MAY 2023

Tracking Error 101: The Intuition Behind Measurement and Control

Hamish Seegopaul, Managing Director, Global Head Index Product Innovation, Qontigo

QONTIGO AXIOMA | DAX | STOXX

Table of Contents

1. Why revisit?

Life, some have said, is the sum of all our choices. In this way, life is very similar to portfolio construction. When we design a benchmarked portfolio, every design choice that takes us away from the benchmark has a consequence, and every consequence has a risk. Tracking error (TE) – the humble statistic that serves as a measure of this risk – is having its day in the sun again due to its use in passive sustainable investing. It has become so prevalent, however, that its meaning can often be obscured. Intuitively, we know that a portfolio with a 1% tracking error is "closer" to its benchmark than a portfolio with a 5% tracking error, but how should we interpret the degree of difference? What does it imply for the future? And can we control this? This article provides a high-level refresher of what tracking error means, and how we can embed it directly into portfolio construction.

2. Necessary trade-offs

Diversification is the only free lunch in finance, though it is a meal that can be known to disappear (sometimes when you need it the most). What we can be certain of is that any deviation from a benchmark – in terms of either components or weights – will lead to the possibility of overperformance or underperformance. While it is impossible to predict with certainty what the result will be, tracking error helps us in ascribing a likely range of possibilities.

In the world of benchmarked[1] sustainable investing in particular, the trade-off between benchmark deviation and risk can be clearly illustrated (as shown in Figure 1). The "greener" a portfolio becomes, the larger the deviation from its benchmark is and the more risk (of over- or underperformance) is incurred.

Figure 1: STOXX ESG indices: ESG scores vs. tracking error.

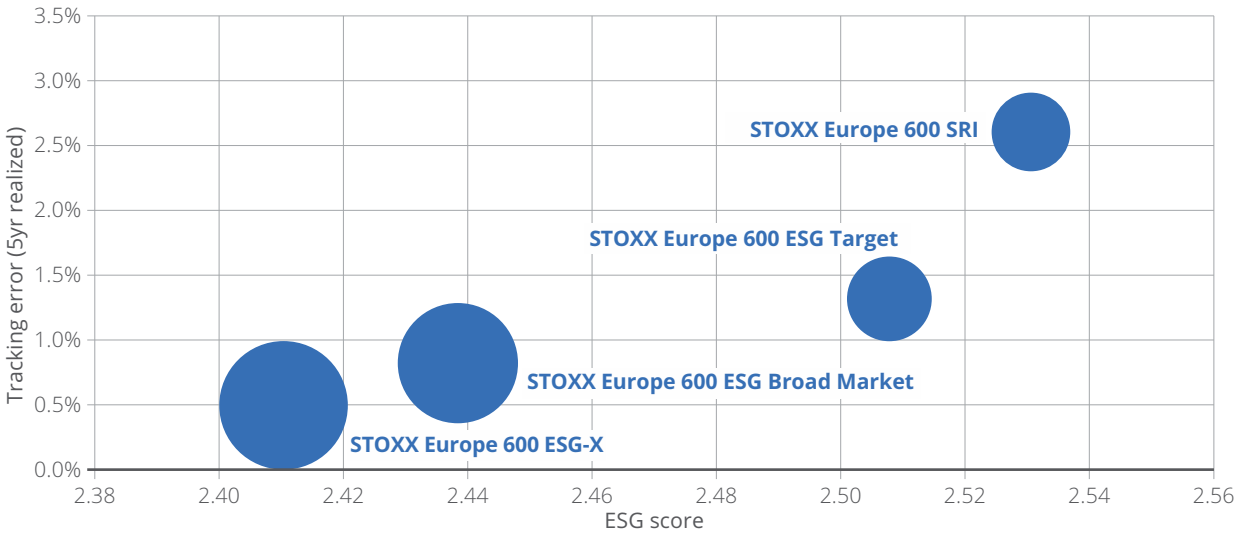

Sources: Qontigo, ISS. Simulated data prior to launch. Average ESG scores and realized tracking errors calculated over a five-year period from April 2018 to April 2023. Realized tracking errors use monthly returns and are relative to the STOXX Europe 600.

