SEPTEMBER 2023

What Are the Odds?

Getting a better read on portfolio risk-return metrics

Olivier d'Assier, Head of APAC Applied Research, Axioma

AXIOMA

Introduction

The genesis of this paper comes from the many conversations I, and others, have had with our clients and the obvious need not to miss the forest for the trees. Risk management solutions can deliver a myriad of quantitative reports at any level and users often get bogged down with very detailed statistics about all aspects of their strategy, bypassing the important insight provided by summary metrics on their portfolios.

There are several assumptions behind this paper. The first assumption is that what distinguishes an investor from a gambler is that the former bets only when the odds are in their favor. This obviously introduces the need to measure those odds. Our tools here will be the ex-ante active risk and projected active returns of an optimal portfolio at the time of construction.

The second assumption is that the risk and reward metrics we will be using to define these odds, represent the mean and standard deviation (or variance) of a distribution which, rightly or wrongly, is assumed to be normal. The validity of this framework, therefore, relies on the investor's ability to distinguish between normal and non-normal distributed returns for the investment instrument (portfolio) in a particular market environment before deciding whether or when to rely on this analysis.

The third assumption is that we agree on the ex-ante Information Ratio (IR) as an acceptable measure of the average risk-reward trade-off to be expected from an optimal portfolio over the long-term. An IR of 0.00 means that the probability of this portfolio delivering positive (or negative) active returns is exactly 50% (i.e. the manager has no above-average skills but is also not a reverse indicator). Under the assumption of normally-distributed returns, an IR of 0.5 translates into a 69% probability of outperforming the benchmark (i.e. the manager has more skill than 69% of his or her peers).

This paper will therefore lie at the intersection of quantitative analysis and probability theory. Using top- level ex-ante metrics, we will first provide a framework for classifying the nature of the investment opportunity provided by a portfolio within the context of portfolio selection. Next, we will propose a framework for risk budgeting driven by the market's level of reward for risk taking at different points in time. Given this information, managers will be able to put their portfolios in the context of what is to be expected under a normal market environment and treat any deviating results as a warning sign about either the environment, the optimality of their portfolios, or both.

What is...?

Predicted active risk (a.k.a., ex-ante active risk): Sometimes called the tracking error for strategies that very closely track an index. This is expressed as an annualized standard deviation of the active return.

Information Ratio (IR): The ratio of annualized expected residual returns to residual risk, a central measurement for active management. A manager's value add is proportional to the square of the information ratio.

Downside risk: Aggregates the risk from the losses below a minimum threshold that investors may have. Standard deviation measures of risk only look to compute the risk that is under the bulk of the (normal distribution) curve, but investors also need to consider about the negative returns in the left tail.

Upside potential ratio: Like its downside risk counterpart, this measure looks at the size and probability of returns above a certain investment goal. Investors should only choose investments where the upside potential is larger than the downside with (i.e. risk-adjusted returns).

A Portfolio Risk, Return, and Probabilities Analysis Framework

Consider what the summary active metrics mean in the context of the probability of outperforming the benchmark over the long term. With our portfolio construction tool, Axioma Portfolio Optimizer, we provide investors a Risk Summary report. If you look at a sample report for a hypothetical portfolio (Portfolio A), you can see that it's returned a predicted active risk of 5% and an expected active return of 1.25%. Taking these top-level statistics and assuming normally distributed returns for this portfolio, we generated 10,000 randomly simulated active returns to get the distribution shown in Figure 1 for Portfolio A.

