securities traded in thick markets, the measurement of alpha and beta is straightforward and well understood. Alpha is the constant and beta the 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, market valuations of venture-backed companies only occur episodically, as additional rounds of investment occur. The time between valuation events is endogenous and random. Second, important information is missing for 74 percent of venture fundings. Companies disseminate information with more enthusiasm and frequency when the values are high than when they are low. Earlier we described how our unique second-look data source provides a direct solution to this sample selection problem.

7.1 Earlier measures of venture’s financial characteristics

The most prominent study of the return and risk of venture investments is Cochrane (2005). Cochrane 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, our framework does not take logs, so 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 various claimants on venture companies. Financial valuation is only appropriate for the limited partners, whose role is to provide financial capital. The entrepreneurs 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, preferences, and dilution protection that the limited partners enjoy in the distribution of cash.

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. 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.

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. It does not, however, appear to yield a reliable measure of the excess return to venture.

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 is within the confidence interval of their estimate.

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 risk-adjusted return to venture investments.

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.

7.2 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,

$$r_{v,t} = \alpha + \beta r_t + \epsilon_t. \quad (11)$$

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.

The problem with 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. Endogeneity induces a negative correlation between ϵ_t and r_t and so biases the estimate of β downward and α upward. The source of the bias is the following: Investments with longer holding periods tend to have lower per-period returns, so they have negative values of the idiosyncratic component, ϵ . Longer holding periods raise the cumulative return r_t . Thus the two variables are negatively correlated.

In fact, CAPM regressions on holding-period returns give absurd estimates. We earlier noted that Cochrane obtained a gigantic α in this framework. We obtained results of the same character, though not as extreme, which we do not report.

7.3 Marking to market

The problem of endogenous holding periods would be avoided by marking investments to market at regular intervals and applying the conventional CAPM over these fixed, non-endogenous intervals. We lack accurate methods for marking venture investments to market between valuation events, but we can show that approximations to intermediate market values essentially solve the CAPM problem. We take the value of the venture investors' holding to change along a straight line between valuation events.

Given approximate market values in each month, we calculate a monthly return index using the same principles as used for market indexes such as the S&P 500—in each month, we calculate the proportional rise in value of an investment at the beginning of the month spread across all available companies in proportion to their values.

The resulting index inherits leads and lags from the interpolation process. Following Dimson (1979), Geltner (1989), Asness, Krail and Liew (2001), Jones and Rhodes-Kropf (2004), and Getmansky, Lo and Makarov (2004), we account for the leads and lags by incorporating leading and lagging values of the market return in the CAPM regression:

$$r_{v,t} = \alpha + \sum_{\tau=-L}^L \beta_{\tau} r_{t+\tau} \epsilon_t. \quad (12)$$

A second consequence of interpolation is to make the disturbance in the CAPM regression—the idiosyncratic component of the venture index—serially correlated. We model the disturbance as a second-order AR process.

We interpret α in the same way as in the standard CAPM and take $\beta = \sum_{\tau=-L}^L \beta_{\tau}$ as a measure of risk akin to the standard beta. To validate this interpretation, we carried out the following experiment: From data on the Nasdaq index for the period from 1971 through the end of 2006, we constructed a large number (about 27,000) of random, endogenous holding period returns that roughly mimicked the timing of the returns in the venture data, from one funding round to the next and from the last funding round to exit. We then built an index for Nasdaq inferred from the holding period returns in the same way as for our index of venture, with linear interpolation for values between the beginning and end of the holding period. Details of the simulation appear in Appendix 2.

Table 6 shows the results of estimating the CAPM parameters for the Nasdaq index directly from the original data and from the index that simulates the construction of our venture index. The estimates of α are similar within sampling variation and the estimates of β are exactly the same. The main consequence of forming the holding-period returns and then extracting the coefficients from the index is a considerable loss of precision—the standard errors of α is twice as high in the venture simulation as in the direct estimate, and the standard error of β is 7 times as high. The results are still informative, however. Notice that even the standard estimator for α has quite a wide confidence interval even when estimated over the period since 1971. One should not have high expectations about the precision of an estimate of α for venture estimated over a shorter period.

