PRELIMINARY DRAFT
COMMENTS WELCOME

Learning by Investing: Evidence from Venture Capital

Morten Sorensen[†]

April, 2006

Abstract: Venture capital investors (VCs) create value by actively exploring new investment opportunities to learn about their returns. A free rider problem reduces exploration in a fully informed market, but the organizational structure of VCs permits them to internalize the benefit of exploration. I present a basic model of learning, based on the statistical Multi-Armed Bandit model. The value of investing in an industry consists of both the immediate return and the option value of learning about the industry. The model is estimated, and it is found to explain VCs' observed investment behavior well. VCs that explore more have higher return.

Venture capitalists (VCs) invest in privately held entrepreneurial companies, and are actively involved in the monitoring and management of these.[1] The literature has identified a number of ways VCs add value. VCs may screen bad companies, and their direct involvement with portfolio companies is valuable in several ways (Lerner (1995)): they help bring products to market faster (Hellmann and Puri (2000)); they replace inefficient managers (Hellmann and Puri (2002)); through their syndication partners, they

---

[†] University of Chicago, Graduate School of Business. Email: sorensen@uchicago.edu. I would like to thank Jean-Pierre Dube, Gunter Hitsch, and Francisco Perez-Gonzalez for helpful suggestions and discussions. I am grateful to participants at the Chicago GSB Faculty Lunch for helpful comments and to Susan Woodward of Sand Hill Econometrics for generous access to data. The usual disclaimer applies.

[1] See Gompers and Lerner (1999), Gorman and Sahlman (1989), and Sahlman (1990) for descriptions of the VC industry and the role of VC investors in general.

provide access to a network of customers and suppliers (see Hochberg, Ljungqvist and Lu (2005)); and the reputation of an established VC help certify the value of a young entrepreneurial company (Megginson and Weiss (1991)).

I identify an additional source of value generated by VCs. I argue that the organizational structure of VC investors permits them to internalize the value of learning about new technologies, which is socially valuable. When investing in uncertain industries and technologies, the return to an investment consists of both the immediate return and the option value of learning. The option value is indirect value that arises when the acquired knowledge help improve future investment decisions. An investment in an unknown technology may have low expected immediate return, but a large option value, since a success would spur additional investments. An investment in a well-known technology has low option value, since it's unlikely to alter beliefs about returns and affect future investment decisions. Here, uncertainty is distinct from traditional risk or volatility. It represents uncertainty about the underlying distribution of returns. A new and uncertain technology may turn out to be less risky than a known technology, but investors are initially uncertain about this, and the potential to learn about the distribution of returns generates the option value.

The option value of learning is social value, and a social planner facing the same uncertainty about the distribution of returns would internalize this value. However, in a fully informed decentralized market, a free-rider problem makes it difficult for individual investors to fully internalize this value. With informational spillovers, investors would prefer investments that maximize immediate returns, and leave others to experiment with new uncertain investment opportunities (see Bolton and Harris (1999) and Keller, Rady and Cripps (2005) for theoretical analyses of this problem).

The option value of learning is greater for investing in entrepreneurial companies. The main investors in these companies are venture capitalists, and their desire to internalize the option value is consistent with a number of pervasive organizational features: VCs are structured as long-term limited partnerships and are compensated for long-term performance; they repeatedly invest in projects with great uncertainty; they are in close contact with entrepreneurs, facilitating learning; and they are private investors in privately held companies, reducing informational spillovers.[2] These organizational features satisfy three important conditions promoting internalizing option value: they promote learning about technologies and investments in an uncertain environment; they promote forward-looking investment behavior, allowing investors to internalize gains of future investments; and they reduce informational spillovers, limiting the free-rider problem.[3]

The analysis proceeds by presenting a simple formal model of the option value of learning. This model is based on the statistical Multi-Armed Bandit model. In this model, the returns from investing in an industry are two-fold, and consist of both the immediate return and the option value of learning about the future potential returns in the industry. The model is calibrated and estimated using a dataset with US venture capital investments, and it is found that option values vary between .5% and 10% of the value of the investments. The model appears to describe observed investment behavior reasonably well. VCs are found to internalize the option value of learning to a significant extent. The investors can be classified according to the weight they place on this value, as well as the weight they place on general investment trends, and how "opportunistic" their investment strategies are. The empirical results suggest that VCs who place more weight on option value have better performance; that VCs who follow general trends have lower returns, although the evidence here is more mixed; and finally that VCs who make more opportunistic investment decisions have lower returns.

The Multi-Armed Bandit model addresses a stochastic control problem and is originally described by Robbins (1952). It presents the simplest and most tractable dynamic model of forward looking learning. The name refers to a situation where a gambler faces a

[2] Recently, the Freedom of Information Act (FOIA) has led to the disclosure of detailed information about investments by certain VC firms with public investors. This has led to a vigorous response from the VC industry and may limit future investments by public investors in general.

[3] Manso (2006) analyses a two-player Principal-Agent model where the Principal provides incentives for the Agent for internalizing this value. Bergemann and Hege (1998) consider a dynamic contracting problem in a similar context.

number of slot machines (one-armed bandits). The gambler is uncertain about the distribution of payoff from each of these, but learns about this distribution from repeated play, and his problem is to determine the optimal gambling strategy. The fundamental trade-off is between *exploiting* machines with a known payoff and *exploring* machines with an uncertain payoff, hoping that they will turn out to be more profitable. In the words of Berry and Fristedt (1985) (p. 5), "it may be wise to sacrifice some potential early payoff for the prospect of gaining information that will allow for more informed choices later." Whittle (1982) (p. 210) states that "[the Bandit problem] embodies in essential form a conflict evident in all human action. The "information versus immediate payoff" question makes the general problem difficult."

