

Limitations of Investment Metrics, Part 1

David Working

IRR often does not mean what people think it does.

It's a deceptively simple question: did this investment perform well?

The answer? *It depends.*

Metrics that provide insight into return on investment consider three basic concepts: 1) how much was invested, 2) how much was paid back, and 3) the timing of each.

Unfortunately, it is our experience that the wide variety of investment metrics available to investors, and the complexity of their definitions, have resulted in common misapplications.

But because investors differ in their view of relative importance of these concepts, they use a wide variety of calculations and metrics to deliver an answer to the opening prompt. Unfortunately, it is our experience that the wide variety of investment metrics available to investors, and the complexity of their definitions, have resulted in

common misapplications. There are many scenarios and situations and not every metric is always appropriate.

In this first article of a two-part series, we'll break down some of the most common investment return metrics, and highlight how they can commonly be misused; in part two, we'll suggest alternatives and the underlying rationale.

Acronym Salad: MOIC, CAGR, NPV, & IRR

We'll start off this discussion by defining four lynchpin metrics used for analysis in the investment world: the "multiple of invested capital" (or "MOIC"), compound annual growth rate (or "CAGR"), net present value (or "NPV"), and internal rate of return (or "IRR").

MOIC is the simplest: it compares all returned dollars to all invested dollars as a ratio. For example, an investment that requires \$1 million in upfront investment and over its lifetime returns a total of \$2 million has an MOIC of "2.0x." This calculation eliminates any effects of timing, so in the eyes of MOIC, an investment that returns cash over two years or ten is no different.



$$MOIC = \frac{\sum \text{Returned Dollars}}{\sum \text{Invested Dollars}}$$

CAGR introduces an element of return over investment length, communicated as a percentage. It compares a beginning value to an ending value, and determines a steady annual growth percentage at which the investment would have needed to grow to arrive at the ending value. This is written as an equation below (the variable "T" represents the time period of the investment):

$$CAGR = \left(\left(\frac{\text{Ending Value}}{\text{Beginning Value}} \right)^{\frac{1}{T}} - 1 \right) * 100$$

CAGR provides a smoothing effect for lumpy activities. This can be useful to cut through the noise and see performance over time, as many investments are "two steps forward, one step back" experiences, and CAGR can help make sense of that volatility. However, because of that smoothing, CAGR doesn't value consistency: the difference between an investment that makes its entire return in one fantastic year out of five and an

investment that performs the same way every year can have the same CAGR.

CAGR also struggles to account for inflows and outflows of cash and is a better fit for tracking changes in value. For this reason, CAGR is rarely used as a metric used by a private equity investor to evaluate a prospective investment, but it could be very useful to a family office or a corporate conglomerate looking at the change in value of a subsidiary over time.

NPV introduces the concept of timing of individual cash flows. NPV is the present dollar value of all future cash flows of an investment, both in and out, assuming a discount rate r .

$$NPV = \sum_{T=0}^T \frac{(Cash\ Flows_T)}{(1+r)^t}$$

NPV is useful as it introduces the effect of the time value of money, that earlier returns are more valuable than later returns. Because of that effect, two investments with identical MOICs or CAGRs can have very different NPVs, if one returns some of its dollars earlier than the other.

The best way to envision IRR is as the discount rate that would have rendered the NPV of an investment's cash flows to be zero.

However, NPV fails to relate the present value in dollars to the overall size of the investment. It's impossible to know if an NPV of \$1 million is "good" or "bad" without knowing the overall size of the investment, which requires pairing NPV with other



metrics to drive any useful conclusions. Additionally, the use of NPV relies entirely on the selection of an appropriate discount rate, a choice that is fraught with its own challenges.

IRR, last but not least in our collection of acronyms, offers a solution to NPV's problems. The best way to envision IRR is as the discount rate that would have rendered the NPV of an investment's cash flows to be zero. It's the same math as determining NPV, but changing the variable—NPV sets a time value of money to determine a present dollar value, while IRR sets the present dollar value to zero to determine a discount rate.

$$0 = NPV = \sum_{T=0}^T \frac{(Cash\ Flows_T)}{(1+IRR)^t}$$

IRR appears on its surface to be the pinnacle of investment metrics: a single percentage value that communicates an investment's performance relative to its size and in consideration of time. But as we'll see, even IRR has its limitations, and is often not used appropriately given its

definition. But enough theory—let's move on to more practical (-ish) examples.

There's Always Money in the Banana Stand

Consider a simple investment in a banana stand. You spend \$10 million building the high-speed banana infrastructure in year 0 and make \$2 million each following year from operations. At the end of year 5, you've returned a total of \$10 million and recovered your initial investment, so MOIC will equal 1.0x and both the CAGR and the IRR will equal 0%.

Imagine that at the outset, you know that KKR will be interested in buying the banana stand in a national rollup venture, and you'll be able to sell your interests at the end of year 5 for \$20 million. This is where CAGR and IRR diverge: this assumption boosts the expected IRR to 31% and the CAGR to 25%. Blackstone, however, offers to buy you out a year earlier in year four, and is so excited to own the stand they offer you \$20 million for the buyout, and a \$2 million earnout in year 5 in addition. Both offers

result in the same total return of dollars—\$30 million—but the timing is different. Intuitively, you recognize that the MOIC (now up to 3.0x!) doesn't change between the KKR and Blackstone scenarios, but you'd like to calculate the other metrics.

