

NBER WORKING PAPER SERIES

BENCHMARKING THE RETURNS TO VENTURE

Susan E. Woodward
Robert E. Hall

Working Paper 10202
<http://www.nber.org/papers/w10202>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2003

The views expressed herein are those of the authors and not necessarily those of the National Bureau of Economic Research.

©2003 by Susan E. Woodward and Robert E. Hall. 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.

Benchmarking the Returns to Venture
Susan E. Woodward and Robert E. Hall
NBER Working Paper No. 10202
December 2003
JEL No. G24, G12

ABSTRACT

We describe a new index of the current and historical returns to venture-type capital. The conceptual basis for the index is the value of a continuously reinvested value-weighted portfolio of all venture-backed and similar pre-public companies. It provides a metric for private equity comparable to the S&P 500 for public equity. We build the index from valuations revealed in episodic transactions in the companies' shares – private placements of new rounds of equity funding, IPOs, acquisitions, and liquidations. Our approach to dealing with the episodic nature of the data is similar to the one used in constructing indexes of real-estate value from transaction data for individual properties. We have extended earlier sources of data to deal with selection bias – we tracked down unfavorable valuations that were less likely to be reported in the earlier data. We also use econometric techniques to handle the remaining selection bias. The resulting index has important uses in marking venture portfolios to market and in assessing the performance of venture investments.

Susan E. Woodward
Sand Hill Econometrics, Inc
swoodward@sandhillecon.com

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

I. Introduction

In a brief 10 years, venture investing has grown from a tantalizing sideshow to a serious and institutionalized function of US capital markets. Flows into venture capital funds rose from \$3 billion in 1990 to \$103 billion in 2000. Some of these interests are held directly as stock in the issuing company, and some as limited partnerships in venture capital funds. Holders are primarily pension plans, endowment funds, corporations, and individual investors. Total venture-type private equity now totals roughly \$500 billion in the US. Venture capital is now a major force in launching companies in the public markets. Over the last five years, roughly two-thirds of the companies doing initial public offerings had backing from venture capital funds. Investment by outside investors in pre-public companies is no longer experimental or exploratory, but a permanent feature of the US financial landscape. Venture is here to stay.

Although large sums are now invested in venture, reliable information on the outcomes from these investments is hard to come by. For registered and publicly-traded securities, quantitative performance evaluation is straightforward and nearly universal. Standard indices, such as the S&P500 and the Wilshire 5000, provide measures of overall market measures of value, return, risk, and risk-adjusted return. For non-traded holdings, such as venture, investors have no similar benchmark.

This paper introduces a new index for venture, providing venture a benchmark on the same footing as the S&P500 for traded equity. The index can be used for venture capital the same way market-wide indices are used for traded securities. Because it enables calculation of the mean and variance of venture returns, as well as their covariance with returns on other assets, the index can be used for asset allocation analysis. A second use is portfolio performance evaluation. By comparing portfolio returns to index returns, investors can get a better measure of the performance of particular holdings in comparison

to all venture capital. The third use is unique to venture: The index can be used to update company and portfolio values to obtain more timely estimates of value, and even to predict the returns that venture general partners will be reporting six to nine months forward.

With a reliable benchmark for venture capital and a set of techniques for using it, investors will gain a clearer understanding of their successes and failures. Decades ago, mutual fund performance analysis resembled the analysis now done for venture funds. Funds were ranked by return over a given interval, and little attention was paid to risk because there was no clear standard for measuring it. Things changed in the late 1960s when the capital asset pricing model was introduced, and the focus on performance shifted from pure return to risk and risk-adjusted return. The introduction of better quantitative tools for analyzing venture capital will bring a similar revolution in the assessment of venture performance.

In this paper, we use the term *venture capital* to mean equity in pre-public companies that have outside investors. We exclude leveraged buyouts, management buyouts, and private placements in public companies (PIPES) in our notion of venture-type investments. And we exclude businesses that have no intention or likelihood of becoming public companies.

II. The Challenges of Building the Index

The index is built from company-level pricing data, not fund-level return data. We believe that only the company approach allows creation of a timely, unbiased, monthly index of value that is analogous to traded-market indices such as the S&P500. Only company-level transactions can pinpoint value at a point in time. The companies from which the index is built are those that are (a) privately held, (b) organized as C corporations, and (c) have sold equity securities (mainly convertible preferred stock) to outside investors. To get a sense of how our companies fit into the universe of businesses,

consider the following data: Approximately 22 million businesses in the U.S. file income tax returns. Five million have a payroll. Of these, 1.5 million are organized as C corporations. Approximately 12,000 C corporations are registered with the SEC and trade on stock exchanges. There are at present roughly 20,000 companies represented in the index, over the period starting in 1987, with a total of roughly 65,000 valuation events (private rounds of funding, acquisitions, IPOs, and shutdowns). About 9,000 of these companies are still private and operating.