In economics the Bandit model has been used to describe firms' experimentation with prices to learn about uncertain demand (Rothschild (1974)). It has been applied in labor economics to model employee learning about jobs and job turnover (Jovanovic (1979) and Miller (1984)). Weitzman (1979) considers a related model of research projects and derive the optimal sequencing of these. In venture capital, Bergemann and Hege (1998) and Bergemann and Hege (2005) present theoretical models of staged financing based on the Bandit model. They address the question of how and for how long a VC should finance an entrepreneur given the potential for learning about the quality of the entrepreneurs' project. Bergemann and Valimaki (2006) briefly survey the Bandit literature.[4]

Methodologically, venture capital presents an attractive venue for studying learning and experimentation by firms. Unlike most other settings, data about VC firms contain information about their individual projects (the individual companies receiving financing) and the outcome of each of these. The projects are clearly delineated, and the whole learning and experimentation process is observed. This stands in contrasts to the job turnover applications, where only job changes are observed. As pointed out by Heckman and Borjas (1980), turnover and individual heterogeneity are not separately identified in

[4] Additionally, the Bandit model has found applications in medical experiment design and adaptive network routing.

this setting, but this identification problem is overcome when the full learning history is observed. As a methodological contribution, to my knowledge, this is the first paper to estimate a Bandit model from the agents' full learning histories, and it is the first time the option value of learning is empirically quantified.

To keep the analysis tractable, the Bandit model used here is the simplest model of forward looking learning. As such, it presents a somewhat simplified view of venture capital investments. In particular, the model assumes a stationary environment where investors just learn from their own past investment history (see Axelson, Stromberg and Weisbach (2006) for a first theoretical investigation incorporating the cyclicality of the industry), and the model assumes that projects (entrepreneurial companies) arrive exogenously. Numerical tractability also imposes limitations on the estimation procedure. Option values depend on investors' discount rate and their prior beliefs over returns. However, the option values are numerically difficult to calculate, and in the empirical analysis the discount rate and the prior beliefs are calibrated rather than estimated. More generally, most empirical analysis of entrepreneurial finance suffers from a lack of systematic data. Entrepreneurial companies, by definition, have short operating and financial histories and little information is systematically observed about them. This means that classifications of individual investments and their outcomes are necessarily crude measures, although these are standard in the literature. The paper presents the first investigation of this problem, and it should be read with the above limitations in mind. However, the results are encouraging, and hopefully future extensions of the methodology proposed here can address remaining shortcomings.

I. The Multi-Armed Bandit Model

Following Berry and Fristedt (1985) (ch. 4 and 6), the model presented below is a Bandit model with an infinite horizon, geometric discounting, and independent Bernoulli arms. The original derivation of the Gittins index in the general case is by Gittins and Jones

(1974). The formulation of the Bellman equation is from Whittle (1980), and the case with beta distributed prior beliefs is taken from Gittins and Jones (1979).

In the model a single VC invests in one project in each period.[5] The opportunity cost of investing in an unknown industry is that it prevents an investment in a known industry. This is consistent with anecdotal evidence that VCs' scarce resource is time and not money.[6] The environment is stationary. Only the investor's belief about the return from different industries evolves over time. The outcome of an investment is immediately observed, and the VC's beliefs about returns from investing in different industries are updated using Bayes rule.

A. The Formal Model

The investor faces an infinite sequence of periods, t = 0, 1, … At each t, the investor must choose between K alternatives (here, investments in different industries). This choice is denoted $i(t) \in \{1,...,K\}$. An investment in industry i at time t is either successful or not, as represented by the random variable $y_i(t) \in \{0,1\}$. The investment is successful with probability p_i. The investor is uncertain about this probability, but has prior beliefs $F_i(0)$ about p_i. These beliefs are then updated from realized outcomes of past investments. The updated beliefs before investing at time t are denoted $F_i(t)$, and they are functions of $F_i(0)$ and the history of investment decisions and outcomes up to time t.

The investor's strategy specifies the investment decision in each period as a function of the full history up to this point. Formally, let the strategy at period t be given by the function $s(t):\{1,...,K\}^t \times R^t \rightarrow \{1,...,K\}$, and the full investment strategy is then denoted $S = \left\{s(0), s(1), s(2), ...\right\}$. The investor's problem is to choose the investment strategy that

[5] Whittle (1981) studies a model where projects have Poisson arrival times and finds that an index policy remains optimal.
[6] In their study of VC investment analyses, Kaplan and Strömberg (2004) find that time requirements is a common concern for VC investors when evaluating potential investments.

maximizes total expected return. Let δ be the discount rate, and the investor's problem is then as follows.

$$V = \sup_S E\left[\sum_{t=0}^{\infty} \delta^t y_{s(t)}(t) \middle| F(0) \right] \tag{1}$$

This is a dynamic optimization problem. From a dynamic programming perspective, the investor's belief is the state variable, and the Bellman equation for the problem is the following equation.

$$V\left(F(t)\right) = \max_{i=1,2,\ldots,K} E[y_i(t) \,|\, F_i(t)] + \delta E[V(F(t+1))], \tag{2}$$

where the investors' beliefs develop according to the transition rule

$$F_i(t+1) = F_i(t) \; for \;\; i_t \neq i \tag{3}$$

$$F_i(t+1)(s) = \begin{cases} \dfrac{\int_0^s u f_t^i(u) du}{\int_0^1 u f_t^i(u) du} & for \;\; i_t = i \;\; and \;\; y_t = 1 \\[2ex] \dfrac{\int_0^s (1-u) f_t^i(u) du}{\int_0^1 (1-u) f_t^i(u) du} & for \;\; i_t = i \;\; and \;\; y_t = 0 \end{cases} \tag{4}$$

The transition rule in equation (3) states that the beliefs are only updated for the industry the investor invests in. For the remaining industries, they remain unchanged. Equation (4) reflects Bayesian updating of the investor's beliefs about the distribution of p_i after choosing $i(t) = i'$ and observing either $y_{i'}(t) = 0$ or $y_{i'}(t) = 1$.

The Bellman equation illustrates the basic trade-off faced by the investor. The first term in equation (2) is the expected immediate return. The second term is the continuation value conditional on learning about a given industry. This contains the option value of updating beliefs about the returns from different industries, allowing the investor to make more informed future investment decisions.

