FEDERAL RESERVE BANK *of* ATLANTA

Expected Returns to Stock Investments by Angel Investors in Groups

Ramon P. DeGennaro and Gerald P. Dwyer

Working Paper 2010-14
August 2010

WORKING PAPER SERIES

FEDERAL RESERVE BANK *of* ATLANTA WORKING PAPER SERIES

Expected Returns to Stock Investments by Angel Investors in Groups

Ramon P. DeGennaro and Gerald P. Dwyer

Working Paper 2010-14
August 2010

Abstract: Angel investors invest billions of dollars in thousands of entrepreneurial projects annually, far more than the number of firms that obtain venture capital. Previous research has calculated realized internal rates of return on angel investments, but empirical estimates of expected returns have not yet been produced. Although calculations of realized returns are a valuable contribution, expected returns, rather than realized returns, drive investment decisions. We use a new data set and statistical framework to produce the first empirical estimates of expected returns on angel investments. We also allow for the time value of money, which previous research has typically ignored. Our sample of 588 investments spans the 1972–2007 period and contains 419 exited investments. We conduct extensive tests to explore potential bias in the data set and conclude that the evidence in favor of bias is tenuous at best. Our results suggest that angel investors in groups can expect to earn returns that are on the order of returns on venture capital investments. Estimated net returns are about 70 percent in excess of the riskless rate per year for an average holding period of 3.67 years. This estimate is reasonable compared to Cochrane's (2005) estimate of 59 percent per year for venture capital investments, which tend to be in lower-variance, later-stage projects. Returns have a large variance and are heavily skewed, with many losses and occasional extraordinarily high returns.

JEL classification: G24, G20

Key words: angel investor, expected return, private equity

The authors thank the Ewing Marion Kauffman Foundation for providing the data and Rob Wiltbank for help in clarifying some of the data's features. The authors are grateful for helpful comments from Stephen M. Miller, Bruce Mizrach, and Paul Thistle as well as from participants in Ewing Marion Kauffman Foundation and Federal Reserve Bank of Cleveland Conferences on Entrepreneurial Finance; the European Financial Management Association's Symposium on Entrepreneurial Finance and Venture Capital Markets; conferences organized by the Association for Private Enterprise Education and the Society for Nonlinear Dynamics and Econometrics; and the College of Charleston and the University of Nevada at Las Vegas. DeGennaro thanks the University of Tennessee College of Business and the university's Scholarly Activities Research Incentive Fund for support. Dwyer thanks the Spanish Ministry of Education and Culture for support of project SEJ2007-67448/ECON. The views expressed here are the authors' and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any remaining errors are the authors' responsibility.

Please address questions regarding content to Please address questions regarding content to Ramon DeGennaro, University of Tennessee, Knoxville, TN 37996, rdegenna@utk.edu, 865-974-1726, 865-974-1716 (fax), or Gerald Dwyer, Federal Reserve Bank of Atlanta and the University of Carlos III, Madrid, Research Department, 1000 Peachtree Street, N.E., Atlanta, GA 30309-4470, gerald.p.dwyer@atl.frb.org, 404-498-7095, 404-498-8810 (fax).

Expected Returns to Stock Investments by Angel Investors in Groups

I. Introduction

An angel investor can be defined as a person who provides funds to a private business which does not have publicly traded stock and is owned and is operated by people who are not relatives or friends of the investor. Acting as informal venture capitalists, angels invest billions of dollars in thousands of fledgling companies annually. What returns can these angels expect to receive on their investments? Although previous work has explored realized returns, this paper is the first to obtain estimates of expected returns on angel investments in a form comparable to reported expected returns on stock or venture capital.

Until recently, research on the returns to investments by angel investors and angel groups has been limited because suitably large data sets simply have not been available. For example, Goldfarb, Hoberg, Kirsch and Triantis (2009) have just 32 angel-only investments in their study of private equity, and not all of them are exited investments. The Angel Investor Performance Project recently has produced an informative database on angel investments. In contrast to Goldfarb, Hoberg, Kirsch and Triantis (2009), these data have 588 investments of which 419 are exited investments.

We use these data to explore the expected returns on angel investments. Our paper is similar in spirit to Cochrane (2005), who estimates the returns on venture capital investments, and to Barnhart and Dwyer (2009), who estimate the returns on traded stock in new industries. It differs from Wiltbank (2005) and Wiltbank *et al.* (2008) because we estimate expected returns rather than realized returns. Thus, our paper combines these strands of the literature by estimating expected returns on angel investments. The distinction between realized internal rates of return and expected returns is critical. Realized internal rates of return do not drive financial decisions. Expected returns drive financial decisions.

II. Prior Literature

Angel investors and their investments are not well documented. This is partly because individual investments tend to be informal, so there is little or no documentation or data. Also, practitioners and academics have not reached a consensus concerning even the definition of an angel investor. Shane (2008, 2009) provides a wealth of institutional details along with a good review of angel investors and their investments. Shane shows how the lack of consensus leads to divergent conclusions about the size and nature of the angel market. Our definition is: An angel investor is a person who provides funds to a private business which does not have publicly traded stock and is owned and is operated by people who are not relatives or friends of the investor. This definition is the same as Shane's (2009, p. 14) except that it explicitly limits angel investors to investors in firms that do not have publicly traded stock.[1]

Wiltbank *et al.* (2008) report that angel investors invest their capital directly in early-stage ventures, in many more businesses than do formal venture capital firms and usually in much smaller dollar amounts. They also note that angel investors are often the first outsiders to supply equity capital to entrepreneurs trying to build a business, even before formal venture capital is obtained. Formal venture capitalists invested less than two percent of the total capital in seed-stage companies during the last ten years.

Wiltbank and Boeker (2007a) also report an important trend: individual investors increasingly are forming angel investor groups. Wiltbank and Boeker list several advantages to group membership and investment including shared expertise and diversification. According to

[1] Otherwise Warren Buffett would be an angel investor in Goldman Sachs, which is not the sort of investment generally considered to be an angel investment. We thank Stephen Miller for noting this.

the Angel Capital Education Foundation, about 10,000 accredited angel investors belonged to 265 angel groups as of 2007.

Shane (2005) reports information gleaned from four focus groups of angel investors. These angel investors are diverse in their backgrounds, their motivations for investing and in their investment approaches. Not surprisingly, their investments also are diverse.

Mason and Harrison (2002) and Wiltbank (2005) are among the first to explore the returns on angel investments. Mason and Harrison's data are for the U.K. and Wiltbank's data are for the U.S. Their results are broadly consistent: returns less than the investment predominate with a few spectacular returns of 100 percent net of investment or even much more. Compared to U.K. returns, U.S. returns are more extreme with more returns less than the investment and more net returns above 100 percent.

Much of the research on angel investors examines management strategies. For example, Wiltbank et al. (2008) examine 121 angel investors who made 1,038 new venture investments. They distinguish between prediction and control strategies. Investors who use prediction strategies attempt to predict events and position themselves to capitalize on them. In contrast, investors who use control strategies focus on the subset of events that they believe they can control and optimize accordingly. Sarasvathy (2001) suggests that investors have no reason to predict events to the extent that they can control them. Wiltbank et al. (2008) report that angels who use prediction strategies tend to make significantly larger venture investments. Angels who use control strategies have fewer failures.

Goldfarb et al. (2009) study data from the failed law firm of Brobeck, Phleger & Harrison. The data contain 182 deals involving a first issue of equity made from 1993 to 2002. Of the 182 deals, there are just 32 angel-only deals. Goldfarb et al. report four main results. First, if the capital requirements are small, then the deal can be angel-only, venture capital-only, or mixed. Larger deals include venture capital participation. Second, in first-round financing, called Series A deals, angels almost always take preferred shares. In these data, deals with angel investors have weaker control rights even after controlling for size, age and other variables that might affect risk preferences. Third, among smaller deals, angel-only deals fail the least often. This could be due to what Goldfarb et al. call "inactive" firms, firms that have not officially failed but which have no positive value. Fourth, large deals with only venture capital tend to be more successful than large mixed deals. Goldfarb et al. speculate that some deals require a sufficiently large amount of funds that venture capital cannot be excluded even if such funding seemingly lowers the probability of success.

Goldfarb et al.'s study requires a major qualification: They treat angel groups as venture capital firms. They say that there are "a small number" of such deals. They add that their results are robust to the classification of investor classes. Still, their conclusions may not apply to angel groups.

Moskowitz and Vissing-Jorgensen (2002) study the returns on private equity investment. They report that private-equity returns are no larger and may very well be smaller than returns on public equity portfolios. This is true despite the greater concentration of risk and higher variance of private-equity returns. This suggests a private-equity premium puzzle with private equity paying too little, which is the opposite of the public-equity premium puzzle. For an investor with a relative risk-aversion coefficient of two, private equity held in such concentrated amounts must return ten percent more than public equity portfolios to offer fair compensation for the level of risk borne. Moskowitz and Vissing-Jorgensen report that the shortfall is about $460,000 during the working life of an entrepreneur. They suggest six possible reasons why entrepreneurs might accept this apparent reduction in their wealth. These are optimal contracting and moral hazard, higher risk tolerance, other pecuniary benefits, nonpecuniary benefits, skewness preference and over-optimism.