Building an index from company-level data faces two major challenges. First, events that reveal market values do not occur continuously for private companies as they do for traded stocks, but instead occur episodically, when a company raises new money, goes public, is acquired, or goes out of business. Second, reporting of valuation by companies that complete deals is voluntary. Often companies report that they completed a fund-raising but do not report the value at which stock was sold. The companies that do report valuations are not a random sample of all companies, they are a biased sample. Good news is reported more often than bad.

We have developed tools to solve both of these problems. Intermittent pricing is handled by an interpolation method similar to the one used in repeat-sales analysis. The repeat-sales technique has been enthusiastically adopted in other markets facing similar intermittent pricing challenges. For example, it has revolutionized estimation of nationwide and local housing values, replacing tedious and unreliable individual appraisals. Freddie Mac and Fannie Mae now pool their data and build an index for houses based on this technique. The Federal Housing Finance Board also compiles a repeat-sales index for the nation and numerous metropolitan areas, available on its website (www.fhfb.gov). The firm of Case, Shiller, and Weiss provides local price indices and individual house valuations based on this technique. In addition, the US Commerce Department uses a repeat sales model to construct its price indices for durable goods.

The second problem, bias in the reporting of valuations, is addressed with techniques developed by economists over the last two decades for dealing with selection

bias. These techniques exploit all available information about companies that have done even one round of fund raising, whether or not they report value when a round of funding is done.

These two techniques are applied together to produce an unbiased index of value for venture companies. This index serves as the private company analog of a market-wide index such as the S&P500. It measures return more accurately than do other measures of venture returns that fail to account for intermittent pricing and fail to correct for selection bias. Because the index is calculated monthly, it offers the capacity to measure variance through time and covariance with other assets' returns. While it remains an index of estimated value—as opposed to an index composed of almost continuously reported bona fide trades such as the S&P500—it is more powerful than any other current tool for quantitative analysis of venture holdings.

The effort necessary to build the database for the index was substantial. First, when data are obtained from multiple sources, they must be integrated. In addition, duplicates must be eliminated, small and large inconsistencies in dates of deals, amounts of money raised, and post-money values must be resolved, typos identified and dealt with, and stray irrelevant records—foreign companies, buyout deals, PIPEs, venture funds that are not operating companies—must be eliminated.

As we began the work of assembling the data and building the index, we discovered another source of selection bias: we lacked information about many companies that last raised funds some time ago, but are not reported to have been acquired, gone public, or shut down. We undertook research to learn what had become of companies with a last private round prior to 2001, but no exit. This research confirmed that these companies were more likely to have shut down worthless than other companies. Thus, it was necessary to embark on a systematic program of tracking them down. This research is an important contribution to the integrity of the index and has been in progress for nearly two years.

Our work has some points of contact with John Cochrane's [2003] investigation of the risk and return to venture investments. Cochrane imposes much more parametric structure on the problem than we do. In particular, he works within a model where return ratios are log-normal and models sample selection parametrically. Although an index of returns and thus a cumulative index are implicit in Cochrane's work, they are not the main focus. We have found that the actual distribution of returns departs sufficiently from log-normal to result in some distortion of a returns index inferred from the study of log-returns.

Our work is also related to Quigley and Woodward [2003], who embed the problem of measuring venture returns in the model proposed by Bailey, Muth, and Nourse [1963] and subsequently widely used to infer indexes of housing prices from episodic transactions. Because this method requires the use of logs and therefore an implicit assumption of log-normality, it suffers from the same problem we found with Cochrane's approach. Quigley and Woodward use the adjustment for sample selection from Heckman [1979] and they use the same data underlying this paper.

A. Uses of the Index—Marking to Market

An important use of the index is making more timely estimates of venture portfolio values. Investors can estimate the value of an entire portfolio or particular fund by updating the most recent deal values with the index. This can be done using the broad index for a broadly diversified portfolio, or by applying sector indices for specific parts of a portfolio.

As soon as we had built the index, we compared it to the quarterly averages of returns reported by general partners of venture funds. It is clear that these average GP quarterly returns lag quarterly returns on the index by about six months. Quarterly returns on the index *predict* the average quarterly returns reported by GPs six to nine months forward. GPs use their knowledge about individual companies in their portfolios to estimate value, but they do not use market-wide information. Indeed, without an index like the one we have developed, how could they do so in a systematic way? The index provides

information about valuations on a market-wide basis that is both relevant and more current than what is in the hands of limited partners of venture funds. The index could not predict the average GP quarterly returns if this were not so.

B. Uses of the Index—Portfolio Allocation Analysis

A second use of the index is for portfolio asset allocation analysis. Given the means, variances, and covariances of returns among different asset classes, what is the optimal allocation to venture capital? Given an investor's particular portfolio of stocks and bonds (which may be less-than-fully diversified, and likely to remain so for tax reasons) what is the optimal allocation to venture capital? How many companies or funds does an investor need, and how should investments be distributed across industries and stages, in order for an investor to be reasonably diversified in venture capital? For traded securities, the standard tools for addressing these issues are the means, variances, and covariances of returns. Prior to the creation of an index for venture capital, no similar analysis of venture allocations was possible. The index and its sector indices provide the same tools for portfolio analysis including venture capital that have become the industry standard for traded securities.