B. The Gittins Index

Gittins and Jones (1974) characterize the solution to the above dynamic programming problem in terms of the, now named, *Gittins index*.[7] For each industry, the Gittins index is calculated from the history of investments in this industry, and the formula for the index is

$$v(F_i(t)) = \sup_{\tau} \left\{ \frac{E\left[\sum_{s=t}^{\tau} \delta^s y_i(t)\right]}{E\left[\sum_{s=t}^{\tau} \delta^s\right]} \right\} \tag{5}$$

where τ is a stopping time for investing in just industry *i*. Their Gittins Index theorem states that at each period the optimal investment strategy is to invest in the industry with the highest value of this index. Formally, this is stated as

$$s(t) = \underset{i=1,\ldots,K}{\arg\max}\, v(F_i(t)) \tag{6}$$

Unfortunately, calculating the Gittins index is numerically difficult. There is no known case with a closed form solution, and evaluating the index requires a numerically intensive iteration procedure. One simple case of the Bandit model is analyzed by Gittins and Jones (1979). They also study the case where the outcome is either a success or a failure, i.e. $y_i(t) \in \{0,1\}$, and where each alternative has a constant probability p_i of

[7] Originally, Gittins and Jones (1974) named the index the *Dynamic Allocation Index*.

succeeding. The investors do not know p_i and their prior beliefs about this probability are distributed according to the *Beta*($a_{i,0}$, $b_{i,0}$) distribution.[8] While the assumption of a *Beta* distributed prior is not required by the general theory, it is essential for the calculations to be numerically tractable, and it is sufficiently general to allow for an arbitrary mean and variance of the prior distributions. From Bayes' theorem it follows that the posterior beliefs after observing a number of outcomes from the different alternatives are distributed according to the *Beta*($a_{i,0} + r_i$, $b_{i,0} + n_i - r_i$) distribution, where r_i is the number of successes with alternative *i*, and n_i is the total number of times this alternative has been tried. Let $a_i \equiv a_{i,0} + r_i$ and $b_i \equiv b_{i,0} + n_i - r_i$, and the *Beta*($a_i$, b_i) distribution can be interpreted as the posterior distribution of p_i after observing a total of a_i successes and b_i failures with this alternative, including the $a_{i,0}$ and $b_{i,0}$ from the prior beliefs. As a_i and b_i increase, the mass of this distribution becomes concentrated at $a_i/(a_i + b_i)$, which is the empirical frequency of success and the mean of the *Beta*(a_i, b_i) distribution. With *Beta* priors, Gittins and Jones (1979) derive an approximation of the Gittins index given by

$$v(a_i, b_i) = \lambda_i + \frac{1}{A(\lambda_i) + B(\lambda_i)(a_i + b_i)} \tag{7}$$

Here $\lambda_i = a_i/(a_i + b_i)$ is the investor's expected value of p_i (the mean of the *Beta*(a_i, b_i) distribution), and $A(\lambda_i)$ and $B(\lambda_i)$ are tabulated non-negative functions.[9] The total value of choosing alternative *i* is the value of the Gittins index $v(a_i, b_i)$, and the expression in equation (7) has an appealing economic interpretation. The immediate expected return from investing is λ_i. Here, with a binary outcome, the value of success is normalized to 1, and the investor's expected success probability is λ_i, which then equals the expected return. Clearly, the value of choosing alternative *i* is at least this big. The

---

[8] For the *Beta*(*a*,*b*) distribution, $f(s) = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)} s^{a-1}(1-s)^{b-1}$ for $a > 0$ and $b > 0$. Note that the mean of this distribution is $a/(a+b)$.

[9] Gittins and Jones (1979) calculate this approximation assuming $\delta = .75$, and find the resulting approximation to be accurate to at least four significant digits.

second term in equation (7) is the value of learning more about the distribution of returns from this alternative. This is the option value of learning. For a given λ_i, this value is decreasing in the number of trials of alternative *i*. As n_i increases, $a_i + b_i$ increases and the fraction in equation (7) goes to zero; the posterior beliefs become more certain and concentrated around λ_i, there is then less to learn about p_i for this alternative, and the option value vanishes.

II. Data Description

A. Sample

The data are provided by Sand Hill Econometrics. It contains the universe of VC investments in the US in the period 1987 to 2005.[10] The dataset expands on existing commercially available datasets provided by Thompson Financials (Venture Xpert)[11] and Venture One. These have been extensively used in the VC literature (e.g. Kaplan and Schoar (2005), Lerner (1995). Gompers and Lerner (1999) and Kaplan, Sensoy and Strömberg (2002) investigate the completeness of the Venture Xpert data and find that they contain most venture capital investments, and that the missing investments tend to be the less significant ones.

Since it takes on average 4-5 years for a VC backed company to go public or realize a return, the sample is restricted to investments made up to 2000, thus allowing investors 5 years realize the outcome of their investments. It's common for multiple VCs to invest in the same company, and the sample contains these different investments. When the same VC firm invests in a company over multiple rounds, only the initial investment is included in the sample. VC firms making less than 40 investments are removed from the sample, since their investment histories are too short to make reliable inference about

[10] Sand Hill Econometrics started out with the two existing commercially available databases, Venture Economics and Venture One and then spent substantial time and effort identifying missing investments, and verifying and correcting the information, particularly about companies' exit events.
[11] This database was formely known as Venture Economics and provided by SDC.

their investment behavior. The final sample used for estimation contains information about 17,551 investments in 5,621 individual companies by 210 VC firms.

B. Variables

The main variables in the data are the dates and outcomes of the investments, and the industry classifications of the companies. From these, the investment histories of the individual VC firms are constructed. Summary statistics of all variables are in Table 1, below.