Figure 1. Simulated Distribution of Active Returns for Portfolio A

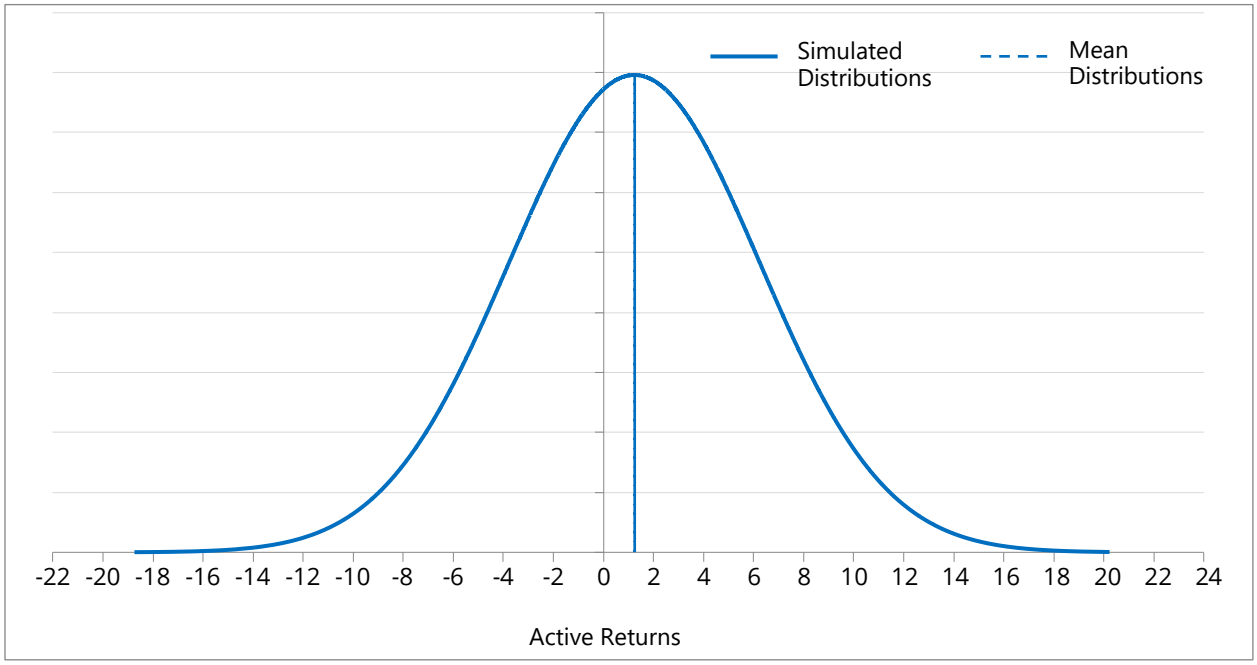

Source: Axioma Portfolio Optimizer

Using the probability distribution of active returns for Portfolio A, we can derive other statistics that may be helpful in our decision to invest in this portfolio or not. From the data in Figure 2 we see that this portfolio has about a 60% chance of outperforming its benchmark over the long-term.[1] The manager has a predicted information ratio of 0.25, putting him or her in the top 60% of their class.

1 We rolled the dice 10,000 times!

Figure 2. Distribution of Active Returns for Portfolio A

Simulation Statistics	Monte Carlo Sim A
Portfolio Predicted Active Return	1.25
Portfolio Predicted Active Risk	5.02
Predicted Information Ratio	0.25
Probability of Out-performing	60%
Probability of Under-performing	0.40
90% CI Upper Bound	7.68
90% CI Lower Bound	-5.18
Average +ve Active Return	4.42
Average -ve Active Return	-3.55
Avg. Gain to Avg Loss Ratio	1.25
Avg. Downside Risk (<10%)	-13.10

Source: Axioma Portfolio Optimizer

Furthermore, we can be 90% confident that the portfolio's active returns will come in somewhere between +7.7% and -5.2%. The ratio of the average positive active return to the average negative one, is 1.25. So, not only do we have a 60% chance of outperforming the benchmark, but when we do, the average positive active returns will be 1.25 times that of the average negative active returns for this portfolio. These look like pretty good odds an investor would be happy with.

But what about the downside risk for Portfolio A? Looking at the average of the active losses in the distribution's left tail corresponding to the first 10% of cumulative probability, we get an average negative return of -13.1% for this portfolio. In other words, if we fall within the bottom 10% of probable returns, we could see ourselves underperforming the benchmark by as much as 13% on average!

What if we had a second portfolio to choose from? Can we use this framework to help us decide between two different investment options? Let's introduce Portfolio B in this process, with a predicted active risk of 3%, an active return of 0.95. Figure 3 shows the active return distribution for both portfolios.

Figure 3. Simulated Distribution of Active Returns for Portfolios A & B

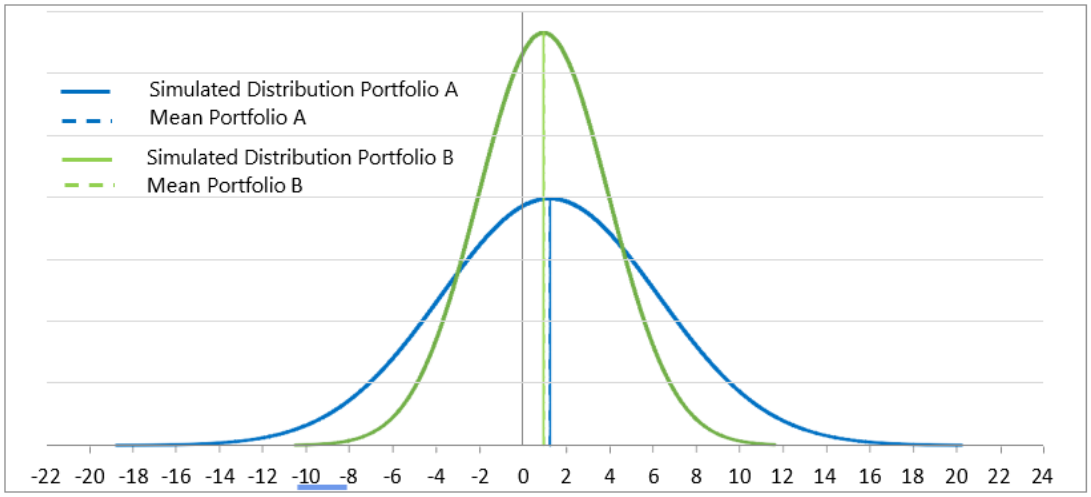

Source: Axioma Portfolio Optimizer

Figure 3 shows the other statistics for both Portfolio A & B side by side. Both have similar statistics when it comes to their probability of outperforming the benchmark, their average positive and negative active returns, and their average active gains-to-losses ratio.

