

ANDREW COLIN

SMOOTHING IN FIA

FLAMETREE TECHNOLOGIES PTY LTD

Why returns do not combine neatly over time

Whether they arise from individual securities, sectors, or portfolio risks, a fundamental property of return contributions is that they compound additively over markets, but geometrically over time. One outcome of this is that, for a portfolio where return is calculated over more than one time period, *cross terms* or *compounding* will distort the results and potentially obscure the true sources of return in a portfolio.

Cross terms

Suppose that a security has a return r_1 over day 1, and return r_2 over day 2. Then its overall return over the two days will be given by $(1 + r_1) \times (1 + r_2) - 1$, or $r_1 + r_2 + r_1r_2$. In other words, its overall return is the arithmetic sum of the two daily returns, plus the cross term r_1r_2 .

Because of this cross term, it is not possible to decompose a return that has been compounded over more than one interval in terms of an arithmetic sum of individual returns. However, it is often highly desirable to adjust returns so that it *appears* possible.

This chapter explains why such an adjustment is not just aesthetically appealing but necessary, particularly for attribution analysis. To this end, we present several algorithms that allow individual return contributions to be adjusted and combined in a self-consistent way so that they combine to the portfolio's known, true performance.

Example 1

Consider a portfolio that holds two securities. The portfolio and security returns are the same over two successive months:

<i>Security</i>	<i>w</i>	<i>r</i>	<i>c</i>
Security 1	40%	20%	8%
Security 2	60%	10%	6%
TOTAL	100%		14%

Table 1: Weights and returns of sample 2-security portfolio over month 1

<i>Security</i>	<i>w</i>	<i>r</i>	<i>c</i>
Security 1	40%	20%	8%
Security 2	60%	10%	6%
TOTAL	100%		14%

Table 2: Weights and returns of sample 2-security portfolio over month 2

Here, w is weight, r is return, and $c = wr$ is performance contribution.

The portfolio's overall return over each single period is calculated by aggregating each security's performance contribution over the portfolio, and then aggregating these returns over time:

$$R = (1 + 14\%) \times (1 + 14\%) - 1 = 29.96\%$$

Can we arrive at the same result by aggregating security returns over time, and then over the portfolio? The weight of each security is unchanged over the two intervals, so at first glimpse this might seem possible.

The aggregated return of Security 1 is

$$r_1 = (1 + 20\%) \times (1 + 20\%) - 1 = 44.0\%$$

The aggregated return of Security 2 is

$$r_2 = (1 + 10\%) \times (1 + 10\%) - 1 = 21.0\%$$

The aggregate return over the whole interval is then $(40\% \times 44\%) + (60\% \times 21\%) = 30.2\%$, which is close, but not identical, to the correct result of 29.96%.

Example 2

Period	R_B	r_{AA}	r_{SS}	Total
Period 1	10%	10%	5%	25%
Period 2	8%	4%	-6%	6%
Total	18.80%	14.40%	-1.30%	32.50%

Table 3: Attribution returns of simple equity portfolio

Consider the attribution returns of the simple equity portfolio shown in Table 3. As before, the total return of the portfolio (32.50%) is not equal to the sum of the individual compounded returns from the benchmark, asset al-

location and stock selection ($18.80\% + 14.40\% - 1.30\% = 31.90\%$).

Another way to account for this discrepancy is to recognise that addition and multiplication are not commutative operations. Returns are combined additively across a portfolio, but multiplicatively over time. Since the return of each security aggregates at a quite different rate from that of the portfolio, there is no reason to expect that they should aggregate to anything like the portfolio's total return. The calculation must be carried out in one particular order to reach the correct overall return.¹

¹ This also ignores the fact that the market weight of each security will probably vary over the calculation interval.

The importance of internally consistent return contributions

Given that we know the individual weights and returns of each security, we can calculate the overall return of the portfolio exactly. What do we lose by not being able to replicate this return by calculating returns in a different order?

From a pure performance measurement perspective, little useful information is lost. It may be useful to see which sectors or securities have contributed the most to the portfolio's performance, but it is seldom critical.

