



SANTÉ VENTURES

Frost Bank Tower
401 Congress Ave., Suite 2950
Austin, TX 78701
(512)721-1200
www.santeventures.com

Why Venture Capital Does Not Scale

Kevin M. Lalande, Managing Director
March 15, 2011

Introduction

No venture fund larger than \$750M has ever returned more than 2.0x to its limited partner investors. Fewer than a dozen funds larger than \$300M have. On the other hand, over 250 funds smaller than \$300M have cleared that same bar (Figure 1)¹. A similar gap exists when performance is measured by cumulative IRR, which is 15.9% in funds smaller than \$250M versus just 4.3% in funds larger than \$250M (Figure 2)².




Figure 1

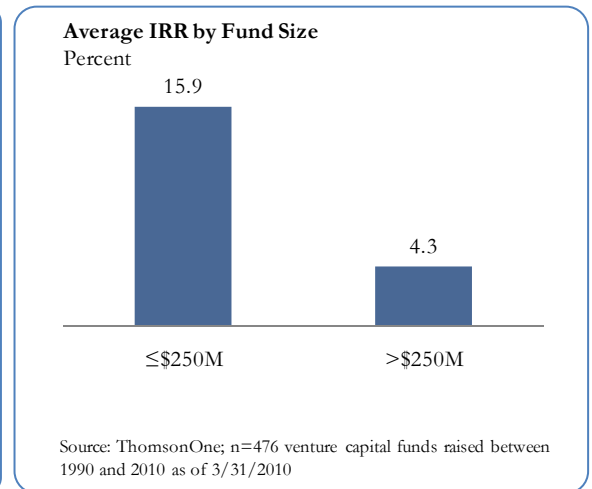


Figure 2

In spite of the clear performance difference by fund size, venture capital has become increasingly concentrated in large funds since 1998, with 50-60% of all new capital raised in the asset class committed to funds larger than \$300M. Before 1998, fewer than 40 venture funds over \$300M had ever been raised. Since then, more than 600 funds have. Unsurprisingly, overall returns for the entire asset class have declined over that same period (Figure 3). This increasing concentration, combined with the systemic challenges in the U.S. market for initial public offerings³, accounts for much of the reason venture capital returns have struggled in the last 10 years.




Figure 3

¹ Preqin Private Equity Performance Analyst

² ThomsonOne Private Equity; Calculated for all U.S. venture funds raised since 1990

³ Beyond the scope of this paper; best articulated in "Market Structure is Causing the IPO Crisis" by David Weild and Edward Kim, June 2010

The inescapable conclusion from this data is that larger venture funds generally underperform smaller funds, despite the presumed advantages of track record and brand name that allowed them to raise such large pools of capital in the first place. *But why is that?*

Commonly cited explanations include: too many dollars per general partner, too many portfolio companies per general partner, and too much management fee per firm. Each of these contains an element of truth, but none fully explains why returns in venture decline with increasing amounts of capital under management, unlike in other private equity asset classes like buyouts or real estate. For example, buyout funds larger than \$750M can return 2.0x or better to their limited partners, as 11% have managed to do since 1990, even though they have similar dollars per general partner and management fee profiles as a comparably sized venture fund⁴.

Exits tend to be small and infrequent

The performance constraint must be more fundamental to the venture capital business model itself. Venture funds generate returns for their investors only when they exit a successful portfolio company through a sale or public stock offering. So to better understand the core issue, we begin by analyzing the empirical data for successful venture-backed healthcare and life science portfolio company exits over the last decade, and then work backward to the implications on fund size and investment strategy. The same analysis and conclusions hold for the information technology sector⁵; we focus in this paper on the healthcare and life science sector given its relevance to our firm's particular investment strategy.




Figure 4

⁴ Preqin Private Equity Performance Analyst

⁵ Supporting data and analysis for IT or for all VC available upon request.

Figure 4 shows the frequency and valuation of every reported exit of a venture-backed U.S. healthcare or life science company in the decade between 2000-2009. Of 534 exits during that period, the average amount of equity invested was \$56M and the average exit value was \$156M⁶. By number of total companies, 45% exited at \$100M or less, 68% at \$200M or less, 83% at \$300M or less and 90% at \$400M or less⁷.