Figure 4. Summary Statistics for Portfolios A & B

Simulation Statistics	Monte Carlo Sim A	Monte Carlo Sim B
Portfolio Predicted Active Return	1.25	0.95
Portfolio Predicted Active Risk	5.02	3.00
Predicted Information Ratio	0.25	0.32
Probability of Outperforming	60%	60%
Probability of Underperforming	0.40	0.40
90% CI Upper Bound	7.68	4.79
90% CI Lower Bound	-5.18	-2.89
Average +ve Active Return	4.42	4.43
Average -ve Active Return	-3.55	-3.50
Avg. Gain to Avg Loss Ratio	1.25	1.26
Avg. Downside Risk (<10%):	-13.10	-7.76

Source: Axioma Portfolio Optimizer

But they differ in two important areas for consideration. Manager B has a slightly better information ratio of 0.32 putting her in the top 63% of her class (and ahead of Manager A). The downside risk as measured by the average bottom 10% of active losses is just 7.76% for Portfolio B versus the 13% for Portfolio A.

How do we choose between those two portfolios? Which is the best investment opportunity?

- If we focus purely on active returns (the predicted active returns of one versus the other), then Portfolio A is the better choice. This is confirmed by the higher expected active return for each portfolio (1.25% versus 0.95%).
- If we focus on selecting the manager with the highest skill (higher IR), then Portfolio B seems more attractive (IR of 0.32 versus 0.25).
- If we focus on downside risk protection (minimizing large losses), then again, Portfolio B seems less risky by virtue of having a smaller average potential active loss across our worst-case (bottom 10%) scenarios (-7.8% versus -13%).

What is the Upside?

Another criterion used in the literature to differentiate portfolios looks at their upside potential ratio.[2] This ratio looks at the portfolio's effective active return divided by its downside deviation from a pre- determined minimum acceptable return level. Figure 5 plots the upside potential ratio for both portfolios along a range of minimum acceptable active returns. What we see is that if our minimum acceptable active return is 0.50%, then both portfolios offer the same upside potential. Any active return lower than 0.5 increasingly favors Portfolio B.

Figure 5. Upside Potential Ratio (Sortino Ratio) for Portfolios A & B

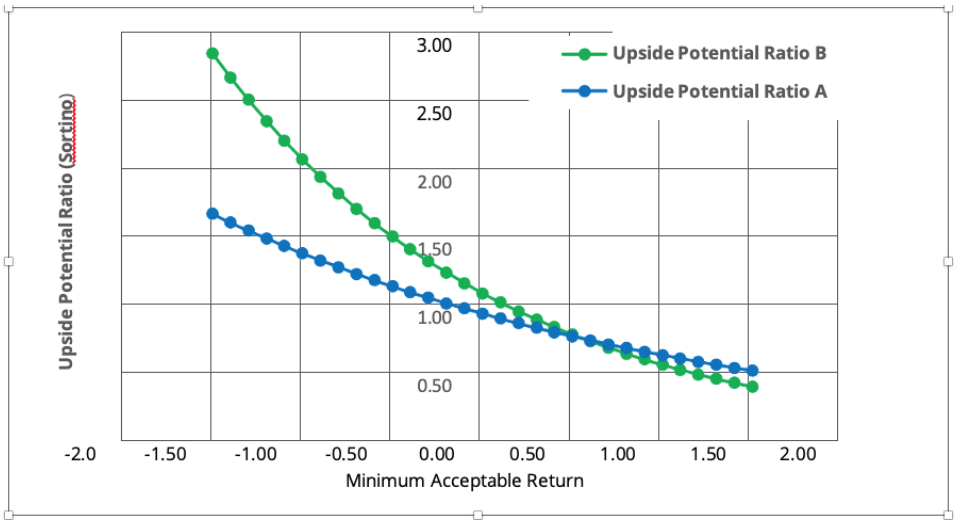

Source: Axioma Portfolio Optimizer

2 Also known as the Sortino ratio.

3 Axioma Portfolio Optimizer's frontier capabilities allow us to compare the results of optimizing portfolios by systematically changing one variable. Using this function, we generated a frontier backtest of 20 portfolios for each month between February 2015 through to September 2017 varying the level of active risk for each portfolio.