However, matters become different when one has a requirement to compare the returns made by different sources of *risk* in the portfolio - in other words, to run an attribution analysis. Example 2 showed that the sum of returns on individual risks need not aggregate to the overall outperformance when compounded over multiple periods.

To take an extreme example: suppose we are considering a set of returns that have been compounded multiple times. If the overall return due to asset allocation was 5 bp, and that from stock selection was 10 bp, but the portfolio's cumulative outperformance was 20 bp, how much did each risk contribute to the total? If the numbers do not match up, it is impossible to say, and this makes the value of the analysis moot.

For attribution analysis, where the active return of a portfolio is decomposed by source of risk, it is therefore highly desirable that the active return contributions from

each source of risk be measured in ways that ensure they aggregate to the total outperformance. Ideally, one should be able to drill down to the individual security level and identify the contribution to return from individual sectors, securities or risks, knowing that return contributions at lower levels will combine to aggregate returns at higher levels.

What is needed is a way to adjust the return contributions of a portfolio so that its overall return can be calculated either by aggregating all security returns by sector and then by date, or by date and then by sector.

Path-independence

As long as the contributions from risk remain relatively unchanged, and their aggregated contributions combine to the overall return of the portfolio at each sample date, we have met the main requirements for attribution; a clear view of the sources of outperformance in terms of the underlying risk factors. It also makes reporting much more straightforward, but that is another issue.

In order to make these contributions act in this way, some sort of rescaling of performance contributions is required. Generally speaking, changing weights for reporting purposes is a bad idea, as this can distort the portfolio's exposures. Rescaling is therefore usually performed on returns or performance contributions, while ensuring that overall portfolio performance remains unchanged.

When smoothing, the first decision should be whether the adjusted returns of contributions r_1 and r_2 should combine *arithmetically* or *geometrically* to the total portfolio return R .

- Arithmetic: numbers add up, but compounding over time is lost:

$$R = r_1 + r_2 \quad (1)$$

- Geometric: compounding over time is retained, but additivity of returns over individual samples is lost:

$$R = (1 + r_1)(1 + r_2) - 1 \quad (2)$$

In both cases, we rescale all contributions so that they aggregate to the correct totals, irrespective of the order in which the return contributions were combined - whether by sector, by source of risk, or by time. In other words, they become *path-independent*.

For the numbers in a performance or attribution report to act in a self-consistent manner, the values must be adjusted so that they combine to the correct totals. The trick is to ensure that the adjustment is not too obvious, and that it does not distort the results so badly that useful information is lost. In particular, security or sector-level performance contributions should have the same sign as they did before the rescaling, and the relative magnitude of these contributions should stay as constant as possible. However, there is no right or wrong way to perform smoothing; it is a matter of preference (or prejudice).²

Many algorithms have been published to perform smoothing. In the author's view, it is more important to understand the reasons *why* smoothing is applied than the specific details of each technique. ? provides useful summaries of many widely-used algorithms.

In the next two sections we will describe two of the most widely used smoothing algorithms, and demonstrate their use on the same data set. They are usually presented in terms of their impact on performance contribution rather than raw return, for the reasons outlined in Chapter 2.

Carino smoothing

Consider the following single-currency portfolio and benchmark, which have weights and returns supplied over four intervals Q_1, Q_2, Q_3, Q_4 :

Q_1	w	W	r	R
Sector 1	30%	10%	-20%	0%
Sector 2	10%	20%	20%	20%
Sector 3	60%	70%	-20%	20%
TOTALS	100%	100%	-16%	18%

² This may seem a rather blunt statement when applied to a rigorous, quantitative discipline as performance measurement. However, smoothing remains a critical requirement for useful and unambiguous attribution results, and it is better to understand why it is needed than to sweep its existence under the carpet.