C. Uses of the Index—Performance Evaluation

We believe that the index is the right starting point for the most informative approach to the evaluation of performance for a venture portfolio. The index enables investors to measure risk-adjusted performance on their venture investments in essentially the same way as they measure it for their traded securities. In other words, investors can now do the same risk-adjusted analysis for their venture holdings that they do for their traded stocks and bonds.

The most basic evaluation compares the changes in the value of the index to changes in the value of the investor's own portfolio. For a portfolio that contains a typical mix of venture investments with respect to industry and stage, this comparison will

indicate whether the investor's own portfolio is performing above or below the average, adjusted for risk.

For a portfolio that is not a representative set of venture investments, we can provide a targeted index. We have developed industry indexes for:

- Information Technology
- Health
- Retail
- Other sectors

We are in the process of developing

- The Venture-backed index, a narrower index including only companies with backing from a venture fund

Company stage indexes for:

- Early stage companies
- Late stage companies

Vintage Year indices for:

- All vintage years 1993 and later

With the sub-indexes, investors can evaluate performance of the parts of the portfolio invested in different industries and stages by comparing them to the relevant index. This analysis will indicate whether the portfolio is performing better or worse than a portfolio invested in companies of a similar industry and stage. Custom industry indexes can be produced upon request.

Investors will likely be interested in the past performance of prospective managers. How has the track record of this group measured up against other managers? Against all venture? Against funds emphasizing the same industry group? Again, comparison to the overall index is the first guide, and then comparison to a sub-index based on sector or stage

will indicate whether the source of over- or under-performance is sector-specific or deal-specific. Venture capital managers often represent themselves as not just good at choosing companies to invest in, but also as adding value to these companies themselves through their involvement in strategic decisions. The index will provide the tools necessary to substantiate this claim.

The venture community—entrepreneurs, investors, and venture capitalists—can gain from the discipline of formal performance evaluation. Some investors are reluctant to invest at all in a sector where no benchmarks or analytics for optimal allocation are available. The index provides guidance on return and risk, as well as correlation with other asset classes. These will give comfort to investors that their investment decisions are consistent with portfolio diversification strategies.

III. Current Practice for Performance Evaluation

A. The Present Focus in Venture Capital

The two most common methods of performance evaluation for venture are the *IRR* or internal rate of return and the *investment multiple*. The IRR is the rate of return that equates the present discounted value of the cash received to the present value of the amount invested. The IRR is sensitive to the timing of investments and returns. The investment multiple is the ratio of the amount received back to the amount paid in. It neglects the timing of the flows.

Investors typically compare the IRR and investment multiple for a fund to data provided by consultants for similar metrics for funds that were started in the same year, usually called the *vintage year* of the fund. Thus a fund started in 1997 will be compared to average results for funds started in 1997. Some investors also make use of market-wide information by comparing their portfolio IRRs to ones calculated from publicly traded stocks.

In our experience, the procedures followed by pension funds and other large fiduciary venture holders are similar. Table 1 shows the first few lines of the most recent disclosure by the largest public entity holding venture investments, the California Public Employees Retirement System (CalPERS).

<i>Fund Description</i>	<i>Vintage Year</i>	<i>Capital Committed</i>	<i>Cash In</i>	<i>Cash Out</i>	<i>Cash Out & Remaining Value</i>	<i>Net IRR</i>	<i>Investment Multiple</i>
1818 Fund II, L.P.	1993	\$75,000,000	\$75,127,168	\$90,273,988	\$119,983,944	11.67%	1.60x
ABS Capital Partners II, L.P.	1996	50,000,000	45,721,976	52,005,512	57,967,781	9.61	1.27x
ABS Capital Partners III, L.P.	1999	75,000,000	61,270,475	6,903,963	35,411,882	-16.96	0.58x
ABS Capital Partners IV, L.P.	2000	75,000,000	30,955,547	6,389,208	28,355,365	-8.21	0.92x

Table 1. Extract from CalPERS's Venture Performance Disclosure (<http://www.calpers.ca.gov/invest/aim/aim.asp>)

Table 1 reveals a central feature of the evaluation process. The difference between the column labeled *Cash Out & Remaining Value* and the one labeled *Cash Out* is the remaining value—the general partner's estimate of the current value of the fund's investment in portfolio companies that have not yet had a terminal event. Even in the case of the 10-year-old 1818 Fund II, there is almost \$30 million of estimated value out of a total of \$120 million. For the funds started in 1999 and later, the estimated value dominates the total value that is reported to the limited partners. The two performance measures both use the estimated value as if it were the equivalent of cash received. Thus the accuracy of the measurement of venture performance rests on the accuracy of the estimates of remaining value. If these estimates can be improved upon, so can the estimates of the internal rate of return.