This analysis confirms that (portfolio) beauty is indeed in the eyes of the beholder and the choice of the best portfolio depends on the investor's stated investment goals. If your goal is to turn a little money into a lot, then Portfolio A has the highest potential. If you are trying to prevent a lot of money from becoming a little, then Portfolio B is a much more reasonable choice.

Summary metrics simply assist investors in screening portfolios promoted as one-size-fits- all, for the one-that-fits-just-them. Once a shortlist has been established with the help of this process, deeper analysis needs to be done using the full set of drill-down reports provided.

A Portfolio Risk-Timing Framework

Portfolio managers often get a mandate with a targeted level of Active Risk as one of the constraints. But does that mean that the manager should maintain that same level of active risk at each rebalancing and across the entire life of the portfolio, regardless of the risk-reward environment in the market? Is there a framework for optimally selecting the best level of active risk to take based on how the market rewards risk-taking at a particular time?

We have all heard of the phrase risk-on / risk-off, so it should not come as a surprise that the optimal level of risk taking should vary with the market environment. Below we will present a framework for identifying the optimal level of active risk at any point in time and use the portfolio analysis from the previous section to validate our results.

Our tool for this analysis will be the efficient frontier and we will assume a manager with a strong level of forecasting skill (we created a track-record[3] for a manager with a look-ahead bias giving him an expected IR of 1.0) to measure how much the market rewards that (top decile) skill at various points in time. We will then use the metrics from the previous section to differentiate between portfolios along the frontier at each point.

We selected five dates during our history and plotted the efficient frontier for each one in Figure 6. At first glance, we can already see how different their trajectories are to each other at these points in time, despite the manager skill being roughly the same throughout. Just as the Sharpe Ratio of the market changes through time, so do the rewards to risk taking based on changes in volatility, dispersion, and asset-to-asset correlation levels.

Figure 6. Efficient Frontiers at five different dates

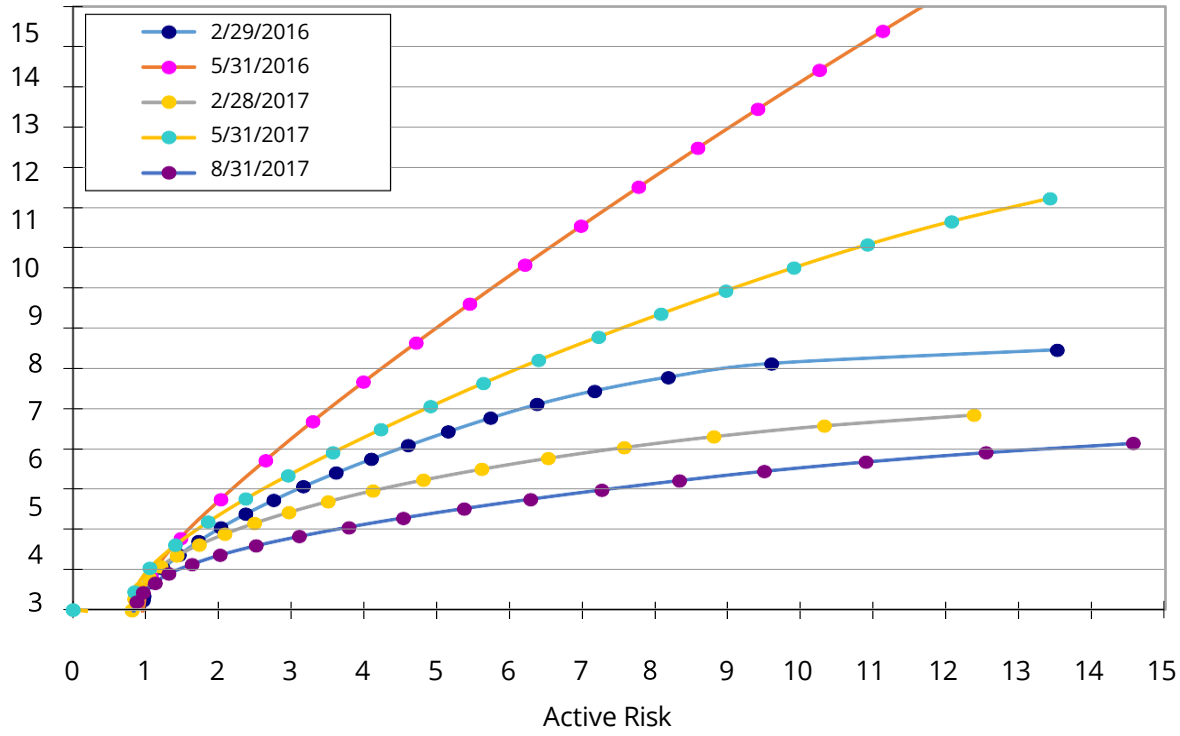

Source: Axioma Portfolio Optimizer

Figure 7 compares the summary statistics for the portfolio with the highest IR in each of the five months (on the left) against a portfolio on the frontier at or near the 3% mandated level (on the right). For each month, we note the predicted active risk and return of the portfolio with the highest information ratio on that date. The sample mean IR reported on the table is close to our full track record of 1.0, but the optimal portfolios vary from an IR of 1.42 in May 2016, to just 0.70 on August 2017, indicating that the market was paying twice as much reward for risk taking on the earlier date.

