

October 08, 2019

Measuring Venture Capital Returns: You Can't Spend IRR

Related Practices

Fund
Grow
Sell
Start
Venture Capital & Private Equity

Some time back, I wrote about how important diversification is for venture capital and angel investors. High risk/reward investing is a power hitter's game: you strike out a lot, and do most of your damage with the long ball. So taking a lot of cuts – doing a lot of deals – is critical to any strategy that relies on talent more than luck. See my previous article, titled “Hard Truths About Angel Investing (And Crowdfunding).”

When it comes to measuring venture investing returns, I've also written about the deficiencies of IRR as a measure of venture investing performance: see my previous article, titled “Measuring Startup Investing Returns.” And still, pretty much everyone does it. Mostly, I think, because it can make even middling performance look pretty good to folks who don't ask too many questions about how it works.

Which leads me to this blog.

Recently, I came across some Pitchbook data about the 10-year performance of a certain vintage and category of venture capital funds. The reported median IRR was 11%. That's not too shabby, even as too many folks have this idea that venture fund investors routinely shoot for IRRs of 35% and up. I mean, if you put in \$1 and made an 11% compounded annual return for ten years you would, at the end of that time, have roughly \$2.84.

Or would you? Let's take a closer look at the data here. Let's say your buddy, Joe Investor, had invested one dollar in that median venture fund (let's call it the Rose Colored Glasses Fund, or RCGF). What would that 11% IRR mean in terms of actual returns, measured the old-fashioned way, in dollars?

As it turns out, the DPI (Distributions to Paid in Capital: i.e. the cash-on-cash return) of RCGF over the ten year period was \$0.78. The TPVI (Total Value to Paid-in Capital) – the DPI plus the (assumed) value of the remaining undistributed portfolio of investments was \$1.53.

Hmmmm. There seems to be some sort of financial slight-of-hand going on here. How could poor Joe Investor make an investment of \$1 in RCGF and, a decade later, a decade in which RCGF reports an 11% IRR on Joe's investment, have just \$0.78 of his \$1.00 back, plus a claim on another \$0.75 of illiquid (and still speculative) assets?

Welcome to the world of venture capital performance reporting (to be fair, all kinds of investment managers play the same sort of game, but still). A world where returns are often as notional as the businesses that generate them.

There's an old economist joke (my undergraduate major) about a physicist, a chemist, and an economist tasked with opening a can of food on a desert island. The physicist (smash the can with a rock) and the chemist (heat the can on the fire until it explodes) propose solutions that get the job done but ruin the food. The economist has a better idea: "first," he proposes, "let's assume that we have a can opener."

VCs reporting IRR results make several key assumptions. Three, in fact, each of which makes for a higher reported IRR, and each of which is, like the economist's can opener assumption, problematic.

First Assumption: The Calculation Doesn't Start Until the Check Clears. Back in the old days, at the birth of the modern venture capital business, when a "Limited" (the generic term for an investor in a venture fund) made an investment of \$X in a venture fund, the Limited sent the fund a check for \$X. No more. Today, it's not actually an investment in the fund, it's a commitment to make an investment (usually with a modest down payment), with the vast bulk of the money only sent to the fund when the fund "calls" for it. Depending on how hot the venture market is, commitments are usually called over several years, in increments that are commonly as small as 5% and seldom more than 25%. The idea is to "call" the capital as needed for investments (or management fees). While there is undeniably a certain logic to that, there is also a problem; a problem that may explain the idea better than the logic. By calling capital only as needed, the fund sets up a situation where it can use the call date for every dollar of capital as the start date for the IRR calculation.

The theoretical underpinnings for calculating the IRR from the call date rather than the commitment date are not without merit. The fund manager's business is picking winners, and the best measure of her/his prowess is the performance of her/his picks from the date of the pick. On the other hand, if I am a Limited, my natural point of reference, performance-wise, is probably the date I became legally responsible for committing the capital.

You can argue either perspective here: both the manager and the Limited can make a decent argument for their preferred "starting date" for the Limited's investment. What you can't deny, however, is that the manager's preference will (discounting the case where the manager is exceptionally bad at picking winners) always lead to a higher reported IRR than the Limited's preference.

Second Assumption: Interim Distributions Continue Growing at Reported IRR. There is an assumption built into the IRR algorithm that cash distributions made within the period under consideration (say for example in year three of a ten year period) continue to earn the same IRR for the rest of the time to the end of the period. Voila, an artifact peculiar to the IRR algorithm that (surprise) tends to boost reported IRRs for higher risk/reward investment vehicles.

Third Assumption: A Deal in the Bush is Worth a Deal in the Hand. Finally, returning to our example, it's interesting that only about one-half of the total returns upon which the IRR figure of 11% is calculated represent actual distributions of capital to the Limiteds. The other half of the "returns" credited to the

Limiteds for purposes of calculating IRR are held by the fund on behalf of the Limiteds, and still speculative (they could increase in value or decrease in value before they are distributed, and may never be distributed, or never be distributed in a liquid form).

If this seems like a minor point, consider that the example median fund has been around for a decade, and has yet to distribute to its Limiteds returns even equal to what those Limiteds have invested. Even as the fund reports that it has generated an 11% IRR for its Limiteds, the Limiteds are not even whole on their investment in terms of actual cash returns.

It's All About Spin. Far be it for me to suggest that the venture capital industry and its private equity brethren might engage in any slight-of-hand in their marketing efforts. IRR is (as noted above and elsewhere blogged about) by its nature a somewhat opaque algorithm, even as it remains popular with the calculator set. When the opacity works in favor of an industry, you'd expect the industry to run with it.

As for me, though, my preference remains as it always has: if you want to see how good a venture investor really is at picking winners, look at the cash-on-cash returns, not the IRR. And, pretty much regardless of the reported IRR, if after ten years of fund life the Limiteds have received back, in spendable cash equivalents, three times what they invested, they have done ok. Certainly better than the median fund in the example Pitchbook data, which is yet to make its Limiteds whole even as it reports a ten year IRR of 11%.

Related People

Paul Jones

Of Counsel

pajones@michaelbest.com

T 608.283.0125

Paul Jones

Of Counsel

pajones@michaelbest.com

T 608.283.0125