[1] It should be noted that benchmarked sustainable investing is completely different to impact investing, where there is often no performance benchmark to speak of. For the purposes of this article, we will use the terms "benchmarked sustainable investing" and "sustainable investing" interchangeably.

Looking at the suite of STOXX Europe 600 ESG indices as an example, this phenomenon of "greening" versus risk can easily be observed by comparing the realized tracking error of each index to its benchmark with its ESG score. While different portfolio construction techniques mean that no strictly linear relationship exists, the trend is clear – the higher the score, the higher the risk compared to the benchmark, as measured by the tracking error. This will be intuitively obvious to many readers, but the question that follows is what, if anything, does this tell us in terms of the future?

3. Deriving meaning

Calculating tracking error is reasonably straightforward with historical returns. Put simply, it measures the degree of difference between the returns generated by the portfolio and those of its benchmark. We have found, however, that there is less discussion about what the results imply.

To visualize what tracking errors mean in practice, we randomly generated portfolio returns over a theoretical 10-year period for differing levels of tracking errors. We assumed that the average outperformance of the portfolio versus its benchmark was 0%, and used a probability distribution to plot relative performances to the benchmark (i.e., the active returns). The graphs below show the simulated portfolio performances for each tracking error scenario.

Figure 2: Randomized active portfolio returns for tracking errors of 1%, 2.5%, 5% and 10%.