The optimal level of active risk for those two portfolios on those dates was 3.30 and 1.63 respectively.

The right side of the table shows the same statistics for a portfolio maintained at (or near) the 3% active risk level. In August 2017, taking 3.11% of active risk resulted in an information ratio that was 16% lower than the optimal solution, with an active risk of just 1.63%. In other words, taking twice as much active risk in an environment that does not pay much reward for risk-taking is just wasting '*risk dollars*'.

Figure 7. Summary Statistics on Optimal Portfolios

Historical Efficient Frontiers	Selected Portfolio Date	Optimal Active Risk	Expected Active Return	Highest Information Ratio	At 3% of Active Risk	Projected Active Return	Projected Information Ratio	IR Delta
Month 1	2/29/2016	2.04	2.05	1.01	3.17	3.07	0.97	-4%
Month 2	5/31/2016	3.30	4.70	1.42	3.30	4.70	1.42	0%
Month 3	2/28/2017	1.43	1.35	0.95	2.97	2.43	0.82	-14%
Month 4	5/31/2017	1.86	2.19	1.18	2.95	3.34	1.13	-4%
Month 5	8/31/2017	1.63	1.14	0.70	3.11	1.83	0.59	-16%
Mean	—	2.05	2.29	1.05	3.10	3.07	0.99	-6%
St. Dev.	—	0.74	1.42	0.27	0.15	1.08	0.32	—

Source: Axioma Portfolio Optimizer

Additionally, if we compare the optimal portfolio on August 31, 2017 with the one managed at a constant active risk level of 3% using the framework discussed above, we see that the lower-risk optimal solution has a better overall value proposition. Figures 8-10 mirrors the analysis done on portfolios A & B in the first section of this paper and confirms that purely from an active return potential perspective, the constant active risk portfolio ("3% AR") has a higher probability-weighted return than the optimal portfolio, but both the average expected active loss and the information ratio of this portfolio make it a less sensible choice (see below).

The optimal portfolio ("Optimal"), with a predicted active risk of 1.63% and expected active returns of 1.14% translates into a probability of beating the benchmark over time of 76% versus 72% for the 3% AR portfolio. Here again, the expected returns of the 3% AR portfolio is higher than for the Optimal portfolio (1.83% versus 1.14%), but all other statistics favor the Optimal portfolio, especially the average downside risk (bottom 10% of probabilities) at -7.3% for the 3% AR portfolio versus -3.8% for the Optimal one (see Figure 9).

In addition, when looking at the two portfolios from the point of view of a minimum acceptable active return, we see that the Optimal portfolio becomes far superior as the minimum acceptable active return target becomes negative, but very little separates them when the target is positive (see Figure 10). This suggests that the higher active risk level of the 3% AR portfolio really does not translate into better odds of beating the benchmark but come with much higher odds of under-performing. Put another way, on that date, the market was simply not paying much reward for risk taking and was potentially offering larger active losses to those who over-spent their risk budget.

Managers must take note of the risk-reward environment in the market at the time of each rebalancing and ensure that their portfolio construction settings are in line with what the market pays for risk-taking at the time. Overspending on the risk side when the rewards just aren't there, leaves the portfolio open to downside risk with little or no upside, and summary risk statistics do not adequately tell that story.