Of course, this histogram includes only those companies that were successful enough to exit; it does not show those that went bankrupt or the so-called “living dead,” which have neither gone out of business nor exited through M&A or IPO. A total of 2,599 U.S. healthcare and life-science companies received venture capital between 1994-2003, which is a reasonable estimate of the corresponding peer group of companies since the average hold time at exit is 6 years. Thus, only 1 in 5 venture-backed healthcare and life science companies funded from 1994-2003 went on to exit in 2000-2009⁸.

The data underscore what we already know intuitively from experience: exits are few and far between in the venture business, with risk of capital loss and long gestation periods associated with each individual investment. Unlike the buyout industry, where exit valuations scale linearly from millions to billions based on the size of cash flows in the underlying companies, the venture industry tends to produce companies whose exit valuations cluster below \$300M because they are most often priced based on growth projections and proprietary technology, rather than on cash flow, given the earlier developmental stage of venture-backed companies. This clustering phenomenon is a critical distinction between venture and buyout, and the resulting average expected value of a successful exit places fundamental constraints on the upper limit of venture capital per company and on venture capital fund size, above which performance begins to drop sharply.

Obviously, average expected exit value constrains capital invested per deal. It does not make sense to invest into projects that require \$56M of equity⁹ into a company in a sector where the typical successful exit is \$156M, and where at least a third of the companies lose capital (expected value of 1.30x return¹⁰). This is an obvious point when stated in terms of averages, but many venture capitalists ignore the historical empirical data and focus instead on statistically improbable outliers – like the \$400M+ exits that represent less than 10% of all successful outcomes, and where the expected value of that same \$56M investment would be 3.33x. Although any single winning investment may be much larger than the average of \$156M, it is improbable that any given portfolio of winning deals will beat the 10-year historical averages.

Instead, the better investment strategy is to accept the historical average as a fixed constraint, and then limit portfolio companies to only the sub-set of high-quality projects whose capital requirements are also low enough to generate attractive multiples at average exit values, or less than \$25M in the case of a \$156M average exit (expected value of 2.91x). This isn't easy. In the life science sector, for example, most technologies require significantly more capital than that to commercialize. But given the historical average data, that just means that most life science technologies are not good venture investment candidates, regardless of how interesting the underlying market or technology appears. The alternative—mistakenly setting the assumed exit valuation too high—leads to the classic Lake Wobegon effect, where all winners have to be above average in order for the overall

⁶ We use the word “average” for linguistic simplicity, but technically calculate the 5% trimmed mean of the dataset, which ignores the highest 5% exit values.

Including these “long-tail” events produces a straight arithmetic mean of \$171M.

⁷ *ThomsonOne Private Equity*; Calculated based on all US venture-backed healthcare company exits from 2000-2009

⁸ *ThomsonOne Private Equity*

⁹ *ThomsonOne Private Equity*; Average total equity raised across all US venture-backed healthcare company exits from 2000-2009

¹⁰ Assuming investors own 70% of company. $EV = \$156M * (1-0.33) * 0.70 / \$56M$

portfolio to return an acceptable multiple because individual investments were underwritten that require too much capital.

Average exit valuations limit venture fund size

Less obviously, average exit value also constrains venture capital fund size. In order for a portfolio of venture investments to perform overall, winning companies must return “big” multiples. *But big relative to what?* **Winning deals must generate large returns as a fraction of total fund size, rather than as a multiple of capital invested.** This is an unintuitive but important point. For example, how relevant is a 15x return on a \$10M investment? The answer depends on fund size. In a \$100M fund, two of those deals would return 3x gross to the limited partners. But a \$1B fund, it’s a rounding error – unless it can be repeated 20 times. Because exits are few and far between in venture, high-impact deals are defined less by multiple of investment capital returned than by the fraction of total fund size returned, which is in turn a function of three variables: exit value, fund size and ownership percentage. The historical average expected exit value curve should be accepted as a fixed constraint and fund size and ownership percentage should therefore be designed accordingly to optimize the overall chances of delivering good fund performance.