Each company is classified as belonging to one of six industry groups. These are: “Health / Biotech,” “Communications / Media,” “Computer Hardware / Electronics,” “Consumer / Retail,” “Software,” and “Other.” The corresponding industry controls are named: *I_Health*, *I_Comm*, *I_Comp*, *I_Cons*, *I_Soft*, and *I_Other*. These broad industry classifications are aggregated from 25 minor classifications. The aggregation is necessarily somewhat arbitrary, but the intent is to classify companies evenly into industries where experience with one company helps predict outcomes from investing in other companies in this same industry.

The data contains information about the stage of development of the companies at the time of each investment. Companies are either at the early stage or the late stage, where late stage roughly corresponds to companies having regular revenues. The binary variable *Stage* equals one for investments in late stage companies.

The outcome of each investment is given by the binary variable *Succ*. An investment is successful, if the company subsequently goes public or is acquired. While it is true that VC investors realize most of their return from a few successful investments, this is clearly a coarse outcome measure. A more attractive measure would measure success in dollars or percentage return, but given the available data, it is difficult to refine the measure, and the binary outcome measure is common in the literature. Gompers and Lerner (2000) compare this measure to a broader range of outcome measures and find they lead to qualitatively similar results. For each investor, the variable *SuccRate* measures performance as the number of successful investments divided by the total number of investments in the sample.

The total number of investments each year in each industry is calculated and the corresponding variable is named *IndustryPerYear*. This captures general investment trends, such as the general shift towards computer related investments in the late nineties. It varies from a low of 24 investments in the “Consumer / Retail” industry during 1989 to a high of 1,714 investments in the “Computer Hardware / Electronics” industry during 2000.

C. Constructing Expected Return and Option Values

One consequence of the numerical difficulty of calculating the Gittins index and the option value from the investors’ preferences, is that these preferences are calibrated rather than estimated. While it may be possible to extend the estimation method to recover these preferences, it would add substantial numerical complexity to the estimation procedure, and it goes beyond the scope of this analysis.

For each investor, at the time of each investment, the investors’ expected return and option value from investing in each of the industries are calculated. To do this, the investors start out with prior beliefs about the distribution of p_i for each industry. These prior beliefs are unobserved, and I assume they follow a $Beta(a_{i,0}, b_{i,0})$ distribution, with $a_{i,0} = 1$ and $b_{i,0} = 19$. This roughly corresponds to each investor having observed one previous success and 19 previous failures in this industry, and means that the investor’s initial expected success probability is $\lambda = 1/20 = 5\%$. This may seem low compared to the empirical success rate of 50%, but many VC investors are specialized investors, and the prior expected return is then the expected success rate of an investor specialized in industry i_1 and switching to industry i_2. The returns in the industries they have specialized in are likely higher, but with a higher prior probability, the model would be unable to capture the specialization. To allow the investor’s beliefs over outcomes to “burn-in” and to decrease the dependency on the assumption about the prior beliefs, the initial ten investments made by each investor are only used to update the investor's beliefs and are not included in the estimation. This means that an investor that does not invest at all in an industry during the first ten investments has an expected return of 5% from investing in this industry. For the industries the investor does invest in, the expected return is likely higher (if the investments were successful). After each investment, $a_{i,j}$ and $b_{i,j}$ are updated, where $a_{i,j}$ contains the number of successful investments by investor j in industry i, and $b_{i,j}$ contains the number of failed investments by the investor in the industry. For each investment, the investor's expected success probability is calculated as $\lambda_{i,j,t} = a_{i,j,t} / (a_{i,j,t} + b_{i,j,t})$. The option value of investing is calculated as $OptVal_{i,j,t} = \left(A(\lambda_{i,j,t}) + B(\lambda_{i,j,t})(a_{i,j,t} + b_{i,j,t}) \right)^{-1}$, from equation (7). Thus, the total value of investing in the model is $v_{i,j,t} = \lambda_{i,j,t} + OptVal_{i,j,t}$.

**** TABLE 1 ABOUT HERE ****

III. Empirical Bandit Model

A. Empirical Specification of Bandit Model

Let the investors in the sample be indexed by $j = 1, ..., J$. At the time of each investor's investment, the empirical specification of the value of investing in industry i is the following.

$$v_{i,j,t} = \beta_0 + \lambda_{i,j,t}\beta_1 + OptVal_{i,j,t}\beta_2 + IndustryPerYear_{i,t}\beta_3 + X'_{i,j,t}\beta_4 + \varepsilon_{i,j,t} \qquad (8)$$

Here $\lambda_{i,j,t}$ is the immediate expected return, $OptVal_{i,j,t}$ is the option value, and $X_{i,j,t}$ contain control variables. The construction of the immediate return and the option value is described above. It is important to note that these values depend only on each investor's beliefs about the distribution of outcomes, and not on the true success probability p_i. Thus, the analysis does not require p_i to be constant across investors. The investor then chooses to invest in the industry with the highest total value. When the error term, $\varepsilon_{i,j,t}$, follows an Extreme Value distribution,[12] the probability of observing investor j investing in industry i' at time t is given as follows.

$$\Pr[i_{j,t} = i' \mid F_{i,j,t}] = \frac{\exp(v_{i',j,t})}{\sum_{i=1,\dots K} \exp(v_{i,j,t})} \tag{9}$$

This empirical model is equivalent to the Multi-Nomial Logit model (also known as the McFadden Choice model, see i.e. McFadden (1972)). The estimated coefficients of equation (8) are reported in Table 2.

**** TABLE 2 ABOUT HERE ****

B. Empirical Results

Consistent with the model, Specification 1 in Table 2 shows that the estimated coefficients on both $\lambda_{i,j,t}$ and $OptVal_{i,j,t}$ are positive and significant. Investors prefer to invest in industries where they expect a higher immediate return, represented by $\lambda_{i,j,t}$, and they prefer industries with higher option value, represented by $OptVal_{i,j,t}$.