B. Role of Estimated Remaining Value

As Table 1 demonstrates, for unfinished funds, calculations of IRRs and other performance metrics rely on funds' estimates of remaining value. Even relatively mature funds (10 years or more) often have a substantial fraction of their value in companies that have not had a terminal event and are still operating as private standalone companies.

Each quarter, the general partners for most venture partnerships report estimated values for the remaining companies in their portfolio. Valuation techniques vary: most apply the price from the last valuation event to all of their holdings and extend it to the present; some write down companies based on insider bad news, a few adjust value upwards for companies meeting business milestones. Some funds systematically assign conservative remaining values, in which case their performance is understated until all of the companies have terminal events. We have observed none who uses any sort of market-wide information to adjust values.

The values arrived at by the GPs as described above are precisely the values we were referring to earlier when we noted that GPs do not use market-wide information to update their estimates of remaining values. Comparing reported averages of GP quarterly returns across many funds to returns on the index, we find that on average, the GP valuations are smoothed, too high, and about six to nine months old. For individual funds, as opposed to averages built from biased samples, we find reported returns are about right when we compare them to the returns for their portfolio companies, but they are smoothed and, as with the averages, about six months old. They are too low in a rising market, but too high in a falling market. This evidence confirms our belief that GPs use company-specific knowledge to update company values, but not information about venture returns in general. By incorporating market-wide information, investors can obtain more accurate estimates of portfolio value.

C. Performance Evaluation for Publicly Traded Securities

Benchmarks for traded securities have reached a highly refined stage of development. Benchmarks are derived from state of the art finance theory. A leader in this field is Barra, a company founded by finance economist Barr A. Rosenberg. Because traded equities have known prices at frequent intervals, evaluation focuses on the *excess returns* earned by a portfolio on a daily, weekly, or monthly basis. The excess return is the capital gain plus dividend received during the period divided by the price at the beginning of the period, less the current safe interest rate. For brevity, we will call this measure the return, as the subtraction of the risk-free interest rate has little material effect. A given portfolio may have hundreds or even thousands of returns. Evaluation involves comparing these returns to a benchmark derived from the returns of traded equities in general. In venture, the similar evaluation problem is addressed by computing a measure for a fund, such as the IRR or investment multiple and comparing that measure to an average of similar measures for funds in general.

For traded equities, a more powerful approach now dominates. According to finance theory, the returns to portfolio i in period t , $r_{i,t}$, are related to the returns of the market in general, $r_{M,t}$, by the following equation:

$$r_{i,t} = \alpha_i + \beta_i r_{M,t}$$

This equation encapsulates the Capital Asset Pricing Model or CAPM. The coefficient β_i measures the covariation of the portfolio return with the market return. It shows the percent return expected in a period when the market delivers a one-percent return. β_i measures the *financial risk* of the portfolio. In finance theory, risk relates to covariation. A risky stock is one that moves more than point for point with the market; investors require a higher return. The most desirable stock would insure the holder against movements of the market and would have a negative β_i . The CAPM's predictions for average returns are

$$\bar{r}_i = \alpha_i + \beta_i \bar{r}_M$$

Thus the average return \bar{r}_i is higher for riskier stocks with higher β_i .

The β_i of a portfolio is the value-weighted average of the β s of the stocks in the portfolio. These β s are provided by many financial services firms. Further, financial science has developed more elaborate versions of the CAPM with factors apart from the market return. All these models deliver the capability to calculate the *average risk adjustment*, $\beta_i \bar{r}_M$ or its generalization, over a given span of time.

The CAPM makes the sharp prediction that the intercept, α_i , should be zero. This is obviously the case for a portfolio that matches the market, where $\bar{r}_i = \bar{r}_M$ and thus $\alpha_i = 0$ and $\beta_i = 1$. It is a fundamental conclusion of finance theory that, in a market where all traders share the same information, $\alpha_i = 0$ no matter what is the risk of the portfolio. A single trader with better information or superior ability to analyze information may achieve a positive value of α_i . After calculating the β_i of a portfolio, one can calculate α_i from

$$\alpha_i = \bar{r}_i - \beta_i \bar{r}_M$$

This is the *risk-adjusted average return*. It is the standard metric of performance for portfolios of traded securities. It takes full account of the fact that a manager can earn a higher return by taking on more risk. A positive risk adjusted return signals strong performance, while a negative return signals below-market performance.

If outside measures of β_i are not available, another procedure gives essentially the same estimate of the metric α_i . Treat the equation $r_{i,t} = \alpha_i + \beta_i r_{M,t}$ as a regression and estimate the intercept and slope by standard regression methods. The intercept is the risk-adjusted average return.

The basic lessons of portfolio evaluation derived from modern finance theory are (1) the appropriate measure of performance is the average over a longer time interval of returns measured over relatively brief intervals, and (2) the risk adjustment $\beta_i \bar{r}_M$ or its generalization is the appropriate way to standardize the performance metric for the critical factor of risk.