To illustrate, consider the impact of exiting a 50% ownership position in a company valued at \$200M on two venture funds of different sizes: Fund A of \$100M and Fund B of \$500M. In Fund A, this single exit returns 100% of the entire fund and one-third of the total capital required to generate a 3x gross return to LPs. In Fund B, the same exit returns just 20% of the entire fund and one-fifteenth of the total capital required for a 3x return. Thus, Fund B must generate *five times as many* winning companies—or it must produce winners that are *five times larger*—than Fund A. In practice, however, the number of individual portfolio companies is not five times larger. Instead, Fund B typically writes bigger individual checks into a portfolio of 25-50 companies whereas Fund A writes smaller checks into a portfolio of 15-25 companies, which means Fund B is implicitly relying on a higher hit rate and on higher than average exits than Fund A to generate similar performance.

Another way to look at the impact of fund size is to consider how large (and correspondingly rare) an exit is required to return 1x the entire fund. As calculated above, at 50% ownership, a \$100M fund requires a \$200M exit, or about a third larger than the 10-year average exit of \$156M, to return 1x the entire fund. On the other hand, a \$500M fund requires a \$1 billion exit – which would rank in the 99th percentile of all successful exits in U.S. life sciences over the last decade. Referring back to Figure 4, 164 companies exited for \$200M or more (69th percentile and above), while only 4 exited for more than \$1B. *Given that, which fund is more likely to produce winners that return 1x the entire fund?*

A venture fund should be no larger than 1.0-1.5x the average equity value at exit of venture-backed companies in its investment sector. Fund performance drops off precipitously as the ratio of fund size to average exit value increases above that range, as evidenced in Figure 1 above, because returning a given multiple of capital to limited partner investors requires either increasingly improbable mega-exits within the portfolio (i.e., “grand slams”) or an increasingly unrealistic fraction of the total portfolio of companies to develop into successful investments (i.e., “batting average”).

In addition to fund size, the percentage ownership in the underlying portfolio companies is also critically important to fund performance. For example, 1x the \$100M in Fund A could be returned by exiting a single company with 60% ownership for \$167M (63rd percentile), or by exiting three companies with 20% ownership for \$167M each (63rd percentile x 3 independent events), or by exiting a single company with 20% ownership for \$500M (95th percentile). Although these three examples are mathematically equivalent, they are not statistically equivalent as the first outcome is much more probable than the other two.

Exit Value Required to Return 1.0x Fund
\$ Millions

Own	Fund Size \$ 100		Fund Size \$ 500	
	Exit Value	Percentile	Exit Value	Percentile
10%	\$ 1,000	99	\$ 5,000	99
20%	\$ 500	95	\$ 2,500	99
30%	\$ 333	86	\$ 1,667	99
40%	\$ 250	77	\$ 1,250	99
50%	\$ 200	69	\$ 1,000	99
60%	\$ 167	63	\$ 833	99
70%	\$ 143	57	\$ 714	98

Table 1

Venture capital fund performance is therefore highly sensitive to two design variables: 1) the ratio of fund size to average exit value, and 2) the average percentage ownership of its portfolio companies. Stitching both of these parameters together, the following table shows the difference in the number of portfolio companies that must succeed to generate a gross return of 3.0x on Fund A of \$100M versus Fund B of \$500M fund at a given average exit value and a given average portfolio company ownership percentage¹¹. *In other words, how many deals have to work to return 3x?*

Number of Successful Deals Required to Return 3.0x Gross Fund

Fund Size = \$100 Million									Fund Size = \$500 Million								
Average Ownership	Average Enterprise Value at Exit (\$Millions)								Average Ownership	Average Enterprise Value at Exit (\$Millions)							
	\$ 50	\$ 100	\$ 150	\$ 175	\$ 200	\$ 250	\$ 300	\$ 400		\$ 50	\$ 100	\$ 150	\$ 175	\$ 200	\$ 250	\$ 300	\$ 400
15%	33	17	11	10	8	7	6	4	167	83	56	48	42	33	28	21	
20%	25	13	8	7	6	5	4	3	125	63	42	36	31	25	21	16	
30%	17	8	6	5	4	3	3	2	83	42	28	24	21	17	14	10	
40%	13	6	4	4	3	3	2	2	63	31	21	18	16	13	10	8	
50%	10	5	3	3	3	2	2	1	50	25	17	14	13	10	8	6	
60%	8	4	3	2	2	2	1	1	42	21	14	12	10	8	7	5	
70%	7	4	2	2	2	1	1	1	36	18	12	10	9	7	6	4	