In the following specifications, I control for other factors that may affect investment decisions. First, it may be that investors learn from the investments made by other

[12] The Extreme Value distribution (of type I) has distribution function $F(s) = e^{-e^{-s}}$, it has mean of roughly 0.577 (equals the Euler-Mascheroni constant), and it has variance $\pi^2 / 6$.

investors or are driven by general trends in the market. Specification 2 controls for the total number of investments made each year in each of the six industries. It is found that this has a small but positive and significant effect on the investment decisions. More importantly, including the year control in the regression does not eliminate the effects of $\lambda_{i,j,t}$ and $OptVal_{i,j,t}$, and the predictive power of the learning model is not diminished.

Further, it may be that investors are inherently timid or face costs of switching into new industries. Then they would remain invested in a few industries, and the option value could merely be an (inverse) proxy for the number of investments an investor has made in each industry. To test this "switching cost" alternative, Specification 3 includes the variable *IndustryPerInvestor* which contains number of investments each investor has previously made in the industry. The estimated coefficient is positive and significant, but again the coefficients on $\lambda_{i,j,t}$ and $OptVal_{i,j,t}$ remain positive and significant. Note that the Bandit model does predict a positive effect of total industry investments. When investors receive positive information about the returns in a given industry, they tend to invest there, and it is not surprising that the estimated coefficients of $\lambda_{i,j,t}$ and $OptVal_{i,j,t}$ change when *IndustryPerInvestor* is included in the regression. But, $\lambda_{i,j,t}$ and $OptVal_{i,j,t}$ are not statistically eliminated, and the learning model is not rejected relative to the "switching cost" hypothesis. The final specification includes all regressors, and the results remain unchanged. I conclude that the learning model does capture at least some aspects of VCs' investment behavior reasonably well.

In the first specification, the estimate of the coefficient on $\lambda_{i,j,t}$ is 4.480 and the estimated coefficient on $OptVal_{i,j,t}$ is 38.177. The Bandit model predicts that these two coefficients should be equal, and that the immediate payoff and the option value should be equally valuable for investors. The greater coefficient on $OptVal_{i,j,t}$ suggests that investors place more weight on the option value than predicted by the model, or, equivalently, follow more explorative strategies than the model predicts. Across all specifications, the coefficient on $OptVal_{i,j,t}$ exceeds the coefficient on $\lambda_{i,j,t}$. There are (at least) four possible explanations for this. First, it may be that the prior distributions are not sufficiently dispersed. If actual priors are more dispersed than the priors used to calibrate the learning model, this would lead the investors to experiment more than this model predicts. Second, if the investors' discount rates are smaller (δ closer to one) than in the calibrated model, this leads investors to assign more value to returns in the future and this increases the option value relative to what is predicted by the model. Third, the model may be structurally misspecified in the sense that investors have access to additional information or the environment may be changing, leading them to explore in different ways than the model predicts. Finally, investors may be behaving suboptimally relative to the model. Distinguishing these explanations is interesting, but it is infeasible in the current setup, and the analysis below does not hinge on this distinction.

IV. Classifying Investment Strategies

As shown above, the Bandit model presents a reasonable model of VCs' investment decisions. Now this model is estimated for each investor individually. The regression estimated for each investor specifies the value of an investment as follows.[13]

$$v_{i,j,t} = \beta_{i,1}\left[\lambda_{i,j,t} + OptVal_{i,j,t}\right] + \beta_{j,2} OptVal_{i,j,t} + \beta_{j,3} IndustryYear_{j,t} + \varepsilon_{i,j,t} \qquad (10)$$

The bracket contains the sum of the immediate return and the option value. This is the total value of investing, equal to the Gittins index, $v(F_{i,j}(t))$, in the Bandit model. The second term is the option value, and $\beta_{j,2}$ captures an investor's tendency to place too large or too small a weight on this value. An investor with a high $\beta_{j,2}$ exhibits a more explorative investment behavior. As argued above, there are different possible explanations for this. Greater exploration may be due to an investor having more

[13] In contrast to the specification in equation (8) this equation is estimated without industry specific effects. Since the coefficients are investor specific, adding 5 industry specific effects per investor would add a total of 210 x 5 = 1,050 coefficients to be estimated. Not surprisingly, this substantially reduces the statistical power of the analysis.

dispersed priors, a lower discount rate, sub-optimal behavior, or a combination of these factors. Regardless of the explanation for the difference in investment strategies, $\beta_{j,2}$ does classify investors into those who follow more or less explorative strategies. The resulting ranking of the investors is reasonably robust over different assumptions about their prior beliefs.

For each investor, the coefficients $\beta_{j,1}$, $\beta_{j,2}$, and $\beta_{j,3}$ are estimated. Descriptive statistics of these estimates are in Panel B of Table 1.[14] There is wide variation in the estimates, but the largest variation is caused by investors that make relatively few investments in the data, and where the estimated values of $\beta_{j,1}$, $\beta_{j,2}$, and $\beta_{j,3}$ have substantial estimation errors. To control for the uncertainty in estimating these coefficients, each investor is weighed according to the precision of these estimates below.[15]

In the standard Multinomial Logit model, the scale of the parameters in normalized by fixing the variance of the error term. This normalization is less convenient here, and a renormalization where $\beta_{j,1}$ is normalized to one allows for interpreting the estimated coefficients. This corresponds to estimating the model fixing the scale by normalizing the first coefficient rather than the variance of the error term. The remaining coefficients are renormalized accordingly.[16]

$$\widehat{\beta}_{j,2} = \beta_{j,2} / \beta_{j,1} \qquad \widehat{\beta}_{j,3} = \beta_{j,3} / \beta_{j,1} \qquad \widehat{\sigma}_{\varepsilon'} = 1 / \beta_{j,1} \tag{11}$$

[14] In this table, the estimate of $\beta_{j,1}$ is denoted *Gittins* (*Est*) and its variance is *Gittins* (*Var*). The estimate of $\beta_{j,2}$ is called *OptionValue* (*Est*), and the estimate of $\beta_{j,3}$ is *IndustryPerYear* (*Est*).