Our paper is similar in spirit to Cochrane's paper (2005), which examines the distribution of returns on venture capital investments from 1987 through June 2000. He finds that his sample of venture capital investments has expected proportional returns on the order of 59 percent per year. He concludes that venture capital investments and the smallest NASDAQ stocks have roughly similar returns and return volatilities during his sample period. Barnhart and Dwyer (2009) show that expected returns to investors in stocks in new industries are positive and approximate those of market returns. *Ex post* returns reflect infrequent but very large gains and frequent but smaller losses, and the payoffs are broadly consistent with a log-normal distribution of expected returns.

III. Data

This paper uses data from the Angel Investor Performance Project (AIPP).[2] Wiltbank and Boeker (2007a) report that these data contain survey responses from 86 angel groups totaling 539 investors who had made 3097 investments. Exits had been achieved for 1137 of those investments. Not all of these investors provided data for the variables used in our analysis. We find 588 useable investments in the data. Of these, 419 have exited. Wiltbank and Boeker discuss non-response bias and other possible problems with survey data but provide evidence that the AIPP data are relatively free of such problems. They conjecture that the most likely source of bias is that respondents might report good investment outcomes and neglect poor outcomes. Although the response rate by groups is 31 percent, the average response rate of members of

[2] Information and the underlying data are available at http://www.kauffman.org/aipp.

those groups is only 13 percent. With such a low response rate, the concern is that responders may not be representative. For example, responders may be those with higher returns. Wiltbank and Boeker, however, report that the distribution of reported returns is similar for groups with response rates of 60 to 100 percent as for groups with a much lower response rates. If the distribution of returns is the same across these groups, then this provides little evidence of correlation of return with reporting frequency.

The potential for sample-selection bias compared to other samples remains, though. The AIPP angels are not representative of angels in general because all of the AIPP angels are members of angel groups and are accredited investors. The restriction of any conclusions to angel groups when the data cover only angel groups is inevitable. In addition, AIPP investments are all equity, whereas Shane (2009, p.80) finds that 40 percent of the dollar value of angel investments is debt. Debt investments by angel investors may generate lower expected returns than the AIPP data represent. We attack the issue of selection bias along two lines. First, we compare the return multiples from the AIPP to other reported measures. For the most part, the AIPP numbers are similar to those reported by other researchers. Second, we compare other AIPP variables such as the proportion of IPOs, buyouts and failures to other datasets. We find only relatively small deviations between the AIPP and other reported values.

Because these tests are lengthy, we provide details in Appendix I and a brief summary in the text. There is some evidence of selection bias, but the evidence is far from overwhelming. The most convincing evidence of bias is the high percentage of AIPP investments that result in an IPO, which may inflate the reported performance of the angel groups that participated in the AIPP.[3] However, implied internal rates of return and return multiples are comparable with

[3] A recent paper by Ball, Chiu and Smith (2008) identifies forces that drive the choice between venture capital exits by IPO or by acquisition. Such forces could drive angel exits, too, and these could account for the differences

figures from other sources, especially on a value-weighted basis. In addition, the data survive several checks designed to detect bias. For example, we might expect larger deals to be less prone to bias because survey respondents are more likely to remember those deals correctly. We find, though, that multiples from deals larger than the median do not differ statistically from those smaller than the median. Deals with more than one angel group member might be more likely to be reported accurately, yet we find no evidence from the number of co-investors that the annual return multiple is biased upward.

The financial variables in the AIPP are the key for our purposes. Because the AIPP reports information about initial investments, intermediate cash flows and exit values, we are able to compute both total returns and internal rates of return. From these, we can then estimate expected returns.

Figure 1 diagrams cash flows between the angel investor and the entrepreneur in these data. Flows above the horizontal time line are cash flows from the entrepreneur to the angel investor. Flows below the line are cash flows from the angel investor to the entrepreneur. Investments in the project made after the initial investment are called follow-on investments. The AIPP calls the first of these *follow1invest* and the second *follow2invest*. In both cases the data contain the dollar amount and date of these investments. The AIPP sums any subsequent investments and reports the total as *followxinvest*. For these, the dollar amount is available but not the year of the investment. These three cash flows are reported 117, 24 and eight times, respectively. For exited projects, these cash flows occur 106, 19 and seven times, and all reported values are positive. *Midcash* is the total dollar amount of any cash flows from the project to the angel investor before the exit date. The AIPP does not report the year of these cash flows. Midcash is positive for 78 investments, and all of these are exited investments.

---

between the exit proportions of Band of Angels and the AIPP dataset. We have no way to control for these factors.

Most previous work on angel investments uses base multiples, defined as total cash flows to the investor divided by total investments. This has the advantage of simplicity but ignores the opportunity cost of the timing of cash flows. We use present values for all computations. We do this by discounting these cash flows from the year in which they occur to the year of the initial investment at the risk free rate. This means that all expected returns are excess returns: returns in excess of the risk free rate.

Using the risk free rate may seem strange given that angel investments are risky. Our use of the risk-free rate does not mean that our calculations assume that the investor would invest in a risk-free asset if the angel investment were not available. The purpose of discounting is to bring all cash flows to a common date. Appendix II contains the details of the present value calculations.

IV. Issues in Estimating Expected Returns

The purpose of our paper is to estimate the expected returns on angel investments. This is not as simple as it might seem because the observed returns on angel investments are influenced by choices that investors make. What sorts of the choices are they?

A. The Investor's Exit Rule

We assume that investors maximize their expected utility of consumption, which is increasing in the payoff from the investment. We assume that the investment lasts a finite period.

The investor's choice is an application of investment theory tracing to Frank Knight and Irving Fisher. The expected present value of utility of the cash flows is

$$U = \mathrm{E}\, e^{-\delta T} u\left[S\left(T\right)\right] - \mathrm{E} \int_{t=0}^{T} e^{-\delta t} u\left[I(t)\right] dt \qquad (1)$$

where E denotes expected value, $e^{-\delta t}u\left[.\right]$ is the investor's utility function for each period discounted t periods by the discount rate δ, *S(T)* is the value of the investor's payoff at *T*, *I(t)* is the value of the investment at *t*, the date when the investment begins is *t*=0 and *T* is the project's termination date. The investor chooses *I(t)* and *T* to maximize this expected present value. For simplicity, we set *I(t)* to a constant *I* determined at the start of the project, in which case equation (1) is

$$U = \mathrm{E}\, e^{-\delta T} u\left[S\left(T\right)\right] - u\left[I\right] \qquad (2)$$

Making *I(t)* variable would complicate the algebra without adding any insights. Maximizing equation (2) with respect to *T* implies that the partial derivative is set to zero, which implies after some rearrangement that

$$\frac{\mathrm{E}\left[\dfrac{du\left[S\left(T\right)\right]}{dS\left(T\right)}\dfrac{dS(T)}{dT}\right]}{\mathrm{E}\, u} = \delta. \qquad (3)$$

This states that the expected increase in the utility of the value of the project relative to the expected utility of the value of the project equals the discount rate at the end of the project. If the utility of the value of the investment is increasing at a decreasing rate ($d^2u\left[S\left(T\right)\right]/dS\left(T\right)^2 < 0$) and the value of the investment as a function of time is decreasing at a decreasing rate ($dS\left(T\right)/dT < 0$), then this solution for the stopping time is a maximum. Researchers commonly suppose that marginal utility is decreasing in wealth. Even so, it is possible and even plausible that the value of the investment increases at an increasing rate before it starts to increase at a decreasing rate, which could result in the expected increase in marginal utility increasing for a time before it decreases.

Discounting the future value *S(T)* does not change the basic message. With the value *S(T)* discounted at the risk-free rate r with $\bar{S}(T) = e^{-rT} S(T)$ and $u' = du[S(T)] / dS(T)$, equation (3) evaluated at T is

$$\frac{u'}{u} e^{r(T-t_1)} \left[r\bar{S}(T) + \frac{d\bar{S}(T)}{dT} \right] = \delta. \tag{4}$$

Although more complex than Equation (3), this represents exactly the same condition evaluated at T.

The choice to exit with zero payoff is not necessarily included in this characterization of the optimal exit time. That choice also will be made if it maximizes expected utility. Exiting the project with a zero payoff can be included by generalizing the condition for exit to

$$\begin{aligned} &\bar{S}(T)\left(\frac{u'}{u} e^{r(T-t_1)} \left[r\bar{S}(T) + \frac{d\bar{S}(T)}{dT} \right] \right) = \delta \\ &\bar{S}(T) \geq 0 \end{aligned} \tag{5}$$

which includes the prior solution for a positive ending value as a special case.

The overall implication is that the exit decision is exogenous. The entrepreneur maintains his investment in the project until it no longer earns the opportunity cost of the investment. Although the form of the exit may vary, the timing of the exit is exogenous.

B. Computing Expected Returns

Though the computation of internal rates of return from actual investments is straightforward, the calculations are more subtle and complex in the context of expected returns with different project lives. Average internal rates of return per year are average returns per project lasting T years. This is given by $\mathrm{E}\, R = \mathrm{E}\left[S(t+1) / S(t)\right] - 1$. For angel investments in Wiltbank and Boeker's (2007b) data, average internal rates of return are on the order of 27 percent annually. Expected returns may well be similar to these actual returns, although this is not inevitable.