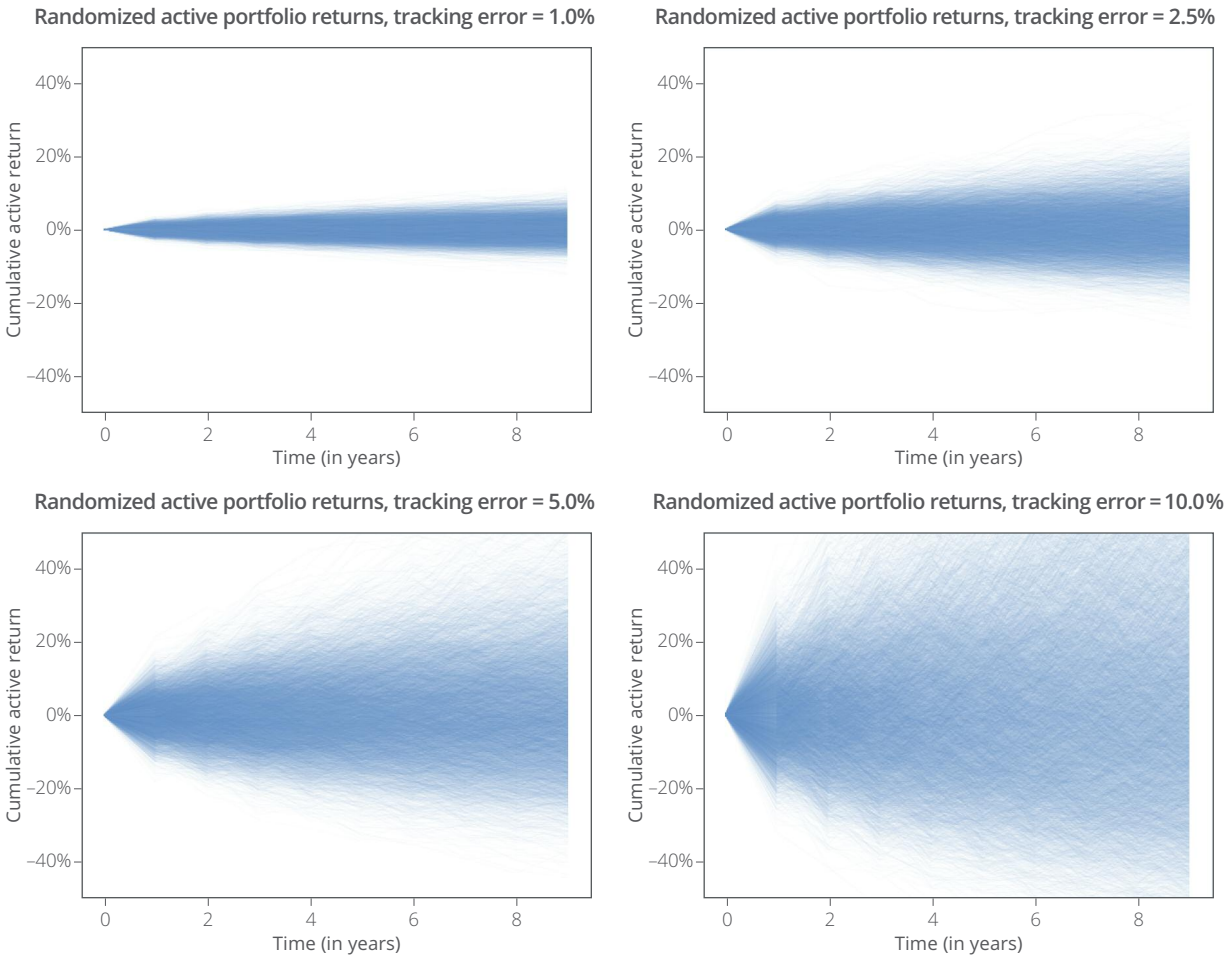

Source: Qontigo. Randomized annual active performances based on normally distributed log returns and zero alpha; 10,000 simulated portfolios per scenario.

As can be seen in the graphics above, varying the tracking error did not (initially) influence the average returns, but rather the range of potential outcomes.

To put this in terms of probabilities, if your portfolio had a tracking error of 1%, you could be reasonably (95%) certain that the relative returns of your portfolio to its benchmark fell in a range of –2% to +2%[2] on an annualized basis. If you had a portfolio with a tracking error of 5%, however, you would be much less (only 31%) certain that the relative returns were in the same –2% to +2% range. To be clear, we are not discovering anything new here, but rather reinforcing the idea that tracking error can be interpreted as a probability distribution.

The number of assumptions needed to arrive at these figures means that they should be taken with a grain of salt (see Section 6 on limitations). Nevertheless, they are powered by basic statistics and specifically by the normal distribution. Tracking error figures are reported as one standard deviation moves relative to the benchmark. In other words, roughly 68% of the values for TE distribution would be +/- one standard deviation, 95% would be within two standard deviations, 99.7% would be within three standard deviations, and so on.

Using these implied probabilities, we can also look back on the theoretical probability that a specific return was exceeded – in statistical terms, this is expressed as the survival function (or 1 – the cumulative density function). The figure below compares 1% and 0.5% TE portfolios. We can observe there was a 16% chance that the active return was greater than or equal to 1% – this is a one standard deviation move. For the 0.5% TE portfolio (which is intuitively more like the benchmark), the tracking error implies that the chance of exceeding a 1% outperformance dropped to just 2%.

Figure 3: Normal distribution implied probabilities of exceeding return levels.