Figure 8. Simulated Distribution of Active Returns for 3% and Optimal Portfolios

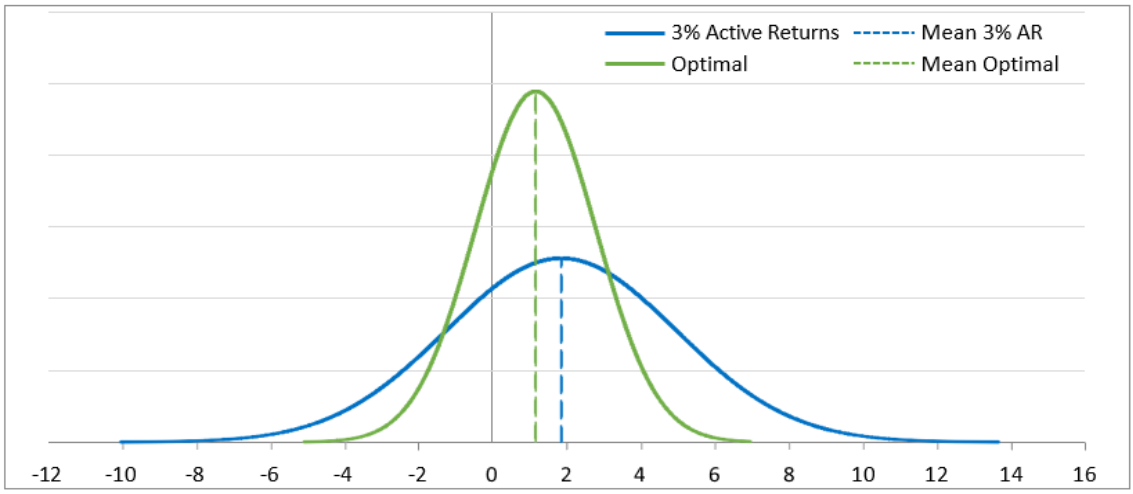

Source: Axioma Portfolio Optimizer

Figure 9. Summary Statistics for 3% AR & Optimal Portfolios

Simulation Statistics	3% AR	Optimal
Portfolio Predicted Active Return	1.83	1.14
Portfolio Predicted Active Risk	3.11	1.63
Predicted Information Ratio	0.59	0.70
Probability of Out-performing	72%	76%
Probability of Under-performing	28%	24%
90% CI Upper Bound	5.82	3.23
90% CI Lower Bound	-2.16	-0.95
Average +ve Active Return	3.26	1.83
Average -ve Active Return:	-1.91	-0.95
Avg. Gain to Avg Loss Ratio:	1.70	1.91
Avg. Downside Risk (<10%):	-7.27	-3.84

Source: Axioma Portfolio Optimizer

Figure 10. Upside Potential Ratio (Sortino Ratio) for 3% AR & Optimal Portfolios

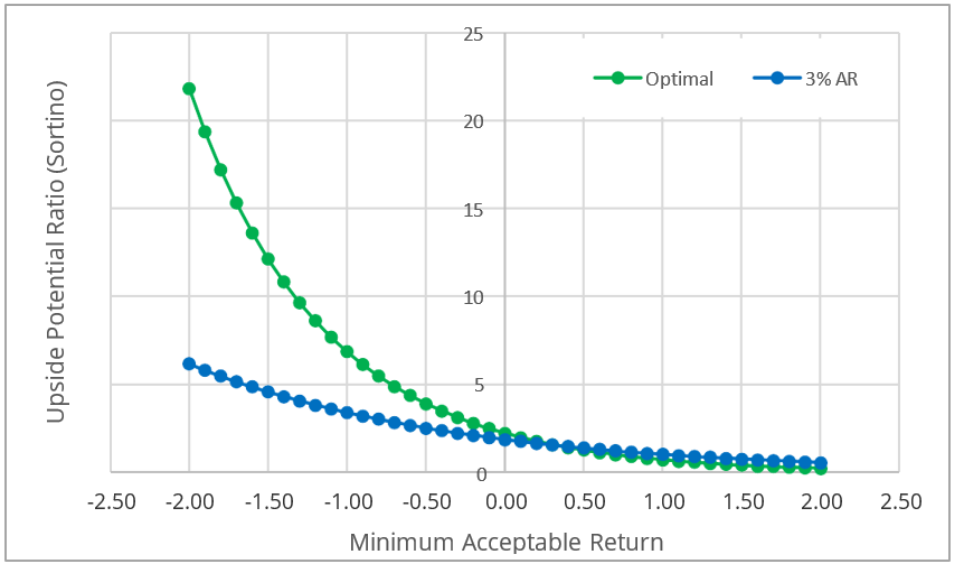

Source: Axioma Portfolio Optimizer

Put together, this analysis tells us that the Optimal portfolio has an upside potential that is almost 12 times larger than its downside risk compared to the constant active risk portfolio's 7.5 times.

Aligning Skill with Mandates

This analysis can also be used in reverse by the portfolio managers during mandate negotiations, to set realistic expectations for investors. Summarizing the track record of our fictional top decile manager, we get the statistics on the left side of the table in Figure 10.

Figure 10. Summary Statistics on full track record and proposed new mandate

Historical Summary	Mean	St. Dev.	New Mandate
Active Risk	2.15	0.50	3.00
Active Return	2.27	0.90	3.50
Information Ratio	1.02	0.19	1.17
Probability of Reaching New Target IR			**44.31%**

Source: Axioma Portfolio

On the right is a proposed new mandate from an investor requiring the manager to target an active risk level of no more than 3% and deliver 3.50% of active return, translating into an expected IR of 1.17.

Should the manager accept this mandate and what are the odds that he or she will be able to deliver these returns and earn a performance fee?