D. Limitations of Current Methods of Venture Performance Evaluation

Performance evaluation for venture investing has not tracked modern developments in portfolio evaluation of traded securities because no periodic benchmark of venture value has been available. Investors have used investment multiples and IRRs. The information provided by an IRR is useful. But the IRR has some limitations.

First, the comparison of one fund's IRR with an average IRR benchmark cannot account for risk. In a rising market, is an investment over-performing, or just more risky than average? And is an under-performer actually over-performing, when properly adjusted for risk? A true risk-adjusted performance captures this, and an IRR comparison does not.

Second, the IRR benchmarks are stale because they are computed from stale valuations. The IRRs are computed from VC estimates of value, which lag market-wide changes by roughly 6 months. This phenomenon is established by analyzing the time series properties of the index returns in comparison with reported averages of GP quarterly returns. The latter are sometimes call *time weighted IRRs*.

Third, the computation of internal rates of return is inconsistent across portfolios holding still-private companies because there is no generally accepted method for estimating the remaining value of these companies or for computing partial internal rates. Practice varies with respect to this key element.

Fourth, current available IRR benchmarks—vintage year averages—are biased upwards because they miss data from the less successful investments. A company-based approach to benchmark construction better lends itself to bias correction because

individual companies generally seek publicity when they raise money even if they do not report the valuation implied by the terms of the funding, and because additional research can be done to establish the status of those who cease reporting. Companies generate public information while funds do not. This information makes it possible to correct for bias from unreported values.

Fifth, with the exception of the pooled-cash-flow IRRs, current benchmarks are not value-weighted. Tiny funds carry just as much weight in the average as do big funds.

In sum, the traditional venture benchmarks of vintage year IRR and investment multiples are not as informative as the benchmarks for portfolios holding traded securities. They lack the ability to adjust for the risk of the portfolio under evaluation. They suffer in two respects from being based on data from venture funds—they suffer lags in valuation and they fail to include disappointments. And for the most part, current benchmarks are not value-weighted. Investors can improve their assessments of portfolio performance by using the index to evaluate the performance of their venture holdings. This approach and the concepts behind it are not new: only the benchmark is new. Investors already use public market indices to evaluate the performance of their traded securities in the same way they can use the index.

IV. Calculating the Benchmark

This section describes in more detail the techniques for benchmark construction. The our approach to building a benchmark for venture performance incorporates the methods now generally used for evaluating portfolios of public equities, described above. The conceptual basis for our approach, therefore, is to construct a time series of the average return on an unbiased, value-weighted, market-wide portfolio of venture companies.

To build the index, we cumulate the individual company monthly returns into indices of cumulative returns. The index (as well as the sector indices) has the interpretation of the value that an investor would have in a given month as a result of investing \$100 in January 1989 and reinvesting each month in proportion to the values of the companies at the beginning of the month. Thus each index is a value-weighted index of cumulative returns. To put the point in reverse, the monthly returns can be calculated as the rates of growth of the index.

A. Steps in Constructing Average Returns

This section describes the construction of the total index. Construction of sub-indices is the same, with the inclusion of companies fitting the definition of the sub-index.

Step 1 We obtain or estimate a value for every company in every month from its entry into the data to its exit or to the present, in the case of no exit. For exited companies that report value at every round, we use the post-money value of the round at the report date. Between report dates, we interpolate value for the company using an interpolation index, from the post-money value of the prior round to the pre-money value of the next round. For missing data—either missing valuations for known valuation events or for companies of unknown status—we use econometric methods to impute values. One equation fills in missing data for rounds of venture financing where the data sources do not give the valuation of the company at the time of the round. Another fills in terminal values (actually current values – or at least as current as the end of the period covered by the index) for companies that have not had an identified terminal event.

For exited companies with some reported rounds and some (or possibly none, in the case of some unreported-value acquisitions) non-reported rounds, we use reported values at the round date when they are reported, and we estimate value for unreported-value rounds. For non-exited companies, we use either reported value or estimated value and interpolate as above between rounds.

For dates after the last reported round, we estimate value for every month, depending on company characteristics and the time elapsed since the last round of funding.

A “lost company” is one that has not exited but also has not done a new round of private funding for several years. We have done substantial research on these companies and have found with that many have exited, either by shutdown, acquisition, or IPO. We use the results of our research to estimate the likelihood of these outcomes for each company based on whether it was last known to be in development, shipping, profitable, amounts raised, and its other information. The more recent the company’s latest round, the lower is the likelihood that its current status cannot be established.

We interpolate company value between actual (or estimated) valuation event dates to assign every company a value for every month from its entry into the database until its last known round of funding (or for lost companies, its estimated exit date). To do this we use both pre-money valuations (value before new money raised is added) and post-money value (after new money is added).

Step 2 We build a value-weighted index from these values. Our aim, as is the aim of indices for traded securities, is for the index to report the value of a portfolio that 1) invests in every round of every company in proportion to its value, and 2) rebalances when companies exit (takes the resulting cash and re-invests it in the remaining companies) and when their values change.