Figure 5

At a \$150M average exit value and a 40% average ownership position, Fund A needs only 4 winning companies in its portfolio to return 3x gross. Fund B needs 21. The latter is clearly unlikely, both as a fraction of the total portfolio and of the total number of exits in the entire U.S. over a 5-year time period (8% assuming the historical rate).

Notice also that both funds are extremely sensitive to average ownership percentage: as it drops from 40% to 15%, the number of winners required for Fund A grows to a less plausible 11 companies and for Fund B grows to a plainly unachievable 56. So in addition to keeping the fund size small, maintaining high ownership positions is also important. Doing so requires the fund to adopt an early-

¹¹ Assuming all other losing investments return just 0.5x the fund in aggregate.

stage entry strategy, since apart from distressed situations, there is almost no way to get 30-50% ownership positions when investing for the first time at the Series B round and beyond. It also puts significant downward pressure on the total capital requirements of the underlying portfolio companies in a relatively small fund. For example, if Fund A were to build a portfolio of 12 equally sized positions at 40% average ownership, then each company could consume only \$16-18M in total equity capital from entry to exit¹². This is challenging but can be managed through rigorous portfolio selection—that is, investing in only the sub-set of otherwise interesting opportunities that also meet the constraint of total capital required.

A common concern with a portfolio selection strategy that screens on capital efficiency is based on the mistaken belief that the most successful companies (i.e., the companies most likely to exit well above the mean) will also tend to raise the largest amounts of capital. In reality, however, the correlation between total amount of venture capital raised and ultimate exit value is very low (Figure 7). Similarly, there are plenty of examples of companies that raised less than \$25M in total equity and went on to exit for \$100-500M, as shown in the same scatter plot.

Of course, a large venture fund or one with a low average ownership percentage could still get lucky and return 3x based on one or more statistical outliers (e.g., 15% of a company that exits for \$2B). *But luck is not a strategy.* Well-designed venture funds should be no larger than 1.0-1.5x the average historical exit value of companies in their sector and should own an average of 30% or more of their underlying portfolio companies. That way, they don't have to hope to get lucky with a 99th percentile exit by valuation in their portfolio in order to produce reasonable overall returns. Creating winning companies is hard enough without having to also get lucky on a string of unlikely exit valuations to compensate for too big a fund or too much capital per deal.




Figure 6

enough without having to also get lucky on a string of unlikely exit valuations to compensate for too big a fund or too much capital per deal.

Why large venture funds continue to be raised

So, if it is structurally improbable that large venture funds will produce 2x or better returns regardless of prior track record in smaller funds or brand advantages, why do these large funds continue to be raised? The answer on the general partner side is simple: it's in their economic interest to do so if they can. As illustrated in Figure 7, the net present value of management fees *alone* for a fund at \$500M exceeds the *fees and carry* of a fund at \$100M, *even if it returned 4x*. This is not a particularly extreme example by fund size, so you can see how a general partner could rationally choose to operate a larger fund that is structurally likely to return just 1.0-1.5x rather than a smaller fund with the potential to return 2-3x or better.

¹² Assumes other investors accounted for equal funding and another 40%. Assumes management fees are re-invested. Ignores debt and non-dilutive financing in portfolio companies.