<i>Measure</i>	<i>Description</i>	<i>Standard CAPM</i>	<i>Venture simulation</i>
α	Pure excess return, percent per month	-0.29 (0.14)	-0.44 (0.29)
β	Risk, beta	1.26 (0.03)	1.26 (0.21)

Table 6. Simulation Test of Interpolation Procedure

7.4 Index of returns to venture investments

Let s be the share of all venture investors resulting from the investment made at the beginning of an interval together with the investments made in earlier rounds. Venture starts the interval with a position worth $s(v + f)$. This position is worth $s'v'$ at the end of the interval, where s' is the share of the same investments at the end of the interval and v' is the pre-money value at the next valuation event (another round of venture investment or an exit). The share s' is normally the same as s , but may increase because of anti-dilution provisions or decrease because of the exercise of options by entrepreneurs or employees.

To calculate our venture index, we approximate the value of the venture position during an interval by linear interpolation between the beginning and ending value. The denominator of the proportional change in the index is the sum of the values of all venture positions in one month and the numerator is the sum of the values of those positions in the next month. Thus the index measures the value-weighted change in all venture holdings over the month. Our calculation is analogous to the construction of standard stock-market indexes, except that we use approximate values based on linear interpolation instead of reported transaction values when those are absent.

7.5 Results

The method developed earlier in the paper answers the question of whether the extra return that venture enjoys adequately compensates investors for the extra risk. The answer is yes, for the years taken together. Though the risk, measured by beta, is substantial, alpha is slightly positive, though easily within sampling variation.

Table 7 gives the results of estimation of the CAPM regression extended for time aggre-

<i>Measure</i>	<i>Description</i>	<i>Value</i>
α	Pure excess return to venture, percent per month	0.17 (0.59)
β	Venture's beta (sum of lead and lag coefficients)	1.65 (0.35)
σ	Standard error of the regression, percent per month	0.66
	First autoregressive process parameter	0.84 (0.07)
	Second autoregressive process parameter	0.08 (0.07)

Table 7. Estimation Results for Extended CAPM Regression

gation and Figure 5 shows the lead and lag coefficients with bars showing plus and minus one standard error. The point estimate of alpha is about a sixth of a percent per month or 2 percent per year, but the 95-percent confidence interval runs from -12 percent per year to 16 percent per year. Beta is reasonably precisely estimated at 1.65.

Figure 6 shows the value of the venture return index and the fitted values from the CAPM regression. The difference is the idiosyncratic return to venture in general. By far the largest idiosyncratic movement occurred during the tech boom in 1998 through early 2000. Another significant positive venture-specific return occurred in 2004, though this finding will be refined as additional data become available about the companies whose values are imputed as of this writing.

7.6 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.

Figure 7 shows the three available indexes of returns to venture. The three show similar movements over the period for which they overlap. They have similar average quarterly




Figure 5. Lead and Lag Coefficients in the Extended CAPM Regression

returns over the common period from the first quarter of 1989 through the first quarter of 2005: 5.2 percent for our index, 4.7 for the Venture Economics index, and 5.2 percent for the Cambridge Associates index. As noted earlier, the reasonably close agreement between our measure based on company-level flows—including extensive adjustments for preferences, dilution, and GPs’ charges—and the amounts received by limited partners is a confirmation of the accuracy of our adjustments.

Table 8 shows quarterly CAPM regressions for the Venture Economics and Cambridge Associates indexes. They agree broadly with the earlier findings. Neither shows statistically unambiguous evidence of excess returns after risk adjustment. The estimates of β are somewhat higher than our earlier figure of 1.7, but within sampling variation.

8 Concluding Remarks

The venture capital institutions of the United States convert ideas into functioning businesses. We show that the process contains an important bottleneck—for good reasons based




Figure 6. Index of Venture Returns and Fitted Value from CAPM Regression

mainly on moral hazard, the venture contract cannot insure entrepreneurs against the huge idiosyncratic risk of a startup. Risk-adjusted payoffs to the entrepreneurs of startups are remarkably small. Although our results are based entirely on the venture process, we believe that no other arrangement is much better at solving the problem of getting smart people to commercialize their good ideas.

On the other hand, no similar bottlenecks appear to inhibit the efforts of two other groups involved in the venture collaboration. General partners are sufficiently diversified across companies that their earnings are not nearly as exposed to idiosyncratic risk as entrepreneurs. Investors receive risk-adjusted returns comparable to those available from the stock market and other investments. Investors supply capital elastically to venture-backed companies at market rates.