A large proportion of angel investments return nothing (see DeGennaro, 2010). Standard Brownian motion for the value of equity is inappropriate for angel investments because the value of the equity can never be zero if it can be represented as a logarithm. To circumvent this complication we start the analysis from standard Brownian motion for the value of the firm's *assets A(t),* rather than the equity. The firm's assets consist of debt *D* and equity *S*. For analytical purposes, we assume that the firm ceases operations if the value of the firm's assets is less than the value of the debt. This guarantees that the value of assets is non-negative and except for a set of measure zero, positive. The associated diffusion equation and conditions are

$$dA(t)/A(t) = \mu dt + \sigma dB(t), \quad A(t) = D(t) + S(t), \quad A(t) \geq D(t), \tag{6}$$

where μ is the expected return on the assets, σ is the underlying return volatility and *B(t)* is a standard Wiener process. Ito's lemma implies that

$$d\ln A(t) = \left(\mu - 0.5\sigma^2\right)dt + \sigma dB(t), \quad A(t) \geq D(t), \tag{7}$$

This can be written in terms of the ratio of the value of equity to debt, *f*,

$$d\ln S(t)\left(1+f\right) = \left(\mu - 0.5\sigma^2\right)dt + \sigma dB(t), \quad A(t) \geq D(t), \tag{8}$$

which can be rewritten as

$$d\ln S(t) = \left(\mu - 0.5\sigma^2\right)dt - d\ln\left(1+f(t)\right) + \sigma dB(t), \quad A(t) \geq D(t), \tag{9}$$

If we assume that the ratio of the value of debt to equity primarily is a drift term such as $d\ln\left(1+f(t)\right) = \mu_f dt$ and that the mean equity return is $\mu_s = \mu - \mu_f$, then the evolution of the value of owner's equity is given by

$$d\ln S(t) = \left(\mu_s - 0.5\sigma^2\right)dt + \sigma dB(t), \quad A(t) \geq D(t). \tag{10}$$

These equations indicate that the expected return for firms that end with a positive value is $\exp\left(\alpha + 0.5\sigma^2\right)$, where α is the log return for firms that end with positive value. If firms that are terminated with a value of zero have the same variance of returns as other firms but a different ex post realization,[4] then an estimate of the expected return is

$$\mu_s = p\exp\left(\alpha + 0.5\sigma^2\right), \tag{11}$$

where p is the probability that a firm will end with a positive return.

Maximum likelihood is a natural estimator of the expected return. Campbell, Lo and MacKinlay (1996), Gourieroux and Jasiak (2001), and Tsay (2002) provide the maximum likelihood estimators of these parameters, which are

$$\hat{\alpha} = \frac{1}{T}\sum_{t=1}^{T} r_t$$

$$\hat{\sigma}^2 = \frac{1}{T}\sum_{t=1}^{T}[r_t - \hat{\alpha}]^2 \tag{12}$$

$$\hat{\mu}_s = \hat{\alpha} + (\frac{1}{2})\hat{\sigma}^2$$

where $r_t = \Delta \ln S_t, T = \sum_{i=1}^{N} T_i$ and t = 1, ..., T. The index t spans all project-years, so that an observation is the log return for a year for a project.[5] Barnhart and Dwyer (2009) provide details.

One benchmark for the return is the average annual return on stocks, which is consistent with using each year of a project as an observation. If one is interested in comparing the expected return in a typical angel investment for a year to stock returns, the returns should be weighted by the number of years that the investment is ongoing. One ten-year project earning 10 percent annually is equivalent to 10 one-year projects earning 10 percent. This is in contrast to computations by Wiltbank and Boeker (2007b), for example, who weight each investment the

[4] We offer empirical support for this assumption in the section that reports estimated expected returns.

same regardless of duration.

V. A First Look at the Data

Table 1 presents summary statistics on the underlying data used in this paper. We use 588 observations from the AIPP data set that report the dollar amount and year of the initial investment in each project. Of these 588 investments, a total of 419 are exited and 169 are not exited as of the end of the data in 2007. If the project is exited, then the observation contains the year of exit and exit cash. Investments made in the project after the initial investment are called follow-on investments.[6] The data contain the dollar amount and date of up to two such investments. The AIPP also reports the total dollar amounts of any subsequent angel investments but does not report the year or years of those investments. There are eight projects with more than two investments. We know that any additional investment occurs after the second investment and before the project is exited or else before the end of the dataset in 2007. The data also include the dollar amount of any cash flows from the project to the angel investor before the exit date but provide no information on the date. Such cash flows are reported 78 times. We know that these mid-cash payments occur between the year of the initial investment and the year of exit or the end of the dataset in 2007 but not the exact date.[7]

The maximum invested in projects by a single investor never is more than $5.1 million and the mean is only about $147 thousand. Even this number is misleadingly large: the median amount invested is only about $49 thousand. These are quite small investments by venture

---

[5] This is different than Wiltbank and Boeker (2007b), who average net arithmetic returns across projects.
[6] The AIPP calls the first of these follow-on investments *follow1invest* and the second *follow2invest*. There also is a *followxinvest* for the sum of any additional investments.
[7] Projects with follow-on investments or mid-cash flows tend to be a bit longer than those without them. Investments with mid-cash payment last an average of 4.7 years compared to 3.3 years for those with no mid-cash payments (*t*-statistic = 3.6). Projects with one follow-on investment last 4.2 years versus 3.4 years for those without (*t*-statistic = 3.1).

capital standards. Cochrane (2005) reports an average of $6.7 million per round in his sample of venture capital investments.

The total cash received by investors from exited projects ranges from zero to $33 million and also is heavily skewed to the right, with a mean of $490 thousand and a median of only $42 thousand. Total investments and cash received take no account of the timing of the cash flows and neither does the base multiple commonly used by practitioners. Because of its wide use, we include the base multiple in Table 1 even though it takes no account of the time value of money. The base multiple computed as the present value of receipts relative to the present value of investment, or *PV base multiple*, shows the ratio of receipts relative to investment with both valued at a common date. The median PV base multiple is less than one, indicating that more than half the projects return less than their investment. However, the mean PV base multiple is 7.37, well above one. The maximum PV base multiple is 1,333, which is extraordinary.

The net returns and log returns tell a similar story of frequent low returns and infrequent high returns. The net return is the internal rate of return on the project, with all funds discounted back to the investment date at the risk-free rate. The mean net return of the 408 completed projects with a duration of a year or more is 29 percent, but the median is -2 percent with a range from 0 to 3,241 percent. The log returns, which necessarily exclude the completed projects with zero payoff, have a positive median which is less than the mean, but not by as large a value as the net returns.

Aspects of the investments themselves are interesting. The data indicate that investors spend a fair amount of time on due diligence before investing. There is positive skew in these numbers, as there is for payoffs. The mean number of hours of due diligence is 65.4, with a median of only15 and a reported range from zero to 5,000 hours. Surprisingly, the data provide no support for the notion that angel investors tend to invest more time on due diligence as investment amounts increase. The correlation of due diligence and the initial investment is only 0.04 with a *p*-value of 0.68 and the correlation of due diligence and total investment it is only 0.03 with a *p*-value 0.71. What *does* attract an angel investor's attention is the percentage of his wealth that he invests in angel investments. The correlation between the amount of due diligence and the percentage of an angel's wealth invested is 0.15 (*p*-value 0.09). Angel investors apparently take extra care to evaluate projects if they are likely to be disproportionately invested in startup and early-stage companies relative to the rest of the market portfolio.

For projects that include stage information, almost half are startups. Seed and startup companies are more than three-quarters of all projects. Turn-around firms and late-stage companies are the smallest proportion of the investments. The data also indicate that investors are engaged with the firms in most instances, with over three quarters of investors interacting with the firm once a quarter or more and 23 percent of investors interacting with the firm weekly or daily.

These data also support the conventional wisdom that members of angel investor groups are seasoned investors with extensive business experience. Angel investors who reported participating in AIPP projects have been making angel investments for an average of about 11.3 years. This masks a broad range of experience as angels, though. The angel investors in the data have as few as one year's experience to as many as 49 years' experience as angels. They average 16.2 angel investments each and on average have exited from about 7.1 investments.

Angel investors clearly do not completely diversify away the firm-specific risk in these investments. The mean percentage of wealth invested in angel investments is 13 percent, with a median of 10 percent, reflecting a positive skew. The range of the percent of wealth invested is from one to 85 percent. Some angel investors need not be concerned with the idiosyncratic risk of these investments; others are not likely to ignore the idiosyncratic risk of their angel investments.

These angel investors are highly educated. Almost 70 percent of those who reported their academic degrees have advanced degrees. Of the 132 angels who reported their academic degrees (not including one “Other”), 40 had bachelors degrees, 67 had masters degrees, 10 held law degrees and 15 had earned a Ph.D.

How do these angel investors exit from their investments? Table 2 gives the answer. Of the 408 exited projects that last at least one year, 114 have the unhappy result that the firm ceases operation. Selling the firm is by far the most common way that angels earn a return; in 180 cases another company buys the firm and in an additional 20 cases other investors buy the firm. About a quarter as many firms – 56 in all – exited by means of an initial public offering (IPO). The fraction of IPOs for these exited firms lasting at least one year is about 13.5 percent, which is relatively high compared to other estimates of the fraction of angel investments that culminate in IPOs. We investigate this in the appendix. Here, we note that 13 of these IPOs have investments in the same company in the same year and exits in the same year. Thirteen members of a group saw a deal, bought in at the same time and later reported the deal to AIPP. If this single deal had exited some other way, then the IPO percentage would drop to just over 10 percent of the exited projects and a little more than seven percent of the full sample.

Table 2 also shows the net annual returns by type of exit. Not surprisingly, investments in firms that cease operating have the lowest net return. These firms lose about 93 percent of their value annually. Firms that have an IPO or are bought by another operating firm have the highest average returns, about 94 percent annually.