To CAGR, the timing difference in the cash flows between the two scenarios doesn't matter. All returns (and costs) get treated the same regardless of timing.

You quickly pull out your abacus to estimate the new CAGR, and generate an even more exciting CAGR of... still 25%?

To CAGR, the timing difference in the cash flows between the two scenarios doesn't matter. All returns (and costs) get treated the same regardless of timing. Your IRR calculation, however, will result in a new figure of 37%, an increase in value purely from Blackstone's buyout occurring a year sooner. At a discount rate of 7%, Blackstone's offer is worth about \$1.0 million more to you in present value than KKR's offer. Time value of money matters, and CAGR or MOIC alone would not have uncovered this value.

In this, and many scenarios, using IRR offers useful insight. But let's look at a scenario that highlights one of the main drawbacks of IRR. The blessing and the curse of CAGR is that it smooths cash flows over an entire period, so if total cash flows remain the same, then the longer the period, the lower the CAGR. IRR, however, only changes as there are additional cash flows to change it.

Back to the banana stand example, let's say instead of selling your interests to either Blackstone or KKR, you instead effect a seemingly brilliant sale-leaseback plan where you sell the land underneath the stand. This move infuriates the private equity titans and puts an extra \$20 million in your pocket at the end of year 5 as originally planned, while you retain full control of the operating business and achieve a five-year IRR of 20%. Unfortunately, you failed to negotiate a long-term lease for the land before close, the subsequent rent hikes demolish your profitability, and you make no income for the next 25 years. This cuts down your five-year IRR of 31% to a 30-year IRR of... 31%.

For IRR to change in subsequent years, there have to be more cash flows; and for IRR to decline, there have to be more negative cash flows (additional investments).

"But that can't be possible, Dave," is what clients usually say when we illustrate this effect. But it's true. After five years, the discount rate that resulted in the NPV of the banana stand's cash flows equal to zero was 31%; adding additional years of zero cash flows has no effect, positive or negative, on that rate. For IRR to change in subsequent years, there have to be more cash flows; and for IRR to decline, there have to be more negative cash flows (additional investments).

This brain-bending effect comes from the muddling of CAGR and IRR in the minds of investors. In the same scenario, CAGR drops every year there's a zero return, as it's annualizing the total unchanging returns over a growing number of periods. We're used to seeing that effect on a percentage and assume IRR works similarly, hence the confusion. It doesn't mean the 30-year IRR for the banana stand is wrong, it just means it might be misleading.

But Not Everything You Know Is Wrong

IRR is still very useful—it's great for evaluating the near-term profile of a prospective investment, and it's also useful to evaluate the performance of an investment in the last 3-5 years, as an evaluation of transaction price, structure, and company performance.

At this point, it should not be surprising why private equity professionals lean so heavily on a metric that's effective in valuing companies over a 3-5 year period with a contemplated exit. On the other hand, there are situations where IRR is either misleading or woefully inadequate:

- IRR tells us about returns relative to investment size, but it doesn't contextualize across investments, so it can be a poor metric for evaluating opportunity costs. Real-world investors might accept a lower IRR from a \$100 million investment than they would from a \$10 million investment, as it allows them to put more capital to work and earn a far larger sum of actual dollars.
- IRR's value of time can work against it for short investment periods. An investment might look fantastic on paper with a high IRR, until you find out

that it's only going to return a 1.1x MOIC because it'll only be invested for two months before being paid back. You can't spend IRR, so the actual dollar return context is important.

- IRR is a poor measure for evaluating recent performance of historical investments. The banana stand is an extreme example, but recency is a common problem for family-owned businesses. Think of an apple orchard, purchased relatively cheaply 40 years ago, that performed exceptionally well in the 80's and 90's, and has since had difficulty in turning a profit. The IRR for the orchard *looks* excellent, but the original purchase price and early performance is so disconnected from present-day as to render its metrics unusable.
- Finally, IRR doesn't entirely escape the discount rate discussion; a "good" or "bad" IRR requires some understanding of the market and idiosyncratic risk of the target. A 20% IRR over five years might sound great, until I tell you it's from operating an

illegal fireworks factory on the second floor of a steel foundry.

IRR is a very useful tool, especially when evaluating the potential return of investments over a near-term timeframe. However, its position as an industry staple, and its ease of calculation in Excel, has led to the metric being overused.

Until Next Time

The truth is that there isn't a one-size-fits-all investment metric. All of the metrics

discussed in this article are useful and have their place in returns analysis.

And for all our criticisms, IRR is a very useful tool, especially when evaluating the potential return of investments over a near-term timeframe. However, its position as an industry staple, and its ease of calculation in Excel, has led to the metric being overused.

In the next article in the series, we'll investigate the case of the long-held apple orchard, where IRR is a truly poor fit. The best analysis is always going to be asset-unique, and the best metrics aren't going to be solely internal. We'll show how to incorporate market data, including asset-class risk, comparable asset valuations, and ultimately, the market's valuation of the asset, combined with internal metrics to help drive decision-making.

ZacharyScott.

1200 Fifth Avenue, Suite 1500 Seattle, WA 98101

o: 206.224.7380

zacharyscott.com