[15] For investor *j*, $Precision_j = 1/\sigma^2_{\gamma_{j,1}}$. The results are robust to other precision measures.

[16] In Table 1, the renormalized variables are listed as *OptionValue* (*Renorm*), *IndustryPerYear* (*Renorm*), and *StandardError* (*Renorm*), respectively.

Here, $\widehat{\beta}_{j,2}$ measures the relative propensity of an investor to explore; $\widehat{\beta}_{j,3}$ measures the investors' propensity to follow the overall trends in the market; and $\widehat{\sigma}_{\varepsilon'}$ is the standard error of the error term in the renormalized equation (10).[17]

Investment strategies can now be classified along these three dimensions. A high value of $\widehat{\beta}_{j,2}$ corresponds to an investment strategy that places relatively more weight on forward looking learning and exploration. A high value of $\widehat{\beta}_{j,3}$ corresponds to an investment strategy that places relatively more weight on overall market trends. A high value of $\widehat{\sigma}_{\varepsilon'}$ corresponds to an investment strategy with a larger standard error of the error term in the renormalized Equation (10). One can interpret this as the investor being more opportunistic or less focused.

A. Performance of VCs

It is natural to ask how individual VCs' performance depends on their investment strategies, as characterized above. Here, investor *j*'s performance is measured by $SuccRate_j$. Estimates from regressing the success rate on the characteristics of the investment strategy are presented in Panel A of Table 3.

**** TABLE 3 ABOUT HERE ***

---

[17] The extreme value distribution has variance $\pi^2/6$, so formally $SE(\varepsilon') = \left(\pi/\sqrt{6}\right)1/\beta_{j,1}$, where ε' is the error term in the renormalized equation. The $\left(\pi/\sqrt{6}\right)$ multiple is ignored below. For some investors $\beta_{j,1}$ is negative, leading to the slightly curious result that these investors have negative $\widehat{\sigma}_{j,1}$. Naturally, these are investors with short investment histories and large standard errors of the estimate of $\beta_{j,1}$. Excluding them from the sample does not alter the results. In fact, since the investors are weighted according to the standard error of $\beta_{j,1}$, these investors carry little weight in the following analysis.

In Table 3, Specification 1 characterize the investment strategy by $\widehat{\beta}_{j,2}$ and $\widehat{\sigma}_{\varepsilon'}$, and estimate the success rate as a function of these two variables. It is found that investors with higher $\widehat{\beta}_{j,2}$ (*OptionValue* (*Renorm*)), and thus more explorative investment strategies, have higher success rates. Investors with higher $\widehat{\sigma}_{\varepsilon'}$ (*StandardError* (*Renorm*)), that is investors that are more opportunistic, have lower returns.

The second specification also includes $\widehat{\beta}_{j,3}$ (*IndustryPerYear* (*Renorm*)) in the specification. This captures investors' tendency to follow general trends in the market. The coefficient is positive and significant, indicating that investors that follow overall trends more have higher returns. However, some caution is necessary when interpreting this result. The period late in the sample saw a large number of investors entering in "Hot" industries. In the estimates, these investors have low values of $\widehat{\beta}_{j,3}$, since they invest little outside these industries, and they have lower success rates. The estimation here does not control for entry and exit by the investors and, as shown below, the result is not robust to conditioning on each individual investment.

Finally, the third specification includes the total number of investments observed for each investor in the sample. One may suspect that investors with higher success rates can make more investments and gain higher experience. However, the estimate is negative and significant, albeit small in economic magnitude. One interpretation, consistent with Kaplan and Schoar (2005) is that successful investors subsequently raise larger funds and this decreases their overall returns.

In Table 3, Panel B reports estimated coefficients from a Probit model. This model estimates the probability of success for each individual investment as a function of the characteristics of the investor's investment strategy and additional control variables. Again, Specification 1 shows that investments made by more explorative investors are more likely to be successful, and investments made by more opportunistic investors are less likely to be successful. The economic effects are substantial. A one standard deviation increase in *OptionValue* (*Renorm*) or *StandardError* (*Renorm*) both lead to approximately 8% percentage point decrease in the success probability.

In Specification 2, the effect of $\widehat{\beta}_{j,3}$ reverses relative to the estimates above, although the statistical significance of this variable is low. This is consistent with the interpretation that each individual investment, if it is made by an investor that follows overall trends more, is less likely to be successful. However, investors that follow overall trends have higher success rates overall, because they enter in periods with higher overall returns.

The final specification includes a number of additional controls. Controlling for the stage of the company when it receives the investment shows that investment in late stage companies are more likely to be successful, which is not surprising. Again the investors' experience has a negative effect on the success probability. The specification also includes industry and year controls, but the results remain largely unchanged.

V. Conclusion – Summary and Extensions

I propose a new source of value generated by venture capital investors. The organizational structure of VCs permits them to internalize the benefits of learning about new industries and take this into account in their investment strategies. This insight is capture by a simple Bandit model. I estimate this model, and the empirical results suggest that the model does explain investment behavior and learning at the industry level reasonably well. The model allows me to categorize investment strategies as being more or less exploratory, more or less driven by overall market trends, and more or less opportunistic. I find that investment strategies that are more exploratory and less opportunistic perform better. The results are more mixed for investment strategies that follow market trends to a smaller or greater extent.

In addition, the paper proposes a methodological contribution. While the Bandit model has been widely studied in the theoretical literature previous empirical applications have been more limited, and it has primarily been used to derive job turnover rates in the labor literature. While the present implementation is arguably crude, the results are encouraging and further refinements of this model will likely lead to a better understanding of learning processes in venture capital and other markets.

The model estimated here is based on the simple Bandit model. The theoretical literature has explored a number of theoretical extensions of this model that could help extend the empirical analysis.

One interesting refinement of the model considers learning at the technology level rather than the industry level. This presents both an empirical and theoretical challenge. Empirically, this requires classifying the technologies or ideas underlying each investment. Theoretically, it requires understanding investment decisions with informational spillovers between closely related technologies, which presents a greater challenge but also an interesting avenue for future research.