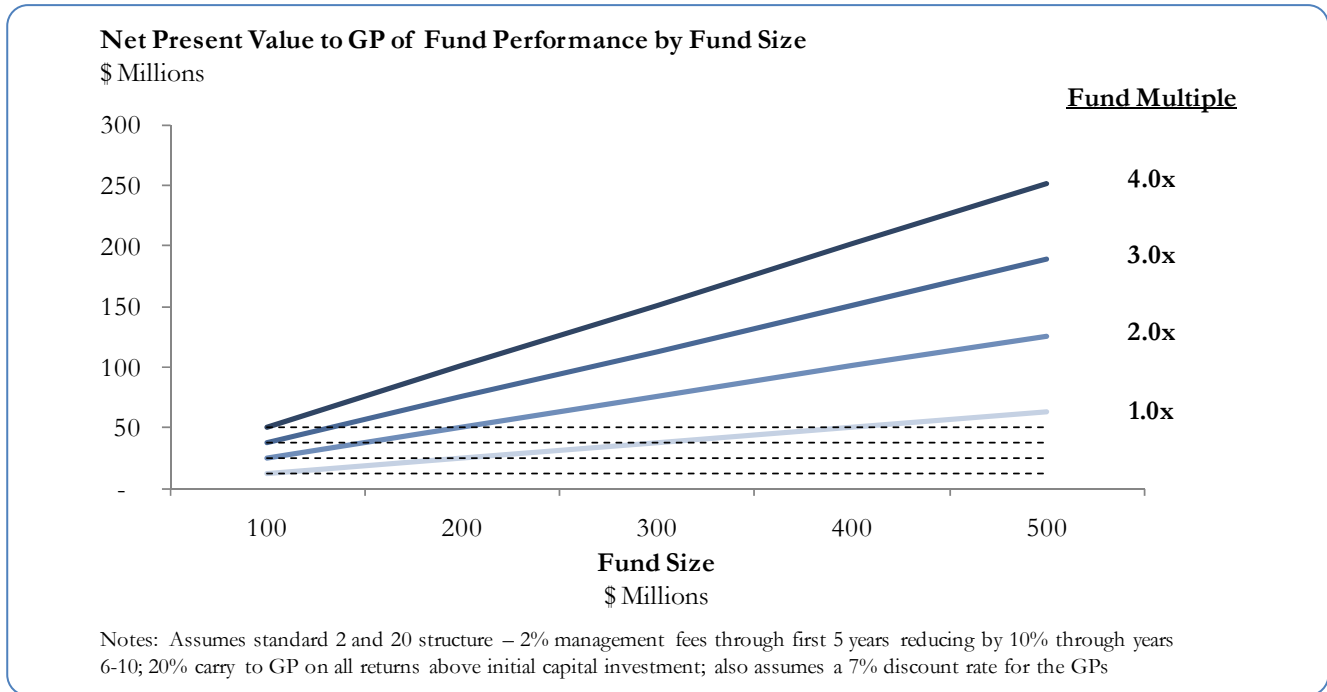


Figure 7

The answer on the limited partner side is more complicated. Venture funds larger than \$300M are a relatively new phenomenon over the last 10 years, so it has taken time for those funds to mature. Now that the results are in and the performance gap is clear, many limited partners are indeed migrating their manager relationships back toward smaller funds. Of course, most of the best-recognized venture firms are still raising large funds for the reason above. *So how are those getting raised?* Fund-of-funds by their very business model aim to balance both access to prestigious brand name firms and long-term investment performance, and those objectives can come into conflict when the best brands raise large funds. But even pension funds and endowments with no particular need for brand proximity can still tend towards large funds for a variety of reasons, including: a need to get a large number of dollars to work without wanting or being able to comprise more than 10% of an individual fund, a misunderstanding of how much harder it is to return 2x on funds larger than \$300M, a perception that smaller funds are riskier, or the perennial human hope that “this time will be different.”

Summary of the main idea

Venture capital performance does not scale well to larger fund sizes because venture-backed exits tend to cluster below \$300M since, unlike in the buyout industry, the underlying companies are rarely valued as a linear multiple of cash flow given their earlier stage of development. Exits are few and far between. So winning deals must generate large returns as a fraction of total fund size, rather than as a multiple of capital invested. Well-designed venture funds should therefore be no larger than 1.0-1.5x the average historical exit value of companies in their sector and should own an average of 30% or more of their underlying portfolio companies. That way, they don't have to hope for empirically improbable exit valuations in their portfolio to produce reasonable overall returns. Luck is not a strategy. Creating winning companies is hard enough without having to also put together a string of unlikely exit valuations to compensate for a fund that is too big, ownership levels that are too small, or individual portfolio companies that require too much capital.