Figure 7. Returns to Venture According to Three Indexes

<i>Index</i>	<i>alpha</i> (percent)	<i>beta</i>	<i>sigma</i> (percent)
Venture Economics	0.9 (2.6)	2.12 (0.69)	9.9
Cambridge Associates	0.0 (3.3)	2.79 (0.78)	11.7

Note: alpha and sigma at quarterly rates

Table 8. Estimates of Alpha and Beta for the Venture Economics and Cambridge Associates Indexes of Venture Returns

References

- Admati, Anat R. and Paul Pfleiderer, “Robust Financial Contracting and the Role of Venture Capitalists,” *Journal of Finance*, 1994, 49 (2), 371.
- Asness, Clifford, Robert Krail, and John Liew, “Do Hedge Funds Hedge?,” *Journal of Portfolio Management*, 2001, Fall, 6–19.
- Bagley, Constance E. and Craig E. Dauchy, *The Entrepreneur’s Guide to Business Law*, West Legal Studies in Business, 2003.
- Casamatta, Catherine, “Financing and Advising: Optimal Financial Contracts with Venture Capitalists,” *Journal of Finance*, October 2003, 58 (5), 2059.
- Cochrane, John H., “The Risk and Return of Venture Capital,” *Journal of Financial Economics*, 2005, 75, 3–52.
- Dimson, Elroy, “Risk Measurement when Shares Are Subject to Infrequent Trading,” *Journal of Financial Economics*, 1979, 7, 197–226.
- Geltner, David, “Estimating Real Estate’s Systematic Risk from Aggregate Level Appraisal-Based Returns,” *American Real Estate and Urban Economics Association Journal*, 1989, 17 (4), 463–481.
- Getmansky, Mila, Andrew W. Lo, and Igor Makarov, “An Econometric Model of Serial Correlation and Illiquidity in Hedge Fund Returns,” *Journal of Financial Economics*, 2004, 74, 529–609.
- 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.
- Repullo, Rafael and Javier Suarez, “Venture Capital Finance: A Security Design Approach,” *Review of Finance*, 2004, 8, 75–108.
- Schmidt, Klaus M., “Convertible Securities and Venture Capital Finance,” *Journal of Finance*, June 2003, 58 (3), 1139.
- Wilmerding, Alex, *Deal Terms: The Finer Points of Venture Capital Deal Structures, Valuations, Term Sheets, Stock Options, and Getting Deals Done*, Aspatore, 2003.

Appendix 1: Adjustment of Ownership Shares in Down Rounds

Most agreements between venture investors and entrepreneurs call for the issuance of additional shares to investors in earlier rounds when the share price in a new round falls short of the price in the previous round—what is called a “down round.” The *weighted-average anti-dilution* provision is the most common form of adjustment. To explain this provision, we let j be the number of the current round and let i range over the earlier rounds. We let $n_{i,j}$ be the number of shares effectively held by round i investors as of round j . “Effectively” means the number of common shares that would result from the investors exercising their conversion rights. Anti-dilution provisions take effect by lowering the conversion price, $p_{i,j}$, and increasing the number of shares, $n_{i,j} = f_i/p_{i,j}$, the investors receive upon conversion. We let $N_{i,j}$ be the total number of shares outstanding at the conclusion of round j .

To identify the investors eligible for the anti-dilution adjustment without knowing actual share counts and share prices, we proceed as follows: The conversion price per ownership share point for earlier investors as of the last round is $f_i/s_{i,j-1}$. The price per ownership point (measured as of the previous round) paid by the new round is the pre-money value v_j . Thus the eligible rounds are those with $f_i/s_{i,j-1} > v_j$. We let A_j designate the set of these rounds and $\sim A_j$ the set of rounds not subject to adjustment, including the common shares.

The weighted-average provision specifies adjustment factors for the eligible earlier investors in the case of a down round:

$$a_{i,j} = \frac{N_{j-1} + n_{j,j}}{N_{j-1} + \frac{p_{j,j}}{p_{i,j-1}} n_{j,j}}. \quad (13)$$

The numerator is the number of shares after round j if the existing shareholders did not receive any new shares. The denominator is the number of shares if the new round had to pay the higher price paid by an investor in round i . The new conversion price is the old price divided by the adjustment factor.