Table 3 shows the net returns and log returns by industry. We aggregate the data into the categories listed in Table 3. Some of the "industries" listed in Table 3 such as "Business products and services" are in the original data and some such as "Internet" are an aggregation of various industries provided on the questionnaires. About half of the exited projects have industry missing and the industry with the highest net return is "Missing", so not too much should be made of this information. Nonetheless, the distribution of investments is interesting. The largest industries in terms of completed projects are software, consumer products and services and business products and services. Business products and services is the industry with the highest mean return and IT services is the industry with the lowest mean return.

Figure 2 shows the log returns for the projects by duration. The underlying data are investments and payoffs, which we have reduced to a common date by the risk-free rate. Figure 2 shows the average log return in excess of the risk-free rate for each exited project by the durations of the projects in years. The range of the log returns decreases as the duration of projects increases, which suggests that the standard deviation decreases. Table 4 shows that the standard deviation does indeed decrease as the duration of projects increases. This is not an artifact of the existence of fewer projects for longer durations. The number of projects with positive payoffs and durations of one to five years is 58, 49, 49, 51 and 33. Actually, the decline in standard deviation is predictable. The annual return for T years can be interpreted as the mean annual return for T years. If the standard deviation of the annual return is constant each year, then the standard deviation of the annualized T-period return is the annual standard deviation divided by $\sqrt{T}$ years.

Does the standard deviation diminish in this fashion? Table 4 suggests a qualified yes. Table 4 shows that the one-year projects' standard deviation is outsized relative to other years.

The standard deviations for project longer than one year ranges from 0.11 to 0.80, which is quite a range but far below the standard deviation for one-year projects. As expected, the underlying annual standard deviations do not tend to decline with duration.

VI. Expected Returns

Table 5 presents the estimates of expected return. The first estimate is for firms that last at least one year. Our estimate of the net return in excess of the riskless rate is 69.9 percent per year. This may seem extraordinarily high if compared to returns on stock that have averaged on the order of 11 or 12 percent per year in recent decades before the recent financial crisis. A more appropriate comparison, though, is Cochrane's (2005) estimate of 59 percent for venture capital. Venture capitalists typically invest after angel investors, when risks tend to be lower. Expected returns would be correspondingly lower.

The AIPP reports the duration of projects which last less than a year as having lasted zero years. The second estimated expected return in Table 5 includes these firms using an estimate of half of a year as the projects' durations. As might be expected, this estimate of 71.4 percent is little different than the expected return based on firms that lasted a year or longer. This is because each observation is weighted by its duration, and six projects that last on average half a year can have little effect given the size of the dataset and the duration of some of the projects.

Our model assumes that projects resulting in payoffs of zero have the same variance as projects that result in nonzero payoffs. How reasonable is this assumption? Two parameters identify a normal distribution. The mean and variance of the lognormal distribution imply a unique associated truncated normal distribution and we can compute the implied fraction of observations at the boundary of zero.[8] We use the data for projects that last a year or longer. The

[8] Stephen Miller suggested this in a seminar at UNLV.

lognormal distribution for project-years with a positive ending value has a mean of 0.128 and a variance of 1.010. The associated truncated normal distribution has a mean of 1.893 and a variance of 3.651. This distribution has 16.1 percent of its mass at zero. Compared to the actual value of project-years with zero payoffs of 10.2 percent (10.6 percent including projects that last less than one year), an estimate of 16.1 percent would not dramatically affect our estimate of the expected return from these investments. We conclude that the assumption that the variances are the same for projects paying zero and projects paying nonzero amounts does not distort our estimate.

VII. Conclusion

Angel investors collectively make extremely large investments in start-up firms. In many cases they are the first outsiders to provide equity to new businesses. Research on angel investments has lagged though, because large data sets have been unavailable or proprietary. The Angel Investor Performance Project now offers the opportunity to conduct research that previously has been impossible.

Previous research has calculated realized internal rates of return on angel investments. Although this is an important contribution, internal rates of return are subject to misinterpretation due to nonlinearities and statistical biases. Perhaps more important though, realized internal rates of return do not drive financial decisions. Expected returns drive financial decisions.

Our results suggest that angel investors earn returns that are similar, at least in broad measure, to the returns on venture capital investments and on new industries. We estimate an expected return of about 70 percent per year in excess of the riskless rate. This compares to Cochrane’s (2005) estimate of 59 percent per year for venture capital.

References

Ball, Eric R., Chiu, Hsin-Hui and Smith, Richard L., Exit Choices of Venture-Backed Firms: IPO v. Acquisition (November 13, 2008). Available at SSRN: http://ssrn.com/abstract=1301288.

Barnhart, Cora, and Gerald P. Dwyer. 2009. "Returns to Investors in Stock in New Industries." Unpublished paper, Federal Reserve Bank of Atlanta.

Campbell, John Y., Andrew W. Lo and A. Craig MacKinlay. 1996. *The Econometrics of Financial Markets*. Princeton: Princeton University Press.

Cochrane, John H. 2005. "The Risk and Return of Venture Capital." *Journal of Financial Economics* 75 (January), pp. 3-52.

Cumming, Douglas J., and Uwe Walz. "Private Equity Returns and Disclosure Around the World," forthcoming, *Journal of International Business Studies.*

DeGennaro, Ramon P. "Angel Investors and their Investments." Working paper, 2010.

Goldfarb, Brent, Gerard Hoberg, David Kirsch, and Alexander Triantis. "Does Angel Participation Matter? An Analysis of Early Venture Financing." University of Maryland Working Paper. 2009.

Gourieroux, Christian, and Joann Jasiak. 2001. *Financial Econometrics.* Princeton: Princeton University Press.

Hamilton, Barton H. "Does Entrepreneurship Pay? An Empirical Analysis of the Returns to Self-Employment." *Journal of Political Economy,* June 2000, 108(3), pp. 604-31

Malkiel, Burton and Atanu Saha. "Hedge Funds: Risk and Return," *Financial Analysts Journal*, Vol. 61, No. 6, November/December 2005.

Mason, C. M. and Harrison, R. T. (1996). Informal Venture Capital: A Study of the Investment Process, the Post-Investment Experience and Investment Performance. *Entrepreneurship and Regional Development 8*: 105-125.

Mason, C. M. and Harrison, R. T. (2002). "Is It Worth It? The Rates of Return from Informal Venture Capital Investments." *Journal of Business Venturing, 17,* pp. 211–236.

Moskowitz, Tobias J. and Annette Vissing-Jorgensen. "The Returns to Entrepreneurial Investment: A Private Equity Premium Puzzle?" *The American Economic Review 92*, No. 4 (Sept. 2002), 745-778.

Phalippou, Ludovic and Oliver Gottschalg. "The Performance of Private Equity Funds." *The Review of Financial Studies / v 22* no. 4, 2009, 1747-1776.

Preston, Susan L. *Angel Financing for Entrepreneurs: Early-Stage Funding for Long-Term Success.* Jossey-Bass, 2007.

Puri, Manju and Rebecca Zarutskie. On the Lifecycle Dynamics of Venture-Capital- and Non-Venture-Capital-Financed Firms. Working paper, June 2009.

Reynolds, Paul D. (2007). *Entrepreneurship in the United States: The Future is Now.* Boston, Kluwer Academic.

Sarasvathy, S. D. 2001. "Causation and Effectuation: Toward a Theoretical Shift from Economic Inevitability to Entrepreneurial Contingency." *Academy of Management Review 26,* 243-263.

Shane, Scott A. "Angel Investing: A Report Prepared for the Federal Reserve Banks of Atlanta, Cleveland, Kansas City, Philadelphia and Richmond." (October 1, 2005). Available at SSRN: http://ssrn.com/abstract=1142687.

Shane, Scott A. "Angel Groups: An Examination of the Angel Capital Association Survey." (January 1, 2008). Available at SSRN: http://ssrn.com/abstract=1142645.

Shane, Scott A. *Fool's Gold? The Truth Behind Angel Investing in America.* Oxford University Press, 2009. New York.

Tsay, Ruey. 2002. *Analysis of Financial Time Series*. New York: John Wiley & Sons, Inc.

Wilson, Fred. (November 29, 2007). "Failure Rates In Early Stage Venture Deals." Downloaded from http://unionsquareventures.com/2007/11/failure-rates-i.php.

Wiltbank, Robert. 2005. "Investment Practices and Outcomes of Information Angel Investors." *Venture Capital* 7 (October), pp. 343-57.

Wiltbank, Robert and Warren Boeker. 2007a. "Returns to Angel Investors in Groups. Ewing Marion Kauffman Foundation and Angel Capital Education Foundation. November 2007." 16 pages.

Wiltbank, Robert, and Warren Boeker. 2007b. "Angel Investor Performance Project: Data Overview." 2007 Kauffman Symposium on Entrepreneurship and Innovation Data.

Wiltbank, Robert, Stuart Read, Nicholas Dew and Saras D. Sarasvathy. "Prediction and Control Under Uncertainty: Outcomes in Angel Investing." Forthcoming, *Journal of Business Venturing*, 2008.