For each month, we compute the total value of the companies that have actual or interpolated values for that month at the beginning of the month. Then, for each company, we calculate the ratio of its actual or interpolated value in a given month to the total value of all companies in the previous month. This is the company's contribution (weight) to the return ratio for the index level for that month. The return ratio of the index in that month is the sum of the contributions of all the companies. Then, to compute the value of the index for each month, we start with 100 as the value of the index for January 1987. For the next month, we multiply the return ratio for that month by 100. For all subsequent months, we multiply the previous month’s index level by the return ratio for the following month.

Step 3 We use the index value calculated in Step 3 as the interpolation index for a second-round calculation of the index. Total return from one event to another is the same, but the route by which the value gets from one event to another has more variation in the second iteration. This step refines the interpolation process and improves the accuracy of the resulting index. This step is then repeated until the index converges, that is, until adjustments to individual company valuations between rounds reaches a small number.

We rejected the more standard ordinary least squares repeat-sales regression approach because our research indicates that the returns to venture are not log-normally distributed. Taking logs of returns and then de-logging the coefficients to recover an index will only yield a satisfactory result if the return distributions are log-normal. The bias inevitably resulting from taking logs can be offset if the distribution is log normal or a member of some other parametric family but not if the distribution is unknown. The data do not pass any test of log-normality. To the naked eye, not only are the distributions fat-tailed, they are mildly bi-modal. To assure that this phenomenon was not coming from dirt in the data (such as typographical errors) we undertook a special review of the 100 companies who contribute the most value (about 35 percent) to the index. All were confirmed to be real events.

Our approach to dealing with sample selection is informed by the literature emanating from Heckman's [1979] famous work, but does not follow a parametric specification as in Cochrane [2003] or Quigley and Woodward [2003]. Our primary line of attack is to track down data previously missing. This approach dominates all others, especially in situations where the identification of the selection effect turns entirely on the functional form of the Mills ratio or its equivalent. Our secondary line of attack is to impute missing data from equations that we believe are not too badly contaminated by selection bias. In particular, we use the results of tracking down of previously missing data as a guide to imputing the remaining missing data.

V. Behavior of the Index

Figure 1 shows the index from 1989 through June 2003. It is the dollar value of an investment of \$100 made in January 1989 and reinvested monthly in all venture-type companies, in proportion to their values in each month. The index reached a sharp peak at about 2300 in mid-2000, fell to about 1200 in 2002 and has risen somewhat since then.




Figure 1. The Index for All Industries

Figure 2 shows the returns underlying the index, stated at annual percentage rates. The return is the increased value at the end of the month of a \$100 investment in all venture-type equity at the beginning of the month, counting cash received from IPOs and acquisitions and counting the increase or decrease in the values of the companies remaining in the venture-type universe. We multiply the monthly return by 12 to state it at annual rates. Returns remained at moderate levels until 1998, when they rose to a sharp peak in early 2000, then plummeted to negative values in 2001 and 2002, and finally regained positive levels in late 2002. The extra volatility in the last few months of the plot

are intrinsic to the construction method—fewer pairs of valuations are available for the last few months, so the returns are necessarily measured with more noise. As more data become available in coming months, the revised index will not have this noise for early 2003, but the index will always have noise in the most recent few months.




Figure 2. Returns for All Industries

Figure 3 shows indexes for four industry groups. The first is information technology, which accounts for about half the value in the venture-type universe. Not surprisingly, the IT index moves in much the same pattern as the index for all venture-type companies. The index for health shows that the sector did not participate nearly as much in the overshooting of value that occurred in IT—a \$100 investment in the health sector peaked at just over \$2000 in late 2000, as against more than \$3000 in IT, but the current value of that investment is around \$1800, compared to \$1500 for IT. The performance of venture-type investments in retail appears to be abysmal. A \$100 investment in this sector made in 1989 only barely appreciated in value on net over the 14-year span. Finally, the index for other sectors (the larger categories of which are entertainment, food, chemicals,

plastics, waste management, and transaction and financial services) shows continuous appreciation at moderate rates, with no overshooting in 1999-2000.




Figure 3. Indexes for Four Industry Groups

Figure 4 shows the estimated value of all companies in the venture-type universe. This calculation includes all of the data that enters the calculation of the index itself plus data on the values of other companies where we observe only a single valuation and cannot incorporate the data in the index. Total value of this type of private equity peaked at about \$1.3 trillion and is currently about \$500 billion.




Figure 4. Value of Venture-Type Equity

Figure 5 shows the total amount raised by venture-type companies over the period covered by our database, starting in 1987. We stress that we measure the inflow of new funding to the companies themselves, not the inflow to venture funds. Fundraising reached a sharp peak at the beginning of 2000 and fell back fairly rapidly to normal levels of about \$2 billion per month.