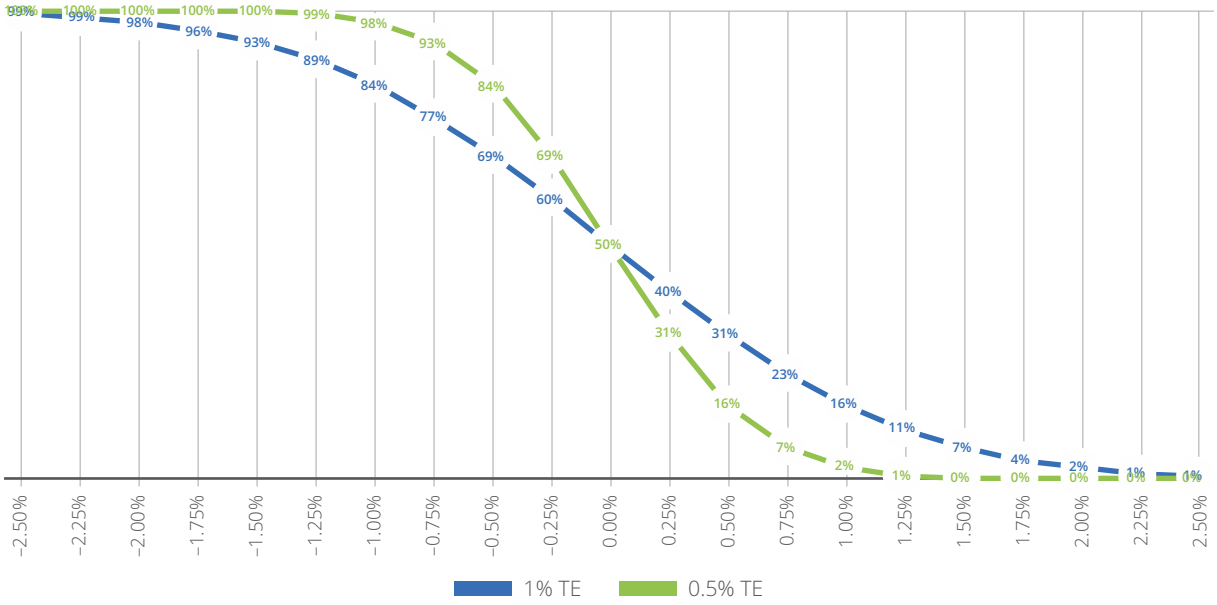

Source: Qontigo.

While the above example focused on an outperformance level, this article will not wade into the debate about whether sustainable investing is a source of alpha. Rather, we would point to sustainability being on a different dimension compared to the typical risk/return matrix. Returns, however, cannot simply

[2] Assuming zero alpha (i.e., a mean active return of 0%) and normally distributed returns.

be ignored. We believe that the increased prevalence of tracking error can partially be explained by one of the most prominent human cognitive biases: loss aversion[3]. Does the underperformance of a sustainable portfolio bring more pain than its overperformance brings joy? Does potential underperformance threaten the continued pursuit of the portfolio's broader objectives? The answer we seem to observe from the investment community is "yes", and we believe that this is partially behind the increased appetite we see to constrain tracking errors in sustainable portfolios.

4. Controlling the future?

No spoiler alert is needed to answer the question of whether we can control the future. The goal of attempting to control tracking errors is not to target a specific return, but, as we have seen, to influence the potential distribution of future returns. It is worth reiterating that if we rely on the normal distribution as our guide, controlling tracking error theoretically limits both our potential upside and our potential downside. Such a move needs to be a conscious choice by investors.

While reported tracking errors are often calculated using historical returns, TE can also be forecasted by calculating the active risk of a set of holdings. Knowing the risk estimates of multiple possible portfolios allows us to select a portfolio that targets a specific level of tracking error while still achieving other objectives.

We can use the following tools to achieve this.

A *risk model* is designed to model the risk of asset returns by reducing the dimensionality of a portfolio with hundreds/thousands of assets to a smaller number of factor loadings. The ex-ante (predicted) risk of the portfolio is then a combination of the portfolio weights, factor exposures and covariance matrices. Reducing the dimensions involved reduces the complexity of modeling large portfolios and can therefore be beneficial. This can be seen from Figure 4 below, which demonstrates the rise in the complexity of calculating risk directly for just two and three assets, respectively.

Figure 4: Calculating active risk using weights.

Portfolio A variance (2 assets) = $w_1^2\,\sigma_1^2 + w_2^2\,\sigma_2^2 + 2\,w_1\,w_2\,\sigma_{12}$
Portfolio B variance (3 assets) = $w_1^2\,\sigma_1^2 + w_2^2\,\sigma_2^2 + w_3^2\,\sigma_3^2 + 2\,w_1\,w_2\,\sigma_{12} + 2\,w_1\,w_3\,\sigma_{13} + 2\,w_2\,w_3\,\sigma_{23}$

w_1 = weight of asset 1; σ_1^2 = variance of asset 1; σ_{12} = covariance of assets 1 and 2

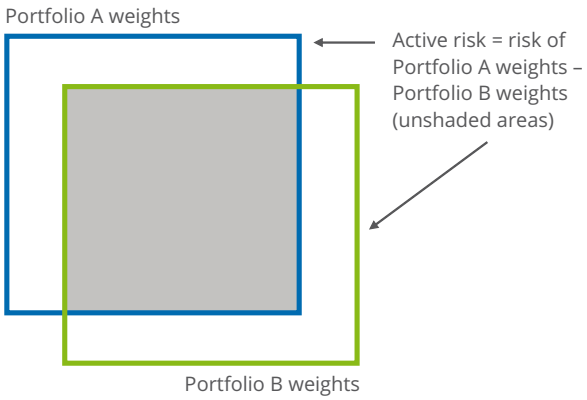

Source: Qontigo.