Table 1
Sample Statistics

	Number of observations	Mean	Median	Standard Deviation	Minimum	Maximum
All projects						
Total investment	588	146,895	48,868	398,958	1,000	5,100,000
PV total investment	588	144,875	47,575	389,932	1,000	5,096,457
Years held so far	588	3.189	3	3.08	0	35
Exited projects						
Total investment	419	169,089	50,000	454,515	1,000	5,100,000
Total cash out	419	490,022	42,052	2,196,936	0	33,000,000
Base multiple	419	11.95	1.00	87.89	0	1,332.80
Years held	419	3.61	3	3.24	0	35
PV total investment	419	166,407	50,000	443,895	1,000	5,096,546
PV cash out	419	394,205	38,113	166,4952	0	21,876,92
PV base multiple	419	7.37	0.94	41.91	0	1,332.80
Net return	408	0.2938	-0.0176	2.2992	0	32.4112
Log return	291	0.2474	0.2154	0.7870	-2.3223	3.5089
Non-exited projects						
Total investment	169	91,869	26,000	194,588	5,000	1,692,000
PV total investment	169	91,493	26,000	192,289	5,000	1,677,931
Years held so far	169	2.14	1	2.34	0	13
Characteristics of investments						
Diligence (hours)	232	65.18	15	346.03	0	5,000
Stage of project at investment: percent of all projects	275	100.00				
Seed	80	29.09				
Startup	124	45.09				
Early	51	18.55				
Late	11	4.0				
Turn around	9	3.27				

Interaction of investor with firm (percent)						
	Number of observations	Mean	Median	Standard Deviation	Minimum	Maximum
	242	100.0%				
Daily	33	11.83				
Weekly	31	11.11				
Monthly	44	15.77				
Quarterly	71	25.45				
Annually	17	6.09				
Rarely or non	46	16.49				
Characteristics of investors						
Years investing	295	11.31	8	10.03	1	49
Number of projects invested in	338	16.17	9	19.49	1	63
Number of projects exited from	362	7.08	2	13.19	0	42
Percent of wealth in angel investments	237	13.05	10	15.22	1	85
Percent with degrees						
Bachelor's	40	30.30				
JD	10	7.58				
Master's	67	50.76				
Ph.D.	15	11.36				

Total investment	Total dollar amount invested in a project.
PV total investment	Total dollar amount invested in a project discounted at the riskless rate.
Years Held	Years between the angel investor's initial investment in the venture and exit from the venture.
Years held so far	Years held for exited projects; the number of years since initial investment in the venture and the survey date for nonexited.
Total cash out	Total dollar amount paid to an investor by a project.
PV total cash out	Total dollar amount paid to an investor in a project discounted at the riskless rate.
Base multiple	Total cash out divided by Total investment.
PV base multiple	PV total cash out divided by PV total investment.
Net return	$(\text{PV Base multiple})^{(1/\text{YearsHeld})}-1$.
Log return	ln(PV base multiple) /YearsHeld.
Years Invested	Years that the angel investor has been investing.
Years Worked Large	Angel investor's years of work experience in firms with more than 500 employees.
Total Investments	Total number of angel investments that the respondent has made.
Total Exits	Total number of angel investment exits that the respondent has experienced.
Diligence	Number of hours of due diligence the angel investor group conducted prior to making the investment.

This table presents summary statistics for our data for all projects, both exited and non-exited. It presents summary data on the stage at which investments are made, interaction of the investor with the firm, and selected characteristics of the investors. There are 588 observations, of which 419 are exited. Of those 419 observations, 11 observations have durations listed in the data as 0 years, for which net returns and log returns cannot be calculated. In addition, 117 projects have durations of a year or more but have payoffs of zero. For these observations, the log return cannot be calculated.

Table 2
Average Net Returns by Type of Exit

Exit Type	Number	Average Net Return
Ceased operating	114	-0.930
Bought by another operating firm	180	0.8412
Bought by investors	20	0.4467
Initial public offering	56	0.9434
Other	9	0.0302
Missing	30	0.4445
Total	408	

Net return (PV Base multiple)$^{(1/YearsHeld)}$-1.

Seventeen projects with durations less than a year are excluded from the averages. This table shows the net return rather than the log return in order to include firms with payoffs of zero. The category "Missing" are firms with computable payoffs and the type of exit is missing in the data.

Table 3
Returns by Industry

Average Net Return

Industry	Number of Observations	Mean	Median	Standard Deviation	Minimum	Maximum
All	408	0.294	-0.018	2.299	-1.000	32.411
Biotechnology	8	-0.067	-0.506	1.131	-1.000	1.941
Business products and services	25	1.808	-0.080	6.945	-1.000	32.411
Communications	4	-0.326	-0.604	0.841	-1.000	0.902
Consumer products and services	21	-0.732	-1.000	0.497	-1.000	0.306
Electronics hardware	17	-0.107	-0.003	0.847	-1.000	1.845
Financial services	5	0.030	0.094	0.661	-1.000	0.830
Healthcare services	8	-0.152	-0.019	0.804	-1.000	1.242
IT services	7	-0.452	-1.000	0.744	-1.000	0.862
Internet	12	0.749	-1.000	0.591	-1.000	0.694
Media and Entertainment	17	0.069	-0.096	1.495	-1.000	5.609
Medical devices and equipment	9	-0.333	0.102	0.640	-1.000	0.265
Software	45	-0.117	-0.820	1.499	-1.000	7.575
Other	21	0.273	0.266	0.813	-1.000	2.824
Missing	209	0.518	0.168	1.803	-1.000	18.390

Log Return including only observations with nonzero payoff

Industry	Number of Observations	Mean	Median	Standard Deviation	Minimum	Maximum
All	291	0.247	0.215	0.787	-2.322	3.509
Biotechnology	4	0.549	0.565	0.451	-0.013	1.079
Business products and services	13	0.916	0.428	1.149	-0.083	3.509
Communications	3	-0.406	-0.840	0.913	-1.020	0.643
Consumer products and services	5	0.112	0.050	0.138	-0.035	0.267
Electronics hardware	12	0.080	0.082	0.600	-0.959	1.046
Financial services	4	0.224	0.149	0.268	-0.006	0.604
Healthcare services	5	0.255	0.224	0.339	-0.049	0.807
IT services	3	0.200	-0.0004	0.370	-0.023	0.620
Internet	2	0.400	0.400	0.180	0.273	0.527
Media and Entertainment	14	-0.042	-0.059	0.688	-1.023	1.889
Medical devices and equipment	5	0.181	0.217	0.064	0.098	0.235
Software	25	0.029	0.241	1.029	-2.213	2.149
Other	19	0.219	0.256	0.550	-1.542	1.341

The mean returns are proportional returns in the first part of the table and log returns in the second part. Hence, there are 408 observations in the first part of the table and 291 observations in the second part after excluding 117 projects with returns of zero. The returns are returns per project, not for project-years as are the returns in Table 5. The industries are our consolidation of the “industries” provided in the survey, which included most of these industries but include for example, “Local bank in Texas”.

Table 4

Standard Deviation of Log Returns by Year

Duration	Number of observations	Mean	Median	Standard deviation	Minimum	Maximum	Underlying annual standard deviation
1	58	0.929	1.338	1.156	-2.322	3.509	1.156
2	49	-0.145	-0.035	0.799	-1.542	1.889	1.129
3	49	0.340	0.484	0.497	-1.717	1.355	0.861
4	51	0.177	0.215	0.275	-0.596	0.842	0.550
5	33	-0.131	-0.285	0.319	-0.357	1.115	0.713
6	15	0.125	0.224	0.714	-2.213	1.045	1.749
7	13	-0.146	0.017	0.572	-2.017	0.326	1.513
8	9	0.198	0.159	0.195	-0.130	0.475	0.552
9	3	0.376	0.240	0.263	0.208	0.679	0.788
10	1	0.024	0.024				
11	4	0.252	0.288	0.110	0.098	0.334	0.364
13	1	0.090	0.090				
17	1	0.083	0.083				
25	1	-0.031	-0.031				
25	2	0.136	0.136	0.141	0.036	0.236	0.704
33	1	0.057	0.057				
Total	291						

All 291 projects that have positive exit values are included in this table. The standard deviation cannot be estimated in the usual manner when there is only one observation and there is no information in the minimum and maximum. The standard deviation multiplied by the square root of the duration provides an estimate by year of the underlying annual standard deviation of the returns.

Table 5

Estimates of Expected Gross Returns

	Number of project-years	Firms closed with zero payoff	Fraction of firms closed per project-years	Expected return conditional on positive payoff	Expected return
Projects that Last One Year or Longer	1,144	117	0.102	1.893	1.699
All Projects	1,147	122	0.106	1.918	1.714

The first estimate of the expected return uses all projects that last a year or longer. The second estimate includes projects coded as lasting zero years. We interpret zero years as a duration between zero and one year; we use one-half year as a point estimate of the projects' durations. The three extra project-years are six half-year projects, of which five fail. In the estimates, each project is weighted by the number of years' duration. As a result, the six projects that are coded as 0 years in the data and given a duration of 0.5 years in the second line of the table add three project-years. The expected return is $p\exp\left(\alpha+\left(1/2\right)\sigma^2\right)$ where p is the fraction of project-years that resulted in a project exiting with a positive return, α is the mean log return and σ is the standard deviation of the log return.