Figure 5. Money Raised by Venture-Type Companies

Figure 6 shows the cash received by shareholders as companies exited the venture-type universe, by way of IPO or acquisition (we are not aware of any venture-type company that has paid a dividend, which would belong in this measure as well). Cash to shareholders reached a high peak during the period of frequent IPOs, in 1999 and 2000 and then subsided to low levels.




Figure 6. Cash Received by Shareholders of Venture-Type Companies

We have examined the returns to venture-type investments by regressing the returns shown in Figure 2 on the monthly returns to the Nasdaq index. Because our index is constructed from returns spanning a year or sometimes more, we expect our returns to lag behind Nasdaq, so we include Nasdaq returns over the past 18 months (there is some evidence of correlation with future Nasdaq returns and with returns after more than 18 months, but not a great deal). Figure 7 shows the regression coefficients and standard errors for the resulting regression coefficients. The coefficients are largest for lags of 2 to 7 months, as expected from the construction of the index.




Figure 7. Coefficients of the Regression of Index Returns on Nasdaq Returns

The constant in the regression is 71 basis points per month or about 8.5 percent per year, with a standard error of 68 basis points. There is mild evidence in favor of the proposition that venture-type investments have higher returns, risk-adjusted, than does Nasdaq, but the magnitude is not as high as some venture boosters have suggested. We also note that our measure of returns does not deduct the fees and “carry” that venture funds impose on limited partners.

The sum of the coefficients shown in Figure 7 is 0.86, implying that the portfolio of venture-type investments captured by the index is a bit less risky than is a portfolio that matches the Nasdaq index.

It is also interesting to compare the quarterly returns on the Sand Hill Index to the quarterly returns reported by the general partners of venture funds. Cambridge Associates calculates an index as the simple unweighted average of the ratio of change in value over a quarter to the valuation at the end of the previous quarter for funds that report to the company. Table 2 compares the two indexes. The Cambridge index is 124 basis points

higher, reflecting the influence of selection and other differences in coverage. The Cambridge index is substantially more volatile, with a standard deviation almost 6 percentage points higher. Our index leads the Cambridge index, in the sense that the forecasting power of our index for the Cambridge index is greater than the forecasting power of the Cambridge index for our index. Adding our index to a univariate forecasting equation lowers the standard deviation of the forecast of the Cambridge index by 172 basis points while the Cambridge index improves a forecast of our index by only 28 basis points. In all of these forecasting equations we use three lagged values of both variables.

	<i>Returns on the Sand Hill Index</i>	<i>Returns from Cambridge Associates</i>
Mean at quarterly rate	4.62	5.86
Standard deviation of quarterly return	8.86	14.59
Standard deviation of one-quarter- ahead forecast, using only lagged values of the same index	3.54	12.05
Standard deviation of one-quarter- ahead forecast, using lagged values of both indexes	3.26	10.33

Table 2. Comparison of the Sand Hill Index to Cambridge Associates' Index

VI. Performance Evaluation for Venture Holdings

The standard metric for traded security investments is the risk-adjusted average return. We use the same metric for venture-type investments—the value of α in the regression equation

$$r_t = \alpha + \beta r_{M,t}.$$

We have dropped the subscript i now that we are focusing on the evaluation of the performance of a given portfolio. r_t is the monthly return to the portfolio and $r_{M,t}$ is an unbiased, value-weighted, monthly index of the market—here, the market of venture-type investments. The index is the first to meet these standards.

To develop data on the monthly returns to a given portfolio, we start with data on valuation events for the portfolio and compute an index for the portfolio using the same basic method as outlined in the previous section. We do not need to impute data for missing round valuations unless this information has not been provided by a venture fund to an investor. (In our experience, even in the best GP reports, about 10 percent of the private rounds and 20 percent of acquisitions lack valuation. Many of these “acquisitions” are transactions in which some intellectual property of the company was sold to another company, but on net assets were worth no more than liabilities.) As we do in constructing the index itself, we use the index to interpolate values for each company in the portfolio between valuation events. We make use of our proprietary data on the current status of individual companies to estimate remaining values and we use the other tools described earlier.

As with traded securities, the regression of the portfolio returns on the index yields estimates of the key parameters β and α . Beta measures for the portfolio’s risk relative to the index and α measures the portfolio’s risk-adjusted average excess return. These metrics are the *lingua franca* for analysts and investors in traded securities. It is our aim to extend their use to venture-type investments.

The key result of this evaluation is the risk-adjusted average return, α , and its standard error, to assess its statistical reliability. We also estimate equations for multi-factor CAPMs, with the returns from the S&P500 or Nasdaq, or both, as additional factors.

VII. Preferences

One feature of the venture investment that deserves further discussion is the role of preferences in determining the value of the financial claims held by the outsiders in a venture-backed firm. The outsiders hold convertible preferred shares, which in principle have a higher value than the common shares held by the insiders. We believe that the preferences have only a small value in practices.