Finally, the model assumes a stationary environment, but finds that investors follow general trends in the market, which supports the view that structural changes and new investments are important determinants of investments. The present model does not capture this in an entirely satisfying way. However, the theoretical analysis points to two reasons why an industry or technology can become attractive. It may be because the immediate returns increase, for example because of improved production methods of proven products. Or, it can be because the option value increases, which would characterize new and unproven technologies with uncertain potential. The analysis here suggests that spot markets are well adapted to respond to increases in immediate returns, while VC investors are probably better at internalizing increases in option value. Either way, the distinction between these two determinants of investment trends seems important for understanding the roles of VCs and spot markets for allocating capital and commercializing innovations.

Bibliography

Axelson, Ulf, Per Stromberg, and Michael Weisbach, 2006, The Financial Structure of Private Equity Funds, *mimeo*.

Bergemann, Dirk, and Ulrich Hege, 1998, Dynamic Venture Capital Financing, Learning and Moral Hazard, *Journal of Banking and Finance*, 22, 703-735.

Bergemann, Dirk, and Ulrich Hege, 2005, The Financing of Innovation: Learning and Stopping, *RAND Journal of Economics*, 36.

Bergemann, Dirk, and Juuso Valimaki, 2006, Bandit Problems, *Cowles Foundation Discussion Paper, no. 1551*.

Berry, Donald, and Bert Fristedt, 1985. *Bandit problems: Sequential Allocation of Experiments* (Chapman and Hall, New York).

Bolton, Patrick, and Christopher Harris, 1999, Strategic Experimentation, *Econometrica*, 67, 349-374.

Gittins, J. C., and D. M. Jones, eds., 1974. *A Dynamic Allocation Index for the Sequential Allocation of Experiments* (North-Holland, Amsterdam).

Gittins, J. C., and D. M. Jones, 1979, A Dynamic Allocation Index for the Discounted Multiarmed Bandit Problem, *Biometrika*, 66, 561-565.

Gompers, Paul, and Josh Lerner, 1999. *The Venture Capital Cycle* (MIT Press, Cambridge).

Gompers, Paul, and Josh Lerner, eds., 2000. *The Determinants of Corporate Venture Capital Successes: Organizational Structure, Incentives, and Complementarities* (University of Chicago Press, Chicago, IL).

Gorman, M, and William Sahlman, 1989, What Do Venture Capitalists Do? *Journal of Business Venturing*, 4, 231-248.

Heckman, James, and George Borjas, 1980, Does Unemployment cause Future Unemployment? Definitions, Questions and Answers form a Continuous Time Model of Heterogeneity and State Dependence, *Econometrica*, 47, 247-83.

Hellmann, Thomas, and Manju Puri, 2000, The Interaction Between Product Market and Financing Strategy: The Role of Venture Capital, *The Review of Financial Studies*, 13, 959-984.

Hellmann, Thomas, and Manju Puri, 2002, Venture Capital and the Professionalization of Start-Up Firms: Empirical Evidence, *Journal of Finance*, 57, 169-197.

Hochberg, Yael, Alexander Ljungqvist, and Yang Lu, 2005, Venture Capital Networks and Investment Performance, *Journal of Finance, Forthcoming*.

Jovanovic, Boyan, 1979, Job Matching and the Theory of Turnover, *Journal of Political Economy*, 87, 972-90.

Kaplan, Steven, and Antoinette Schoar, 2005, Private Equity Performance: Returns, Persistence, and Capital Flows, *Journal of Finance*, forthcoming.

Kaplan, Steven, Berk Sensoy, and Per Strömberg, 2002, How Well do Venture Capital Databases Reflect Actual Investments? *working paper*.

Kaplan, Steven, and Per Strömberg, 2004, Characteristics, Contracts, and Actions: Evidence from Venture Capitalist Analyses, *Journal of Finance*, 59, 2177-2210.

Keller, G., S. Rady, and M. Cripps, 2005, Strategic Experimentation with Exponential Bandits, *Econometrica*, 73, 39-68.

Lerner, Josh, 1995, Venture Capitalists and the Oversight of Private Firms, *Journal of Finance*, 50, 301-318.
Manso, Gustavo, 2006, Motivating Innovation, *mimeo*.
Megginson, William, and Kathleen Weiss, 1991, Venture Capitalist Certification in Initial Public Offerings, *Journal of Finance*, 46, 879-903.
Miller, Robert A., 1984, Job Matching and Occupational Choice, *Journal of Political Economy*, 92, 1986-1120.
Robbins, H., 1952, Some Aspects of the Sequential Design of Experiments, *Bulletin of American Mathematical Society*, 55, 527-535.
Rothschild, M., 1974, A Two-Armed Bandit Theory of Market Pricing, *Journal of Economic Theory*, 9, 185-202.
Sahlman, William, 1990, The Structure and Governance of Venture Capital Organizations, *Journal of Financial Economics*, 27, 473-521.
Weitzman, M., 1979, Optimal Search for the Best Alternative, *Econometrica*, 47, 641-654.
Whittle, Peter, 1980, Multi-Armed Bandits and the Gittins Index, *Journal of Royal Statistical Society*. *Series B (Methodological)*, 42, 143-149.
Whittle, Peter, 1981, Arm-Acquiring Bandits, *Ann. Prob.*, 9, 284-292.
Whittle, Peter, 1982. *Optimization Over Time: Dynamic Programming and Stochastic Control* (John Wiley & Sons, New York).