[3] Coined by Daniel Kahneman and Amos Tversky, encapsulated in their phrase "losses loom larger than gains".

When a portfolio grows to (e.g.) 500 securities, there would be over 125,000 inputs to estimate without a risk model. To return to the introduction of this article, however, nothing is without trade-offs, and performing this dimension reduction intelligently is at the heart of Axioma risk models. To use these models to forecast TE, they can be applied to the difference in weights of a portfolio and its benchmark (i.e. the active weights, also illustrated in Figure 4), which can then give us the active ex-ante risk of a portfolio.

An optimizer is a decision-making tool that can evaluate numerous options to arrive at a solution that best satisfies a given problem. Using Axioma tools and data, we can easily set up problems such as "find the portfolio with the lowest ex-ante tracking error subject to a level of ESG improvement," or "find the portfolio with the maximum ESG improvement subject to a given level of ex-ante tracking error". The Axioma optimizer is agnostic to the datasets used.

5. Real-world example

In their March 2023 paper, "Green efficient frontiers Part 1: Minimizing the risk", our colleagues Melissa Brown and Rob Stubbs illustrated the outcomes that can be achieved by using these tools on an exclusion-only sustainable strategy. They tested three levels of exclusion using a baseline that was defined as:

- No tobacco involvement
- No controversial weapons
- No highly controversial assets, and
- No UNGC non-compliant assets

They then created two portfolios: one that simply reweighted the remaining names ("Exclusions only portfolio") and another that used the Axioma optimizer and an Axioma worldwide risk model to minimize tracking error to the STOXX Developed World Index while disallowing the exclusions ("Optimized portfolio"). Looking at the realized tracking errors for the period from March 2020 to December 2022, they show that the Optimized portfolios have lower realized tracking errors for each scenario.

Figure 5: Realized tracking errors.

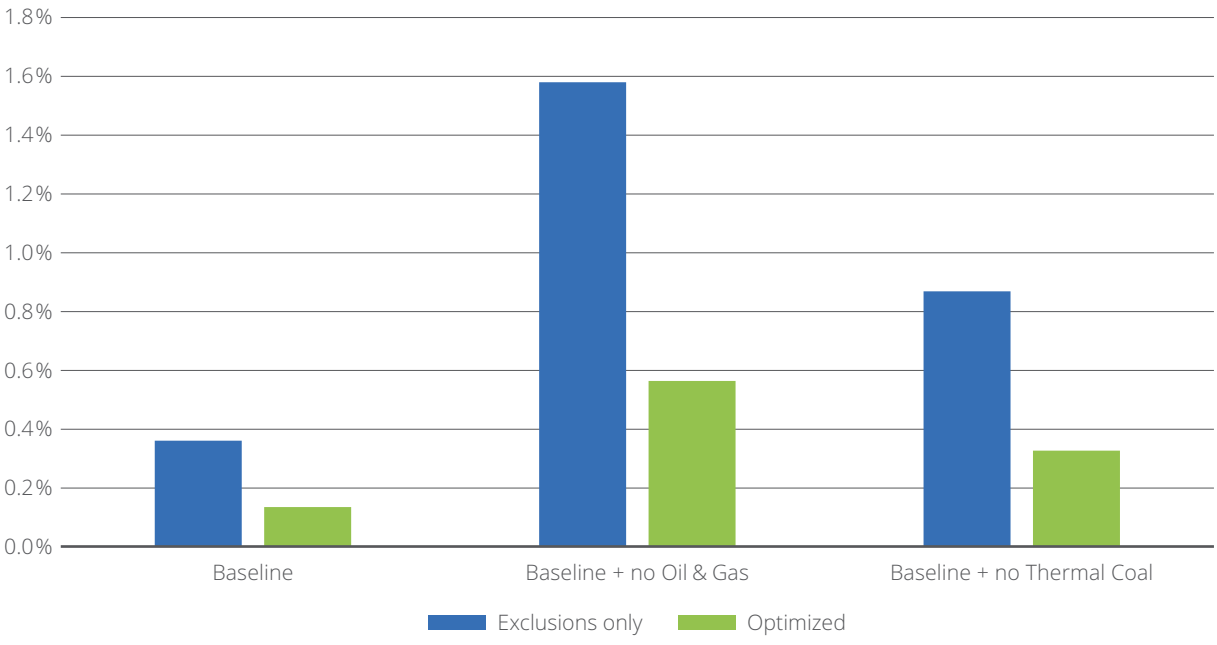

Source: Green efficient frontiers Part 1: Minimizing the risk.

