Revenue management for venture capital

John Hawksley

August 31, 2009

1 Introduction

We present a model for investment selection in venture capital, assuming that prospective deals have known expected value. In particular, we express the screening process for new companies as an optimization problem following the template of revenue management, a subfield of Operations Research.

Revenue management is concerned with maximizing total income from a fixed inventory, such as airline tickets or hotel rooms. In the standard model, an agent has some number of units to sell, and customers arrive one at a time each stating the price he is willing to pay. For example, an airline has to decide whether to sell off its tickets cheaply now, or wait and hope for an influx of business travelers willing to pay full price.

A similar setting applies to the venture capital business, in which there is a fixed time to make a certain number of investments out of a fund. The investment manager chooses between investing in each prospective company (i.e. selling a ticket) or passing on it, and is tasked with dynamically updating his criteria over time.

This optimization problem was heavily studied in the early 1990s, leading to systems that update prices after each sale, depending on the units remaining and time before perish. Behavioral scientists have found that automated systems optimize more effectively than human agents, and in many cases that human errors are highly systematic.

In this paper we walk through the exercise of writing venture capital as a revenue management problem.

2 The Cutoff Point

If we value each company with a single number p, then what a firm needs to know is its **cutoff value**; that is, the threshhold $p^\star$ such that it should invest in any company with value exceeding $p^\star$ and pass otherwise. The magnitude of $p^\star$ depends on a firm's dealflow, the amount of money it needs to put to work, and the time left to invest that money. For example if we have lots of great dealflow and few investments left then $p^\star$ should be very high.

We will calculate the optimal $p^\star$ under some simple conditions. In particular, we assume the following:

1. Prospective deals are independently drawn from a uniform distribution over $p \in [0, 1]$.
2. The utility function of the investors is linear (in reality GP utility is a call function, since compensation depends on a floor return).
3. The number of investments in a fund, and the time restriction are known. Specifically, m is the # of new investments a firm can make out of it's remaining capital. And n is the # of prospective deals the firm will have the opportunity to consider in the timeframe allotted.

These assumptions are not required in principle - we use a simple dealflow distribution for simplicity, and a linear utility function so that $p^\star$ does not depend on previous exits in the fund. Thus $p^\star(n, m)$ depends on n and m only.

3 Solving for $p^\star$

We first comment on the steady-state intuition of our model to show that certain predictions follow plainly, then dedicate most of this section to writing down a finite version.

3.1 Steady-state behavior

Suppose r is the ratio of dealflow to investments, that is $r = \frac{n}{m}$. If we specify r, and let n and m get very large then $p^\star(r) \Rightarrow \frac{r-1}{r}$ since dealflow is assumed to be uniform. That is, the cutoff $p^\star$ is set to exactly the ratio of deals that we have the capacity to accept. This simple equation predicts that increasing dealflow in a venture capital firm has diminishing returns. For example, if r is 10 then $p^\star$ is 0.90, but if r doubles to 20 then $p^\star$ only improves to 0.95 – indicating that it can be counterproductive to add more partners (i.e. dealflow capacity) while holding fund size constant.

3.2 Finite generalization

To determine $p^\star(n,m)$ for any n,m we introduce an intermediate function $f(n,m)$, where f is the expected return on a partial fund with n dealflow and m investments. We can express the relationship between $p^\star$ and f in terms of the following recursion:

$$f(n-1,m) - f(n-1,m-1) = p^\star(n,m) \tag{1}$$

$$f(n,m) - f(n-1,m) = 0.5[1 - p^\star(n,m)]^2 \tag{2}$$

$$f(0,m) = f(n,0) = 0 \tag{3}$$

$$p^\star(n,m) = 0 \text{ if } n = m \tag{4}$$

Proof. Eq. (1) is given by the definition of $p^\star$, namely one invests in a package at time (n,m) iff the value of the partial fund given *non-invest* exceeds by at least $p^\star(n,m)$ the value of the partial fund given *invest*. To see (2), observe that the left-hand side is the expected benefit beyond $f(n-1,m)$ of seeing an extra deal. The probability of investing in the extra deal is $1 - p^\star(n,m)$, and the incremental value is the average net above $p^\star$, or $\frac{1}{2}[1 - p^\star(n,m)]$.

For example, the recursion (1-4) was used to compute $f(100,m)$ for a uniform dealflow distribution normalized to $[-0.85, 0.15]$. We calculate that the optimal fund size in this case is $m = 13$. That is, a firm should plan to make slightly fewer investments than the total number expected to have positive EV (in this case $m = 15$). In other words, it's nearly good enough for a deal to have positive expected return. This is just what we'd hope for, since assigning precise values to early-stage companies is notoriously unreliable, and checking this simpler condition may be significantly easier.

4 Conclusions

We lay the groundwork for a simple model of venture capital investing, based on the language of revenue management. Our model isolates the uncertainty in venture capital to measuring the package value of a company. It also provides a setting in which to behaviorally analyze the decisions of investment managers. For future research, we are interested in (1) replacing the dealflow and utility distributions with empirics via computer simulation, and (2) introducing information uncertainty, which relates to other fields of research.