Figure 1

Potential cash flows from angel investments

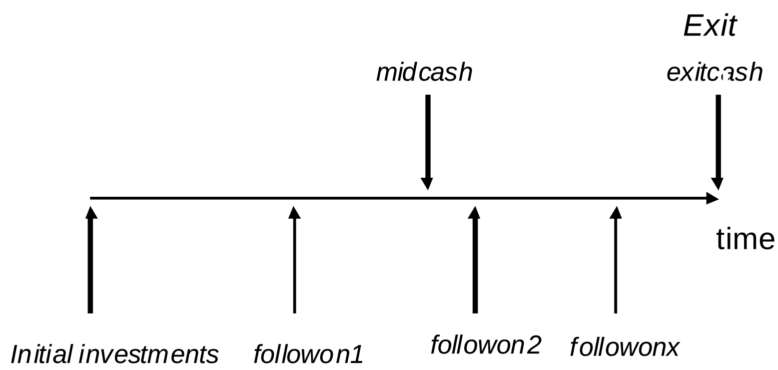

Flows above the horizontal time line represent cash flows from the project (the entrepreneur) to the angel investor. Flows below the line represent cash flows from the angel investor to the project (the entrepreneur). The 419 exited observations in the AIPP data set contain the dollar amount, year of the initial investment, exit year and exit cash, if any. Investments in the project made after the initial investment are called follow-on investments. The AIPP calls the first of these *follow1invest* and the second *follow2invest*. In both cases the data contain the dollar amount and date of these investments. The AIPP sums any subsequent investments and reports the total as *followxinvest*. For these the dollar amount is available but not the year of the investment(s). These three cash flows are reported 117, 24 and eight times, respectively, and all reported values are positive. For exited projects, these cash flows occur 106, 19 and seven times, and all reported values are positive. *Midcash* is the total dollar amount of any cash flows from the project to the angel investor before the exit date. The AIPP does not report the year of these cash flows. Midcash is reported 90 times, of which 78 are positive, and all are in the exited dataset.

Figure 2

Log Returns per Project by Duration

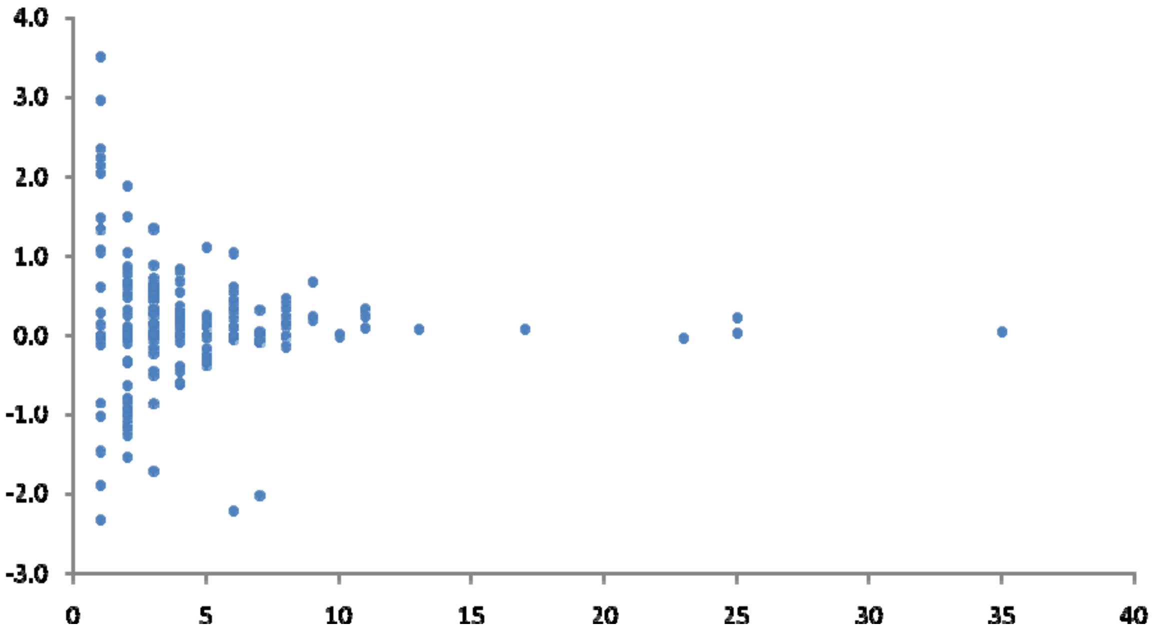

Duration

The log return for each project is on the vertical axis and the duration of the project on the horizontal axis. Dots sometimes overlap and therefore may represent more than one observation.

Appendix I: Are the Data Biased?

Researchers who study stock returns for publicly traded companies are fortunate in that returns data are plentiful. This is not so for angel investments. As is true for other private investments, transactions data are at best difficult to find and potentially unreliable if indeed they can be found. In this section we identify potential problems with the AIPP data, compare values for available returns and other characteristics with other data sources, and speculate on the effect of any biases that may be embedded in the data.

The first problem that users of survey data face is nonresponse bias. Not all people asked to complete a survey do so, and those who complete the survey may differ systematically from those who do not. Wiltbank and Boeker (2007a) report that the response rate of the AIPP is uncorrelated with the return measure, called the base multiple. This argues against the claim that the data are biased and that angels tend to report only good outcomes. Shane (2009), though, remains cautious. Referring to the AIPP, he says, "...we need to treat studies of the performance of angel investments with *extreme* caution." One reason for this is that the survey is of angel groups, not the universe of angel investors. Members of angel groups are not representative of angels in general; Shane reports that members of groups tend to be more successful than individual angels. In addition to selection bias, he notes that a sample of exited investments probably overweights established angels because younger groups tend to have fewer exits.

Malkiel and Saha (2005), who study hedge fund returns using the TASS database (1995-2003), face two additional problems related to selection bias. The first is backfill bias (sometimes called incubation bias) and the second is survivorship bias. The idea behind backfill bias is that selection bias is compounded when those who report the data backfill the data on these funds that have done well. This is because hedge funds that have survived tend to have had good results in the years before the recording period, too. In the case of TASS, there is a related bias: Some funds may have reported data to another service previously. When those funds start reporting to TASS they might not report all of the data that they gave to the previous service. Malkiel and Saha say that the difference between backfilled returns and contemporaneous returns is over 500 basis points, which is statistically significant. Neither of these additional sources of bias is likely to apply to the AIPP, though. The AIPP base multiples that we use are total returns on individual investments, not annual returns on funds. As such, they are not subject to backfill bias.

Survivorship bias arises because returns in the database for any period are those of surviving funds. Funds that fail do not report data. This biases reported returns up. Malkiel and Saha are able to obtain the previous returns from some defunct funds and find that the difference between surviving funds and defunct funds is over 830 basis points, which is statistically significant. The AIPP may well be subject to survivorship bias. If so, then the return multiples are probably too high. To check this, we next compare failure rates in the AIPP dataset with venture capital failure rates, as well as comparing AIPP return multiples and implied internal rates of return to those reported by other sources.

How do AIPP failure rates compare to failure rates on venture capital investments?

Our first bias check in the AIPP data set compares angel investment failure rates with those of venture capital. Angel investments are riskier than venture capital investments, so AIPP failure rates should be higher. Puri and Zarutskie (2009) use the Longitudinal Business Database, which tracks all U.S. employer business establishments, which they match to Venture Source and

VentureXpert. They define failure as leaving the database without experiencing a firm sale or IPO. Their data begin in 1981. After five years, the failure rate of VC-financed firms is 20 percent and after 10 years it is 28.4 percent. These data are right-censored, though. For firms that received venture capital financing before 1998, the figure is closer to 40 percent.[9] The annual failure rates at that point suggest that the total failure rate increases only slightly over the next few years. They also report results for a "one-to-one matched sample." These failure rates are higher, 31 percent after five years and 36.4 percent after 10 years. Again, the evidence suggests that the cumulative failure rate increases only slightly beyond ten years. These figures are much lower than AIPP failure rates, especially given that the mean AIPP project lasts less than five years. Complete failures (returning nothing) comprise 29.6 percent of the AIPP sample and 51.8 percent return less than the invested capital after adjusting for the time value of money (50.4 percent without the adjustment).

Anecdotal evidence from practitioners also supports the conjecture that AIPP failure rates are higher than on venture capital. Fred Wilson, managing partner of the venture capital firm Union Square Ventures, provides anecdotal evidence of even lower failure rates than Puri and Zarutskie's (2009) report for venture capital. Wilson (2007) defines failure as a project that returns less than its required investment and refers to an informal "1/3" rule: "1/3 of the investments will fail, 1/3 will under-perform expectations, and 1/3 will meet expectations." He also states that 20 percent failures, 45 percent under-performing, and 35 percent successes is closer to what a "good early stage investor's track record" should be. These failure rates are far below those of angel investments in the AIPP data set.

Phalippou and Gottschalg (2009) use the Thomson Venture Economics cash flow data set

[9] 9 We thank Rebecca Zarutskie for this information

and the VentureXpert data set to study private equity (not venture capital). Private equity investors provide these data voluntarily. The VentureXpert data have a success rate of 45 percent. The cash flow data set has success rates of 50 percent. These cash flow figures are net of the S&P500 return and net of fees. Phalippou and Gottschalg conclude that the cash flow data set is biased high. The bias in failure rates, though, is probably only about five percentage points. These private equity failure rates are comparable to those in the AIPP data set. Phalippou and Gottschalg say that their private equity sample has poor performance relative to other assets of similar risk, though. If so, then the failure rates of these other asset classes are likely lower than our sample of (presumably riskier) angel investments.

In short, AIPP failures rates are higher than reported venture capital failure rates and much higher than anecdotal evidence of those rates. Private equity failure rates are roughly comparable. Based on failure rates, we conclude that there is little or no evidence of bias in the AIPP sample of angel investments.