If we think back to the probability distributions and our focus on underperformance risk, we can see that we reduced the (historical) likelihood of underperforming the benchmark for the Baseline + no Oil & Gas scenario by 1% or more from ~26% to ~4%[4].

6. Limitations

Tracking error requires a benchmark to be specified. Benchmarks serve as useful barometers for measuring performance and as building blocks for asset allocation. The advent of sustainable investing has also seen some reinvention of what a policy benchmark could look like, for example, with some European asset owners adding sustainable tilts to broad universes. As alluded to earlier, however, benchmarking (or "benchmark hugging") is not suitable for all investments. For those investments where a benchmark is not suitable, tracking error tells you little, as the risk is on a different dimension.

In the case of those investments for which tracking error can be meaningful, the largest limitation is that this statistic is backward-looking. Numerous conditions need to be met for tracking error to give us a sense of future performance:

- returns need to be normally distributed[5];
- risk estimates, and the assumptions behind them, need to be stable;
- and past performance needs to be indicative of future results (something which cannot be warranted).

These are not assumptions that can be taken for granted, and require tracking error to be used as more of a directional than an absolute guide.

Tracking error is also not the only option investors have for capturing the risk of under- and overperformance. Many measures have been created to more precisely capture this, and many also exist for testing whether a dataset can be assumed to be normally distributed in the first instance. In Figure 6 (page 9), we plotted the historical (and largely backtested) rolling returns for a strategy that utilizes a 1% ex-ante tracking error constraint. The returns appear to be only approximately normally distributed. We cannot say that data is not normally distributed,[6] but nor does the plot look like our clean figures in the previous sections.

Portfolio construction is also, of course, not the only way to limit upside/downside risk, as the popularity of derivatives would clearly indicate. Within the portfolio construction world, however, indirect methods can be used to control tracking error – constraining sector exposures, country exposures, minimum/maximum active weights and other techniques will naturally put a "cap" on the level of deviation from the benchmark. It should be noted that this is not just an ad hoc technique, since these types of exposures are major drivers of active risk — in fact, these factors are implicitly controlled when using the direct optimization method. Although the indirect measures can be more straightforward to implement, they can also lead to less precise control over the (ex-ante) tracking error outcome.

[4] Assuming zero alpha (i.e., a mean active return of 0%) and normally distributed annual standard returns.

[5] Many financial simulations rely on the assumption that log returns are normally distributed, but tracking error is commonly calculated with standard returns. Normal distributions also do not account for skewness and kurtosis commonly found in asset returns.

[6] Statistical tests for normality show a p-value >0.05, failing to reject the hypothesis that the returns are normal. Log returns used.

Figure 6: STOXX Global 1800 ESG Target – Rolling 36m annualized alpha observations.