Table 4: Weights and returns for Q_1

Q_2	w	W	r	R
Sector 1	40%	10%	10%	30%
Sector 2	40%	40%	-30%	0%
Sector 3	20%	50%	20%	0%
TOTALS	100%	100%	-4%	3%

Table 5: Weights and returns for Q2

Q_3	w	W	r	R
Sector 1	20%	30%	20%	-20%
Sector 2	10%	20%	-20%	-20%
Sector 3	70%	50%	30%	-20%
TOTALS	100%	100%	23%	-20%

Table 6: Weights and returns for Q3

Q_4	w	W	r	R
Sector 1	20%	40%	20%	20%
Sector 2	20%	20%	30%	30%
Sector 3	60%	40%	10%	0%
TOTALS	100%	100%	16%	14%

Table 7: Weights and returns for Q4

For convenience, the performance contributions (product of weight and return for each sector) from portfolio and benchmark are given in Tables 8 and 9:

c_P	Q1	Q2	Q3	Q4
Sector 1	-6%	4%	4%	4%
Sector 2	2%	-12%	-2%	6%
Sector 3	-12%	4%	21%	6%
Total	-16%	-4%	23%	16%

Table 8: Performance contributions from portfolio

c_B	Q1	Q2	Q3	Q4
Sector 1	0%	3%	-6%	8%
Sector 2	4%	0%	-4%	6%
Sector 3	14%	0%	-10%	0%
Total	18%	3%	-20%	14%

Table 9: Performance contributions from benchmark

The aggregated performance of the portfolio over the four intervals is $(1 - 16\%) \times (1 - 4\%) \times (1 + 23\%) \times (1 + 16\%) - 1 = 15.0572\%$. The aggregated performance of the benchmark is $(1 + 18\%) \times (1 + 3\%) \times (1 + 14\%) \times (1 + 16\%) - 1 = 10.8445\%$. Net outperformance is therefore $15.0572\% - 10.8445\% = 4.2127\%$.

The Carino smoothing algorithm³ is designed to work on a portfolio/benchmark pair, although it is easily adapted for use on a single portfolio, as we show below. Carino's approach introduces a period-dependent factor k , given by

³ Carino, D., *Combining attribution effects over time*, Journal of Portfolio Management, pages 5-14, Summer 1999

$$k = \begin{cases} \frac{\log_e(1 + r^P) - \log_e(1 + r^B)}{r^P - r^B} & \text{if } r^P \neq r^B \\ 1 & \text{if } r^P = r^B \end{cases} \quad (3)$$

where r^P and r^B are the aggregated portfolio and benchmark returns over the entire interval, and a factor k_t for each interval t , given by

$$k_t = \begin{cases} \frac{\log_e(1 + r_t^P) - \log_e(1 + r_t^B)}{r_t^P - r_t^B} & \text{if } r_t^P \neq r_t^B \\ 1 & \text{if } r_t^P = r_t^B \end{cases} \quad (4)$$

where r_t^P and r_t^B are the portfolio and benchmark return over each interval t .⁴ For instance, here

⁴ If you are calculating these quantities in Excel, remember to use the *LN* (natural logarithm, or log to base e) function, rather than the *LOG* (log to base 10) function.

$$k = \frac{\log_e(1 + 15.0572\%) - \log_e(1 + 10.8445\%)}{15.0572\% - 10.8445\%} = 0.885444$$

and

$$k_1 = \frac{\log_e(1 - 16\%) - \log_e(1 + 18\%)}{16\% - 18\%} = 0.999611.$$

To apply Carino smoothing, multiply all performance contributions at time t by k_t/k . The sum of all smoothed performance contributions over all intervals will then give the overall compounded portfolio return over the same period.

For instance, at interval Q1, $k_1/k = 0.999611/0.885444 = 1.128938$. Multiplying every entry in Table 8 by this quantity and adding gives a smoothed return for this period of -18.0630%.

c_P	Q1
Sector 1	-6.7736%
Sector 2	2.2579%
Sector 3	-13.5473%
Total	-18.0630%

Table 10: Carino smoothing applied to Q1

The correction can be applied to per-period portfolio returns, as well as to individual performance contributions, to give the following corrected returns:

Interval	k	r^P	r^B
Q1	0.999611	-18.0630%	20.3209%
Q2	1.005440	-4.5421%	3.4066%
Q3	1.000367	25.9852%	-22.5958%
Q4	0.869587	15.7135%	13.7493%
Total	0.885444	19.0936%	14.8809%

Table 11: Carino factors and smoothed returns

After applying the correction to the aggregated returns shown in Tables 4 to 7, we arrive at the figures shown in Table 11. Although they differ from the actual interval-specific returns, and their sum over the four intervals does not equal the actual aggregated return, the difference between the two corrected aggregated returns of $19.0936\% - 14.8809\% = 4.2127\%$ does equal the true active return, as required.