Plotting the return distribution corresponding to the manager's own above-average skills against the new mandate, we see that the target IR is outside the 50% probability of success, and that the manager has only a 44.31% chance of beating the benchmark (see Figure 11). Under these circumstances, the manager should negotiate a new target, giving both him and his client a higher chance of success.

Figure 11. Forecasted Distribution of Information Ratio versus New Mandate Target

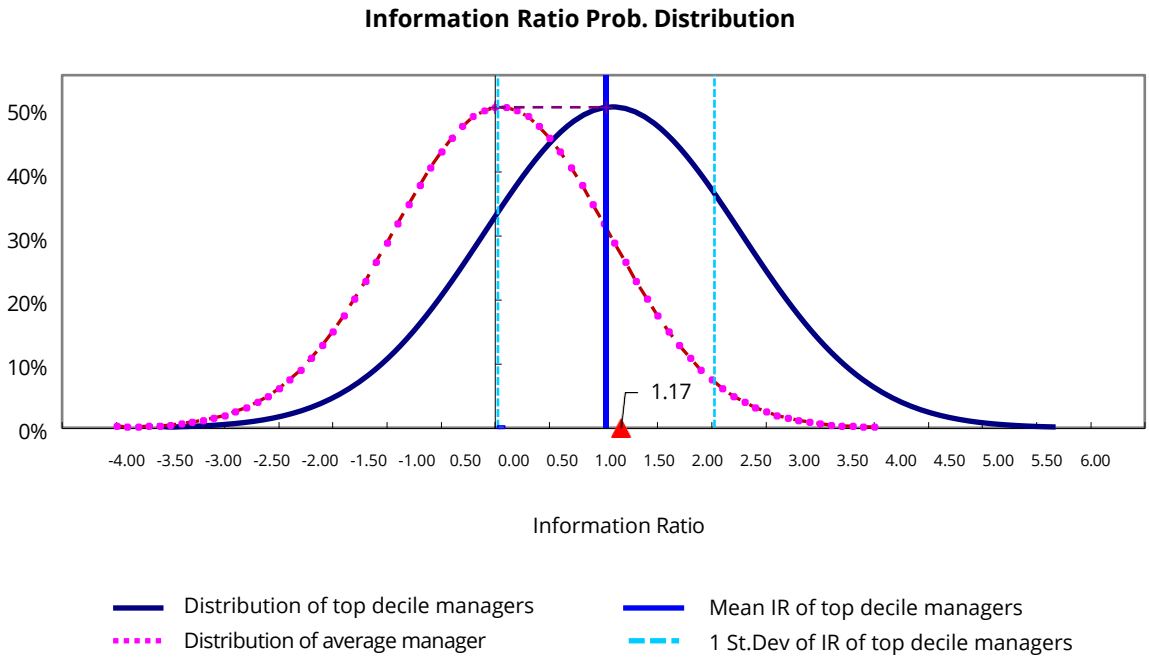

Source: Axioma Portfolio Optimizer

Benefits of Performance Consistency

A manager's skill as reported in the Information Ratio is computed by dividing the active returns of the manager by their volatility. Because returns compound at the rate of time while volatility compounds at the square root of time, as time goes by, a manager's whose skill remains constant will see her odds improve. So, a manager with an Information Ratio of 0.3 has about a 62% probability of outperforming the benchmark in the next period, say one year. For the same level of skill (i.e. in IR of 0.3), but a three-year period, the manager's probability increases to 70%. It increases further to 75% over a five-year period (see table in Figure 12 below). This, of course, assumes the skill level remains constant.

Figure 12. Performance Consistency Benefits through time

Periods	Information Ratios																
(in yrs)	-0.60	-0.50	-0.40	-0.30	-0.20	-0.10	0.00	0.10	0.20	0.30	0.40	0.50	0.60	0.70	0.80	0.90	1.00
0.08	43.1%	44.3%	45.4%	46.5%	47.7%	48.8%	50.0%	51.2%	52.3%	53.5%	54.6%	55.7%	56.9%	58.0%	59.1%	60.2%	61.4%
0.25	38.2%	40.1%	42.1%	44.0%	46.0%	48.0%	50.0%	52.0%	54.0%	56.0%	57.9%	59.9%	61.8%	63.7%	65.5%	67.4%	69.1%
0.50	33.6%	36.2%	38.9%	41.6%	44.4%	47.2%	50.0%	52.8%	55.6%	58.4%	61.1%	63.8%	66.4%	69.0%	71.4%	73.8%	76.0%
1.00	27.4%	30.9%	34.5%	38.2%	42.1%	46.0%	50.0%	54.0%	57.9%	61.8%	65.5%	69.1%	72.6%	75.8%	78.8%	81.6%	84.1%
3.00	14.9%	19.3%	24.4%	30.2%	36.5%	43.1%	50.0%	56.9%	63.5%	69.8%	75.6%	80.7%	85.1%	88.7%	91.7%	94.0%	95.8%
5.00	9.0%	13.2%	18.6%	25.1%	32.7%	41.2%	50.0%	58.8%	67.3%	74.9%	81.4%	86.8%	91.0%	94.1%	96.3%	97.8%	98.7%
10.00	2.9%	5.7%	10.3%	17.1%	26.4%	37.6%	50.0%	62.4%	73.6%	82.9%	89.7%	94.3%	97.1%	98.7%	99.4%	99.8%	99.9%
20.00	0.4%	1.3%	3.7%	9.0%	18.6%	32.7%	50.0%	67.3%	81.4%	91.0%	96.3%	98.7%	99.6%	99.9%	100.0%	100.0%	100.0%