What do others say are the returns on angel investments?

Ideally, we could compare our estimates of expected returns to what other researchers have reported. Unfortunately, this is not possible. To our knowledge, our paper is the first to provide estimates of expected returns derived from reported transactions prices. The few available estimates of expected returns come from surveys, and most focus on investment multiples or internal rates of return. Still, we can compare some of those values to those from our data. If the AIPP data are comparable to what others report, then we can probably conclude that any bias is small – or at least, no worse than any bias in the other data sources.

We begin with base multiples and IRRs. The equally weighted average of the base multiples in our sample of 419 exited projects is 11.95. The average investment period is 3.62 years. This implies an IRR of 98.4 percent for the exited projects -- again, equally weighted. How does this compare to what others have reported? Shane (2009, p. 193) says that, "Given the failure rate of angel investments, successful angels target a thirty times multiple on their invested capital in five years." A multiple of 30 over a five-year period implies an IRR of 97.4 percent, which is remarkably close to the 98.4 percent IRR in our data. This does beg the questions of whether or not the survey respondents are answering a question that makes the two figures comparable, and whether the AIPP respondents are “successful angels” or whether they are more or less “successful.” There is no way to know. Another question is whether respondents who said they expect an implied IRR of 97.4 percent meant that they expect individual investments to earn an average of 97.4 percent, or whether they meant that they expect the total return on all of their investments to be 97.4 percent. Put differently, did they report an equally weighted average or did they report a value-weighted average? The close correspondence between the AIPP’s implied IRR of 98.4 percent and the “successful angels’” implied IRR of 97.4 percent suggests that they reported an equally weighted average.

Still, computing the comparable value-weighted average from the AIPP data is worthwhile. The value-weighted base multiple of the full sample of 419 exited investments is 2.90, which is a bit less than a quarter of the equally weighted average. The value-weighted average for the years the investments were held is 3.67 (obtained by computing the total amount invested in each project, dividing by the total invested in all projects, multiplying each quotient by the years held, then summing the 419 results). The implied IRR in this case is only 33.7 percent, which is probably not too different from the expected returns on other risky equity portfolios during the sample period. This is especially true given that these 419 investments have all exited and probably will have done better than the nonexited investments in the AIPP, which would reduce the implied IRR in the full sample. Based on these calculations derived from base multiples, we conclude that if the AIPP data are biased at all, then the bias is likely to be small.

What other numbers in the AIPP dataset can we compare?

Preston (2007, page 92) gives IRRs and multiples for private equity. We reproduce Table 5.1 below:

Table 5.1, "Rates of Return for Private Equity Investments."

	IRR (percent)	Return on Investment	Implied Holding Period
Seed	60+/year	10x	4.9 yrs
Start-up	50	8x	5.1
Early Stage	40	5x	4.8
Second Stage	30	4x	5.3
Near Exit	25/year	3x	4.9

Preston does not provide holding periods, so we have computed the implied holding periods from the IRRs and multiples. Caveats are in order. First, Preston's table is for private equity. This would include investments by individual angel investors and venture capitalists as well as angel groups such as those in the AIPP. Including individual angel investors would tend to reduce the average returns and venture capitalists, who tend to invest in later-stage companies, would probably do likewise. Second, we do not know whether the figures are equally weighted or value-weighted. Given these caveats, how do these data compare to those from the AIPP?

First, we can see that the average holding periods are all longer than the AIPP data (3.62 years in Table 1). However, the AIPP values are not necessarily comparable to Preston's numbers. All of the holding periods used in the calculation of the AIPP holding period are for exited investments. Because most investments in the AIPP data are relatively recent, the nonexited investments in the sample probably will run longer. If Preston's figures refer to investments over a period long enough to permit a stable proportion of investments to have exited, then the holding periods would indeed be higher than the AIPP figure. Our take is that Preston's reported holding periods are at least roughly comparable to those in the AIPP. Second, the AIPP's weighted-average IRR is about 33.7 percent, which corresponds to Preston's figures for Early Stage or Second Stage investments. But Preston's figures are for private equity. Early Stage or perhaps Start-up investments are a better match for AIPP data because successful angel groups tend to focus on start-ups and early-stage companies. The AIPP's equally weighted average of 98.4 percent would correspond to Seed funding, and although this is half-again as large as the 60 percent minimum that Preston reports, it is probably not too far out of line given that she reports a lower bound and may include investments by those who tend to earn less than angel groups.

Reynolds (2007) reports the results of a survey of informal investors, which he defines as friends, family members and individuals in the personal social networks of the entrepreneurs. These informal investors report expected return multiples that vary by the size of the investment. His Table 8.13 gives the percentage of survey respondents who claim to expect multiples within selected ranges. We can gauge the mean response by multiplying the midpoint of the range by the percentage of respondents who select that range. For investments comparable to the median AIPP investment, informal investors expect a multiple of 6.94 and for the mean the figure is 6.78. These figures are subject to limitations. First, we do not know the expected duration of the investments. Second, the figures might be biased low because the highest range a survey respondent could choose was bounded at 20. Third, there is no way of knowing whether the survey respondents provided an equally weighted expectation or a value-weighted expectation.

Given these caveats, we can say that these multiples are a little lower than the equally weighted averages in the AIPP data and much higher than the value-weighted average. We conclude that Reynolds' results do not suggest that the AIPP data are biased, or at least that they are not biased relative to his. This is especially persuasive given that at least most of Reynolds' informal investors are not members of groups and we know that such investors tend to earn less than angel groups.

Additional information about expected returns can be gleaned from the information that angel groups provide to prospective entrepreneurs. For example, the website of Maine Angels, an angel group in Portland, ME (visited December 11, 2008), says that it typically funds between $100,000 and $2 million, and that, "Candidate companies should also have a high potential for growth and profitability, can provide at least 35% of annual return of investment within five to seven years and have a strategically planned and viable exit." All of the AIPP data are from a period before 2008, so to the extent that expectations shifted between the time of the investments in the AIPP projects and the end of 2008, these data may not be comparable. Still, the mean investment in the AIPP data is $169,089 (median $50,000), and 80 percent of the investments range from about $15,000 to about $300,000. Given that these figures are in nominal dollars, the AIPP data, being earlier, should be smaller, though not by enough to make these investment sizes comparable. A better explanation for the size disparity could be that the Maine Angels' website may refer to the total investment by the entire group. The AIPP data are investments made by individual investors (though the individuals are members of a group). In short, the (individual) AIPP investments are smaller than the average Maine Angel group investment, and the value-weighted AIPP investment tends to return a bit less than the lower bound of Maine Angel's expected amount (33.7 percent vs. 35 percent). Based on this, there is no evidence that the returns on the AIPP data are biased high.

We can also compare the AIPP data on IPOs to those reported by other data sources. IPOs are the gold standard for angel investments, providing the extraordinarily high returns that catch the public's eye. Band of Angels is perhaps the most famous angel group in the United States. According to Band of Angels' website (visited March 1, 2009), Band of Angels has had nine IPOs out of 209 investments (both exited and nonexited). This is only 4.3 percent. In contrast, the AIPP data have 56 IPOs out of 588 investments (419 exited plus 169 nonexited), or 9.5 percent. On the one hand, this is twice as high, but on the other hand, it is at least within an order of magnitude.

In addition, the comparison is not quite as apt as it seems. Band of Angels' figures are for the entire group, while the AIPP data are for individual members within groups. For example, suppose that one member of each of 10 different AIPP groups invests in each of 10 projects and one of these has a successful IPO. The IPO rate for the *groups* is 10 percent. However, if two members of one group invest in the successful IPO and only one in the other investments, then the group success rate is still 10 percent (1/10) while the AIPP success rate almost doubles to 18 percent (2/11). Band of Angels, which reports figures for the entire group, would report a 10 percent success rate regardless of the number of individual angels who invest in each project.

This proves to be important because many angels in the AIPP data did invest in the same IPO. In fact, 13 of the 56 IPO investments are the same deal. All 13 observations are from members in the same angel group who invested in the same company in the same initial year and exited in the same exit year. All earned the same base multiple. Thirteen members of a group saw a deal, bought in at the same time and later reported the deal to AIPP. If we suppose that this single deal had exited some other way then the IPO percentage drops from 9.5 percent to just 7.3 percent. If we treat the 13 observations as a single successful IPO (as Band of Angels would do), then the rate is 7.5 percent. These figures are still higher than Band of Angels IPO rate of 4.3%, but they are considerably closer.

Moreover, the fraction of buyouts – probably the second most profitable type of angel exit – is almost exactly the same. Band of Angels has had 45 "profitable acquisitions," or 21.5 percent of its investments exit via buyout while the corresponding AIPP figure is 22.4 percent (132/588; 73 additional sales resulted in a loss). This comparison is also not quite as apt as it may seem, for the same reasons as for IPOs.

What does this tell us? It seems surprising that the investors in the AIPP would have a higher figure than Band of Angels, and this stands as the strongest evidence that the returns in the AIPP data are biased high. The comparable fraction of exits through buyouts, though, casts some doubt on the conclusion that the returns are biased high.