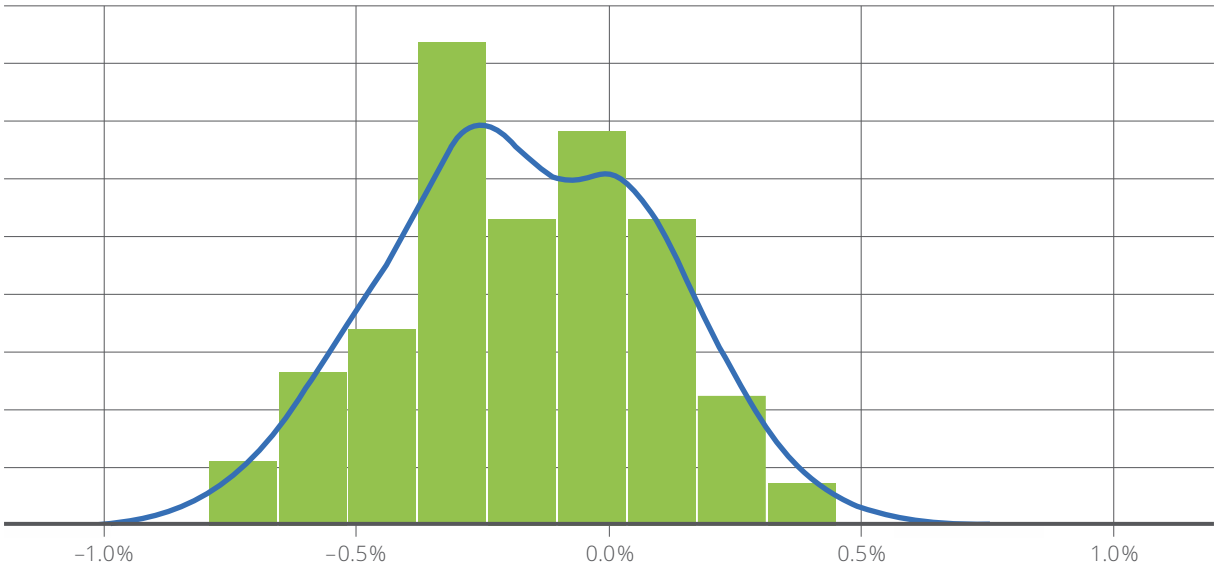

Source: Qontigo. Simulated performance prior to launch. Outperformance vs. STOXX Global 1800 in USD (gross). Monthly observations of log returns for March 2012 to January 2023.

7. Conclusion

Despite its limitations, tracking error is one of the most frequently used statistics in the investment landscape, with broad applications across portfolio management and trading. This article does not offer a new, better, or improved statistic, but rather a reminder of its "roots," how it should be interpreted, and the power of embedding it into portfolio construction. The growth of sustainable investment, combined with the desire to reduce underperformance risk, has created more opportunities for this statistic to help guide the risks investors take.

In this article, we have refreshed the idea that tracking error tells us more than just the "closeness" of a portfolio to its benchmark, and can guide us toward a range of possibilities. There is a danger lurking in this conclusion, though, as it is tempting to use this understanding of TE to expect a singular future outcome. Whether they are deploying this statistic to analyze performance or for portfolio construction, we encourage users of tracking errors to resist this urge, and think probabilistically.

8. Contacts and further information

Learn more about how Qontigo can help you better manage risk and enhance your investment process.
Qontigo.com

Europe

Frankfurt
Mergenthalerallee 61
65760 Eschborn, Germany
+49 69 2 11 0

London
8 Old Jewry
4th Floor
London EC2R 8DN, United Kingdom
+44 20 7862 7680

Paris
19 Boulevard Malesherbes
75008, Paris, France
+33 1 55 27 38 38

Prague
Futurama Business Park Building E
Sokolovska 662/136e
186 00 Prague 8, Czech Republic
+420 228 889 234

Zug
Theilerstrasse 1A
6300 Zug, Switzerland
+41 43 430 71 60

Americas

Atlanta
400 Northridge Road, Suite 550
Atlanta, GA 30350, USA
+1 678 672 5400

Buenos Aires
Corrientes Avenue 800, 33rd Floor
Office 101
Buenos Aires C1043AAU, Argentina
+54 11 5983 0320

Chicago
20 N. Upper Wacker Drive
10th Floor
Chicago, IL 60606, USA
+1 224 324 4279

New York
17 State Street, Suite 2700
New York, NY 10004, USA
+1 212 991 4500

San Francisco
201 Mission Street, Suite #2150
San Francisco, CA 94105, USA
+1 415 614 4170

Asia Pacific

Hong Kong
28/F LHT Tower
31 Queen's Road Central
Hong Kong
+852 3107 8030

Singapore
80 Robinson Road, #02-00
Singapore 068898, Singapore
+852 3107 8030

Sydney
139 Macquarie Street, Level 11
Sydney, NSW 2000, Australia
+852 3107 8030

Tokyo
27F Marunouchi Kitaguchi Building,
1-6-5 Marunouchi Chiyoda-ku
Tokyo 100-0005, Japan
+81 3 4578 6688

QONTIGO
AXIOMA | DAX | STOXX

Part of DEUTSCHE BÖRSE GROUP