Source: Axioma Portfolio Optimizer

This framework confirms that long-term investors have an edge over short-term ones when it comes to the level of manager skill they must find in order to meet their investment goals. A short-term investor with the same investment goals as a long-term investor but an investment horizon of one quarter (i.e. 0.25 years) needs to find a manager with an IR of 0.6 to have the same probability of outperforming the benchmark. In other words, this short-term investor needs to find a manager twice as skillful as the long- term investor.

Conversely, a manager with an IR of 0.3 only has about a 9% probability of underperforming the benchmark in the next two consecutive periods.[4] Investors can use this framework to monitor their managers and see whether they are performing within these probability expectations. The Information Ratio is often backward-looking and the fund may have had a change of manager or the market environment may have changed from the period reported (i.e. from bull to bear market).

However, investors can use the above framework to plot an expected path for their investments and monitor the realized performance against these probability expectations. If, for example, a manager with an IR of 0.3. is found to underperform in three consecutive periods, an event that should occur less than 3.5% of the time, further investigation needs to take place and an understanding of the reasons behind such a deviation from the expected norm must be achieved before a mandate is renewed with this manager (i.e. maybe the IR was overstated…).

4 Calculated using the negative binomial distribution function, returning the probability that there will be X number of failures before the next success given a specified probability of success.

Flexible portfolio construction

All the metrics discussed in this paper can be extracted from the risk analysis reports in our portfolio construction tool, Axioma Portfolio Optimizer. The preference settings in the application can be customized to reflect an investor's targets with regards to confidence levels or size of expected returns/risk they are willing to take. Investors commonly fail to achieve their investment goals because of one of two things. Either their forecast of where returns will come from was wrong, and realized returns bear little resemblance to predicted returns. The ratio of predicted-to-realized returns is called the information coefficient and can be measured in the Axioma platform. The second reason is that although their prediction of where returns would come from was correct, and the ratio of realized-to-predicted returns is high, they had a sloppy portfolio construction process and did not optimally 'transfer' these predicted sources of returns into their portfolio by aligning the sources of risk with those forecasts.

The ratio of predicted-to-implied returns (calculated from the sources of risk in the portfolio) is called the transfer coefficient and can be extracted from the summary reports in Axioma Portfolio Optimizer. The Information Coefficient (IC) and the Transfer Coefficient (TC), combine a measure of how good the manager's forecasting abilities are, with how optimal their portfolio construction process is, to give investors a measure of the manager's Performance Coefficient (PC) which defines the probability that the manager will be able to deliver the investor's targeted returns.

Conclusion

Investors should quantify their investment goals and define their objectives with regards to both the returns they seek to gain as well as the risks they are willing to take in acquiring them. Portfolio selection should then follow a systematic process of projecting each option's risk-return profile and ensuring the final choice is aligned with the investor's risk tolerance. Selecting a portfolio simply based on its probability of outperforming implies an infinite supply of trials, which isn't the case for most investors. Once the savings run out, the trials stop. Using the risk profile of each option to calculate the upside potential AND the downside risk can help ensure sure the asset owner and portfolio manager are comfortable with both, leading to a greater chance of success. In investment, it is easier to live with the regret of the opportunities you didn't go for, than the remorse from the losses of the ones you did; the latter might be a choice you relive over and over again for the rest of your life.

A manager's skill level remains fairly constant over short periods of time, but the risk environment can rapidly change and the market will pay vastly different rewards for risk-taking at different times.

Managers need to ensure that their risk budget is in tune with the risk environment in the market at each rebalancing. Understanding the nature of the opportunity will also help managers and investors structure investment mandates with a higher probability of success, for both sides, during negotiations.

In the words of Fred Schwed Jr., "*Speculation is an effort, probably unsuccessful, to turn a little money into a lot. Investing is an effort, which should be successful, to prevent a lot of money from becoming a little.*"[5] For the latter to be successful, it requires both discipline and transparency. Headline risk characteristics can be used to design a framework such as the ones discussed above, to provide investors with a kind of confidence intervals around their expected investment outcomes. When reality confounds probability, it's time to ask more probing question.

5 Author of the classic "Where are the Customer's Yachts?" (Wiley, 1940).