A venture investment is a contract between outside investors—either individual angel investors or limited partners in a venture fund—and the founders of a new company. Most venture deals have similar terms. Founders hold common stock. Outside investors buy convertible preferred stock, not common stock. The preference requires that the corporation pay available cash to the investors up to a multiple (usually at least double, but sometimes more) of their original investment before any cash goes to the founding shareholders, who hold common shares. Upon a favorable event such as an IPO or a happy sale of the company, the outside investors' convertible preferred stock converts to common stock, equivalent to what is held by the founders.

The primary purpose of the asymmetry of the claims of the founders and the investors is to help offset the danger that the founders might persuade investors to contribute cash to the corporation and then liquidate the corporation in order to gain a share of the cash. The preference cuts off this danger. In principle, the founders will agree to the standard venture deal only if they believe that they have an idea for a company that will generate a payoff well in excess of the amount contributed by the investors.

In the standard venture capital structure, a number of rounds of financing, conventionally labeled A, B, C, and so on, occur before a company becomes self-sustaining or goes out of business. Investors in later rounds typically have the contractual right to receive residual cash to satisfy their preferences before investors in earlier rounds receive any cash. Because of their superior claim, later-round shares are worth more than earlier-round shares, which in turn are worth more than the founders' common shares.

The importance of preferences and the value adjustments they imply rests on how preferences affect investor outcomes. When all shares convert to equally-valued common shares under favorable outcomes, such as an IPO or an acquisition at a good price, the preferences do not matter. If the startup spends all its cash and expires worthless, all shares have zero value and again, the preferences do not matter. Occasionally, a startup reaches a terminal event such as an asset sale at low value or a shutdown before all the cash is spent, in which case the superior claims of the preferred shareholders may influence the distribution of the terminal cash.

When such a liquidation or sale at a low value looms, the original preferences are often re-negotiated to allow a partial recovery of assets. The dynamics are as follows: for the liquidation to occur, a contract-specified majority of the shareholders must agree to liquidate. Generally, this will require the favorable vote of some shareholders—possibly the founders and some outside preferred investors—who would receive nothing if the liquidation occurred according to the strict terms of the preferences. Because these shareholders have an option value from the continuation of the business, they will vote in favor of the liquidation only if they receive enough of the cash from the liquidation to more than compensate for their option value. The resulting bargain diminishes the value of the liquidation preference enjoyed by the preferred shareholders.

The standard practice in the valuation of venture-backed companies is to neglect the differences in value among the rounds of venture shares and between those shares and the common shares held by the founders. We follow this practice, though it slightly overstates our valuation of entire companies. Our empirical study of these biases suggests they are small. The primary purpose of the liquidation preference is not to tailor the risk characteristics of the shares sold to outside investors, but to prevent the founders from taking cash out of the company through liquidation.

VIII. Concluding Remarks

We have built an index for venture capital for two important purposes. First, the index provides a measure of month-to-month market-wide movements in value. Investors can now compare the returns on their holdings with market-wide returns over exactly the same interval, as well as update the estimates of value for their holdings based on market-wide movements in value. Second, the index allows the standard methods for performance analysis for traded securities to be applied to their untraded venture-type investments. In the process, we have overcome two major obstacles: First, hard data on company valuations are available only episodically, when the companies sell new rounds of equity or go public, are acquired, or go under. We have developed a sophisticated interpolation method to deal with episodic data. Second, valuation data are not made public in any systematic way and favorable valuations are more likely to become public than unfavorable ones. We have developed sophisticated methods for imputing values to companies with known rounds of funding but missing valuations. We have also attacked this problem by gathering a large volume of new data.

By applying our tools first to our extensive database of venture-type returns to obtain value-weighted market-wide returns to these investments each month since January 1987, and then to the investment history of a portfolio, we can apply the standard approach offered by finance theory to the evaluation of the performance of the portfolio.

These evaluation methods represent a step forward from practices venture investors were forced to rely upon because they had no regularly reported benchmark of market-wide value. Information about the mean and variation in IRRs among venture funds is better than no information at all, but it is not as informative as what can be learned by applying the standard tools of portfolio analysis with a true periodic index of value.

The index and the tools to apply it can improve venture investors' decisions. They can improve portfolio asset allocations, improve decisions regarding portfolio contributions and disbursements by basing them on more accurate estimates of value,

improve risk management, and give investors a better indication of the relative performance of their holdings as well as the relative performance of the past efforts of prospective asset managers.

References

- Bailey, Martin H., Richard F. Muth, and Hugh O. Nourse. 1963, "A Regression Method for Real Estate Price Index Construction" *Journal of the American Statistical Association*, 4: 933-942.
- Cochrane, John H. 2003. "The Risk and Return of Venture Capital" Graduate School of Business, University of Chicago. October
- Heckman, James. 1979. "Sample Selection Bias as a Specification Error," *Econometrica*, 47(1): 153-161. February.
- Quigley, John M., and Susan E. Woodward. 2003. "An Index for Venture Capital". University of California, Berkeley. July.