TABLE 1: Summary Statistics

PANEL A: Summary Statistic By Company

	Obs.	Mean	Std. Dev.	Min	Max
IPO	5,621	0.167	0.373	0	1
ACQ	5,621	0.327	0.469	0	1
Succ	5,621	0.493	0.500	0	1
Year	5,621	1995.245	4.400	1987	2000
I_Health	5,621	0.175	0.380	0	1
I_Comm	5,621	0.209	0.406	0	1
I_Comp	5,621	0.240	0.427	0	1
I_Soft	5,621	0.178	0.383	0	1
I_Cons	5,621	0.122	0.327	0	1
I_Other	5,621	0.076	0.266	0	1

PANEL B: Summary Statistics by Investor

	Obs.	Mean	Std. Dev.	Min	Max
IpoRate	210	0.202	0.094	0.000	0.523
AcqRate	210	0.299	0.066	0.118	0.492
SuccRate	210	0.501	0.118	0.133	0.864
TotalInvestmentsPerInvestor	210	83.576	48.065	40	260
HH_PerInvestor	210	0.267	0.105	0.174	0.904
Characterizing Investment Strategies					
Gittins (Est)	210	4.331	5.707	-25.617	21.038
Gittins (Var)	210	12.834	19.168	0.236	116.724
OptionValue (Est)	210	42.083	188.961	-469.397	858.634
IndustryPerYear (Est)	210	0.001	0.002	-0.003	0.013
OptionValue (Renorm)	210	-20.511	196.145	-1746.945	554.753
Standard Error (Renorm)	210	0.064	1.092	-10.380	2.794
IndustryPerYear (Renorm)	210	0.000	0.001	-0.011	0.008

PANEL C: Summary Statistics by Investment

	Obs.	Mean	Std. Dev.	Min	Max
Stage	17,551	0.286	0.452	0	1
Year	17,551	1995.241	4.531	1987	2000
Succ	17,551	0.566	0.496	0	1
IndustryPerYear	17,551	286.016	94.354	121	407
InvestmentsPerInvestor	17,551	56.106	47.307	1	260
Lambda	17,551	0.246	0.159	0.038	0.681
OptionValue	17,551	0.009	0.002	0.004	0.014
Gittins	17,551	0.255	0.160	0.043	0.686
OptionValue / Gittins	17,551	0.055	0.032	0.006	0.107

TABLE 2: Aggregate Investment Decisions

The endogenous variable is the investors' industry choice in a Multi-Nomial Logit model, where the possible choices are Health, Communications, Computers, Consumer Goods, Software, and Other. Only investments by investors having made more than ten investments are used for estimation. *Lambda* and *OptionValue* are an investor's expected immediate return and option value of investing. *IndustryPerYear* is the total number of investments in each industry per year. *IndustryPerInvestor* is the total number of past investments the investor has made in each industry. *Previous* is an indicator variable that equals one for the industry of the investor's past investment. *I_XXXX* are six industry fixed effects, the excluded base group is *I_Other*. The observations are weighted according to the precision of the estimate of OptionValue (see text for details). Standard errors are in parenthesis. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	1			2			3			4		
	Coef.	Std.Err.		Coef.	Std.Err.		Coef.	Std.Err.		Coef.	Std.Err.	
Lambda	4.480	(.081)	***	4.035	(.099)	***	0.745	(.199)	***	0.669	(.201)	***
OptionValue	38.177	(4.377)	***	11.577	(4.607)	**	93.089	(6.184)	***	84.757	(6.256)	***
IndustryPerYear				0.001	(.000)	***				0.001	(.000)	***
IndustryPerInvestor							0.045	(.002)	***	0.044	(.003)	***
Previous										0.260	(.019)	***
I_Health				0.764	(.049)	***	0.891	(.049)	***	0.704	(.049)	***
I_Comm				0.856	(.048)	***	1.114	(.047)	***	0.782	(.048)	***
I_Comp				0.478	(.051)	***	0.917	(.048)	***	0.397	(.051)	***
I_Cons				0.924	(.047)	***	0.984	(.047)	***	0.866	(.047)	***
I_Soft				0.776	(.049)	***	0.865	(.049)	***	0.735	(.049)	***
Observations	15,495			15,495			15,495			15,495		

TABLE 3: Investment Strategies and Outcomes

Panel A shows estimated coefficients from OLS regression where each investor is an observation and the endogenous variable is the investor's success rate. Panel B presents marginal effects estimated from a Probit model where an observation is an investment and the endogenous variable is the investment's outcome. *OptionValue*, *StandardError*, and *IndustryPerYear* characterize each investor's investment decisions in terms of dependence on option value, its standard error, and its dependence on number of investments in the industry in same year. These coefficients are renormalized coefficients (see discussion in text). *InvestorExperience* measures the number of previous investments by the investor in each industry at the time of each investment. *Stage* is an indicator variable that equals one for investments in late stage companies. The observations are weighted according to the precision of the estimate of OptionValue (see text for details). Standard errors in parenthesis are calculated with clustering at the company level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

PANEL A: Success Rate by Venture Capital Firm

	1			2			3		
	Coef.	Std.Err.		Coef.	Std.Err.		Coef.	Std.Err.	
OptionValue (Renorm)	0.00018	(.00006)	***	0.00018	(.00006)	***	0.00017	(.00006)	***
Standard Error (Renorm)	-0.03171	(.01245)	**	-0.03846	(.01266)	***	-0.03407	(.01277)	***
IndustryPerYear (Renorm)				12.76616	(4.62939)	***	12.29804	(4.62508)	***
InvestorExperience							-0.00009	(.00004)	**
R^2	0.0091			0.015			0.0194		
Observations	1,260			1,260			1,260		

PANEL B: Success of Individual Investments

	1			2			3		
	dF/dX	Std.Err.		dF/dX	Std.Err.		dF/dX	Std.Err.	
OptionValue (Renorm)	0.00043	(.00013)	***	0.00043	(.00013)	***	0.00038	(.00013)	***
Standard Error (Renorm)	-0.07894	(.02733)	***	-0.07642	(.02670)	***	-0.06066	(.02669)	**
IndustryPerYear (Renorm)				-3.21315	(6.96695)		-8.35237	(7.00199)	
InvestorExperience							-0.00160	(.00013)	***
Stage							0.15299	(.01644)	***
Industry Controls	No			No			Yes		
Year Controls	No			No			Yes		
Observations	17,551			17,551			17,551		