Goldfarb, Hoberg, Kirsch and Triantis (2009, hereafter, Goldfarb et al.) use the records of a failed law firm (Brobeck, Phleger & Harrison; hereafter, Brobeck). Unfortunately, their results are not directly comparable to the AIPP because Goldfarb et al. include angel groups with venture capital firms. Goldfarb et al. say that they conduct robustness checks and their main results are not sensitive to the way that they classify investments as venture capital, angel, founder or family. Still, this means that we cannot directly compare specific figures from Goldfarb et al.'s angel data with the AIPP data because the AIPP is from groups and Goldfarb et al. bury their groups with venture deals. But neither can we compare the AIPP data with Goldfarb et al.'s venture capital results because those figures are dominated by venture capital deals. We can argue that Goldfarb et al.'s angels, who are not members of groups, should be less successful than the investors in the AIPP (Shane, 2009). They probably conduct less due diligence and have lower returns than the investments in the AIPP data. This is especially true if some of Brobeck's deals are loans rather than equity, because the AIPP data are all equity deals. Finally, although Goldfarb et al. say that Brobeck entered deals of high quality, it took positions in some of these deals and in the end, Brobeck did fail. Perhaps the deals were not as high quality as it believed.

What do Goldfarb et al.'s data say? In addition to the classification issues noted above, comparisons are problematic because Goldfarb et al. have only 32 angel investments. Here are the outcomes for angel-only events from Goldfarb et al. and the corresponding figures from the AIPP:

Investment status	*Goldfarb et al.*	*AIPP*
IPO	3.1% (1/32)	9.5% (56/588)
Acquisition	18.8% (6/32)	31.5% (185/588)*
Failure	28.1% (9/32)	25.3% (149/588)**
Non-exited	50.0% (of which 25% are active and 25% are inactive and very possibly failures) (16/32)	28.7% (169/588)
		* Including 20 cases that were "Bought by Investors" would imply 34.9%.
		** Including 73 exits at a loss would imply 37.8%.

The AIPP numbers do not sum to 100 percent because the AIPP data have categories such as "Other." We include investments with missing exit information as being failures. These numbers imply the following percentages for exited investments:

Investment status	Goldfarb et al.	AIPP
IPO	6.2%	13.4%
Acquisition	37.6%	44.1%*
Failure	56.2%	35.5%**

* 48.9% including "Bought by Investors."
** 53.0 including 73 exits at a loss.

Because all of the Brobeck et al. deals are from 1993-2002 we compute similar numbers from 1993-2002 below.

	Goldfarb et al.	AIPP
IPO	3.1%	12.3% (43/351)
Acquisition	18.8%	41.9% (147/351)*
Failure	28.1%	33.0% (116/351)**
Non-exited	50.0% (of which 25% are active and 25% are inactive and very possibly failures) (16/32)	8.0% (28/351) * Including 10 cases that were "Bought by Investors" would imply 44.7%. ** Including 58 exits at a loss would imply 49.6%.

As expected, the 32 angel investments from Brobeck et al. had a substantially lower percentage of IPOs and acquisitions. However, they also had a lower percentage of failures. The Brobeck data have much a lower percentage of exited projects. One interpretation of this comparison is that the proportion of IPOs and acquisitions in the AIPP dataset are too high. Another interpretation is that Goldfarb's sample comprises only 32 deals and that the differences are simply due to chance. The 3.1% IPO rate represents, after all, just one deal. Still other interpretations are that the inferior results trace to Brobeck's angels not being members of angel groups, or possibly that they took debt instead of equity, or the deals were bad and contributed to the firm's failure. If any of these is true, then the superior performance of the AIPP investors is unsurprising.

We can also check whether the AIPP data are comparable to other data along dimensions other than returns. For example, Goldfarb et al. (2009) says that about 18.8 percent of angels in angel-only deals had previously invested in the same company while Mason and Harrison (1996) report 25 percent. The percentage of investments with follow-on funding in the AIPP, which is

19.9 percent (117/588 deals), is comparable to these two figures. In addition, Cochran (2005) reports that venture capital investments that receive follow-on funding tend to have done well. If the AIPP data are biased, then we would expect them to have a higher proportion of follow-on investments. However, they do not. Unless the deals in Goldfarb et al. or Mason and Harrison are a biased standard, this stands as evidence that the AIPP data are relatively free of bias.

The number of coinvestors provides another avenue by which to gain insight. Investments made with a coinvestor are probably less prone to bias because the survey respondent is more likely to be caught reporting biased results. There is, after all, at least one other person who knows about the project (even if he invests different amounts or at different times). In our data this comparison is problematic because there are many missing values for coinvestors. Still, it is worth doing. We find that *t*-tests show that the base multiple is significantly different between investments with coinvestors and those without (t-ratio = 2.76), and the point estimate suggests that base multiples are higher for deals without coinvestors. This is indeed what we would expect if the data are biased. For annual multiples, though, the t-ratio falls to 1.40. Regression analysis tells the same story. When we regress the base multiple on the number of coinvestors, we obtain a significantly negative coefficient (-0.82, *t*-ratio = -1.98). This is consistent with selection bias in the data, which is being mitigated on deals with many coinvestors because the survey respondents are more likely to be caught reporting false data. Yet this evidence of bias also vanishes with annual multiples; the estimated coefficient of -0.70 is not close to being statistically significant (*t*-ratio = -0.94). This suggests that the regression is simply measuring the relation between the number of coinvestors and the duration of projects rather than the relation between the number of coinvestors and the total return multiple. Deals with more coinvestors do, in fact, tend to be shorter. The correlation coefficient between the number of coinvestors and project duration is -0.24 with a *p*-value of 0.004.

We also have deal size. It seems reasonable that larger deals are more likely to be remembered correctly. This alone would not eliminate intentional bias, but it would tend to eliminate unconscious bias. People may still lie, but they are less likely to remember only the good deals if the deals are large. Goldfarb, Hoberg, Kirsch and Triantis (2009) suggest that more sophisticated investors are more likely to do larger deals, too. Arguably, these sophisticated investors are less biased, or perhaps more likely to fear posting inflated results on the AIPP surveys but smaller figures on, say, their tax returns. To check this, we compare multiples on investments that are larger than the median with multiples on investments that are smaller than the median. *T*-tests for both base multiples and annual multiples find no statistical significance. The results show no evidence of bias.

Summary

What then, does the evidence say regarding bias in the AIPP data? Based on the high percentage of investments that resulted in an IPO, we conclude that the data may contain some biases that inflate the reported performance of the angel groups that participated in the AIPP. The IPO comparison is not particularly apt, though, and the disparity in IPO rates in the AIPP dataset is substantially increased by one deal that attracted 13 individual angel investments. In addition, the evidence of bias is hardly overwhelming. First, the fraction of profitable sales is almost exactly the same as Band of Angels. Second, failure rates are substantially higher than the Brobeck data in Goldfarb et al. (2009), tilting the evidence of bias the other way. Third, implied internal rates of return and return multiples -- which are, after all, the fundamental input to our estimation procedure -- are comparable to figures from other sources, especially on a value-weighted basis. Finally, the data survive several checks designed to detect bias. The percentage of follow-on investments in the AIPP data, which are correlated with higher returns at least on venture deals, is comparable to figures that other researchers report. Multiples from deals that are larger than the median do not differ statistically from those that are smaller than the median, and although some tests using the number of coinvestors support the claim that base multiples are biased high, there is no evidence of bias in annual multiples. The reason that base multiples without coinvestors are higher than deals with coinvestors is that deals with coinvestors tend to be shorter. The correlation coefficient between the number of coinvestors and project duration is -0.24, with a *p*-value of 0.004.

Even if the AIPP data are free of bias, this does not mean that they are representative of angel investments in general. This is because the angel groups which participated in the study are not a random sample of all angel investors. Aside from the advantages that members of groups enjoy, the types of investments differ. For example, Shane (2009) reports that many angel investments are debt, whereas all of the deals in the AIPP are equity. Angel investors who are members of groups tend to fare much better than those who are not.

Appendix II: Data

Calculations of present values of the first two follow-on investments:

Because the AIPP data contain only the year of a cash flow, we assume that the cash flow occurs at midyear and therefore treat a year as covering the period from July 1 - June 30. We use end-of-month quotes from CRSP. This means that we use monthly 30-day (approximately) bill rates from the previous months (June 30 - May 31). For example, if the initial investment is made in year 2000 and the project exits in year 2001 then we use Treasury bill rates from June 30, 2000 - May 31, 2001. We adjust these rates for the actual number of days in the month and sum the monthly rates over the period from the initial investment through the year of the follow-on investment (or through the project's exit) to compute the continuously compounded interest factor for the entire period.

If the initial investment year is the same as the exit year then we use rates from the middle of the year (March 31 - August 31), covering the period from April 1 - September 30. This also applies for later cash flows. For example, if the initial investment year is 2000 and the follow-on investment is also year 2000 then we use rates from March 31 - August 31.

Calculation of present values of follow-on investments after the first two:

The AIPP does not report the year of follow-on investments after the first two. Instead, the AIPP sums all remaining follow-on investments and reports this as a single number. The year is not reported. This occurs eight times in the 588 observations and seven times in the 419 exited investments. We assume that these follow-on investments occur midway between the second follow-on investment and the exit year. In one of these eight cases the exit year is not reported and in another the project has not yet exited. To handle these two cases we first calculate the average time between the second follow-on investment and the exit. This is about 2.5 years. We therefore add 2.5 years to the second follow-on investment in these two cases to obtain an estimate of the year of these follow-on investments.

Calculation of present values of midcash payments:

The AIPP does not report the year of midcash payments. We assume that midcash payments occur halfway between the beginning and end of the investment.