The quantity $p_{j,j}n_{j,j}$ is f_j , the amount invested in the new round. The earlier conversion price $p_{i,j-1}$ is $f_i/n_{i,j-1}$. Thus

$$a_{i,j} = \frac{N_{j-1} + n_{j,j}}{N_{j-1} + \frac{f_j}{f_i} n_{i,j-1}}. \quad (14)$$

To reduce the complexity of what follows, we write

$$b_{i,j} = \frac{1}{N_{j-1} + \frac{f_j}{f_i} n_{i,j-1}}, \quad (15)$$

so the adjustment factor is

$$a_{i,j} = b_{i,j}(N_{j-1} + n_{j,j}). \quad (16)$$

The total number of shares of the earlier investors, after adjustment for those who paid more than the current price, is

$$\begin{aligned} \tilde{N}_j &= \sum_{i \in A_j} b_{i,j}(N_{j-1} + n_{j,j})n_{i,j-1} + \sum_{i \in \sim A_j} n_{i,j-1} \\ &= B_j(N_{j-1} + n_{j,j}) + \bar{N}_j. \end{aligned} \quad (17)$$

The ownership share of the new round is

$$\begin{aligned} s_{j,j} &= \frac{f_j}{v_j + f_j} \\ &= \frac{n_{j,j}}{\tilde{N}_j} \\ &= \frac{n_{j,j}}{B_j(N_{j-1} + n_{j,j}) + \bar{N}_j + n_{j,j}}. \end{aligned} \quad (18)$$

This can be written as a linear equation in the unknown $n_{j,j}$. We solve for $n_{j,j}$, multiply by the expression in equation (14), and use the resulting share counts to form the new values of the ownership shares $s_{i,j}$. By enlarging the ownership shares for the investors who paid more than the current share price, the provision reduces the shares of the entrepreneurs and other earlier investors even more than the normal dilution from a new round.

The calculations described above are homogeneous in the numbers of shares, so we can normalize the total number of shares from the previous round at one. After this normalization, $n_{i,j-1} = s_{i,j-1}$. The effect of the calculations, including forming the new shares $s_{i,j}$, is to modify equation (2) to include the rearrangement of equity interests among the existing shareholders that occurs in a down round.

The description of the weighted-average updating in a down round given in this appendix is rather more complicated than in standard references on venture contract terms, such as zzzzz. Those descriptions assume the availability of data on share holdings and conversion prices. Our approach is tailored to our data, which require us to infer these numbers from data on pre- and post-money value.

In the case of the full-ratchet anti-dilution adjustment, the rearrangement of ownership shares can be expressed in the same framework. Those investors who paid more than the current price for their shares in an earlier round receive a proportional increase in ownership (decrease in conversion price) equal to the ratio of the earlier price to the current price. If an earlier round, i , had a higher price, its number of shares becomes $f_i/p_{j,j}$. As before, the current price is $p_{j,j} = f_j/n_{j,j}$. Thus the number of adjusted shares brought into the current round is

$$\sum_{i \in A_j} \frac{f_i}{f_j} n_{j,j} + \sum_{i \in \sim A_j} n_{i,j-1}. \quad (19)$$

Again, we can solve the equation for the ownership share of the new round,

$$s_{j,j} = \frac{n_{j,j}}{n_{j,j} \sum_{i \in A_j} \frac{f_i}{f_j} + \sum_{i \in \sim A_j} n_{i,j-1} + n_{j,j}}, \quad (20)$$

for the new ownership $n_{j,j}$, and then calculate the ownership shares of the earlier investors and entrepreneurs. As before, the total number of shares owned as of the previous round can be normalized at one, so the procedure developed here is a recursion that describes the rearrangement in ownership shares that occurs in a down round because of the anti-dilution provision benefiting earlier venture investors.

Appendix 2: Validation of Estimation Method

For the period from 1971 through 2006, we generated synthetic holdings of the Nasdaq index over periods similar to the venture holding periods. Each holding reached an exit according to the hazard function,

$$h = \frac{1}{1 + \exp(\kappa - \theta R - \lambda \tau)}, \quad (21)$$

where h is the probability of an exit conditional on none having occurred to date, R is the cumulative return ratio for the investment since its beginning, τ is the number of months since the investment was made, and κ , θ , and λ are parameters. A positive value of θ generates the endogeneity of returns and exits characteristic of venture holdings. We chose $\kappa = 5$, $\theta = 1$, and $\lambda = 0.1$. We generated about 27,000 of these investments and then formed

an index of Nasdaq returns based on the investments using the same procedure we used to form the index of venture returns. Finally, we performed a standard CAPM regression of the Nasdaq on the Wilshire and an extended CAPM regression of the index return on the Wilshire, as shown in Table 6.