If it is necessary to ensure that the aggregated smoothed portfolio return equals the unsmoothed portfolio return, use the following simpler expressions to smooth the portfolio and benchmark separately:

$$k = \begin{cases} \frac{\log_e(1+r)}{r} & \text{if } r \neq 0 \\ 1 & \text{if } r = 0 \end{cases} \quad (5)$$

where r is the aggregated portfolio return over the entire interval, and the factor k_t for each interval t is given by

$$k_t = \begin{cases} \frac{\log_e(1+r_t)}{r_t} & \text{if } r_t \neq 0 \\ 1 & \text{if } r_t = 0 \end{cases} \quad (6)$$

Note that the smoothing factors are dependent on the interval over which the return is measured. This means that smoothing must be recalculated if, for instance, an ad-hoc report is required over an arbitrary time interval, which requires in turn that the unsmoothed data be stored.

Other arithmetic smoothing algorithms are described in Bacon⁵, which the reader should consult for detailed analyses of each approach. However, one is so important that we treat it separately in the next section.

⁵ Bacon, C., *Practical portfolio performance measurement and attribution*, 2nd edition, Wiley Finance, 2008

Geometric smoothing

Geometric smoothing should be used when there is a requirement for performance contributions to compound multiplicatively, rather than arithmetically. In this approach, performance contributions are rescaled using the following expression:

$$\hat{c}_i = (1 + c_i) \left[\frac{1 + r}{\prod_{i=1}^n (1 + c_i)} \right]^{\frac{|c_i|}{\sum_i |c_i|}} \quad (7)$$

where for an arbitrary time interval, and a portfolio with n securities,

- c_i is performance contribution from security i
- r is overall performance
- $\hat{c}_i$ is the smoothed performance contribution.

The result of applying this transformation is that

$$r = \left[\prod_i (1 + \hat{c}_i) \right] - 1 \quad (8)$$

To illustrate the algorithm, consider the portfolio shown in Q1 in the previous example. Table 12 shows the various intermediate quantities.

$Q1$	w	r	c	$abs(c)$	$1+c$	c
Sector 1	40%	20%	8.00%	8.00%	1.0800	8.0036431%
Sector 2	30%	-5%	-1.50%	1.50%	0.9850	-1.4993770%
Sector 3	30%	6%	1.80%	1.80%	1.0180	1.8007726%
TOTALS	100%	8.30%				

Table 12: Geometrically smoothed performance contributions

The geometrically smoothed performance contributions in the right hand column are given by

$$c_1 = 1.080 \times \left[\frac{1.083}{(1.080 \times 0.985 \times 1.018)} \right]^{\frac{0.080}{0.113}} - 1 = 8.0036431\%$$

$$c_2 = 0.985 \times \left[\frac{1.083}{(1.080 \times 0.985 \times 1.018)} \right]^{\frac{0.015}{0.113}} - 1 = -1.4993770\%$$

$$c_3 = 1.018 \times \left[\frac{1.083}{(1.080 \times 0.985 \times 1.018)} \right]^{\frac{0.018}{0.113}} - 1 = 1.8007726\%$$

Aggregating these performance contributions geometrically gives

$$(1 + 8.0036431\%) \times (1 - 1.4993770\%) \times (1 + 1.8007726\%) - 1 = 8.3000\%$$

as required.

Repeating the calculation for the successive intervals give the following tables of smoothed contribution returns. The right-hand column is the aggregate return for each quarter, and the bottom row is the aggregate return for each sector.

c_P	Q1	Q2	Q3	Q4	Total
Sector 1	-6.1253%	4.1784%	3.9596%	3.8094%	5.5425%
Sector 2	1.9546%	-11.5463%	-2.0190%	5.7087%	-6.5939%
Sector 3	-12.2345%	4.1784%	20.7533%	5.7087%	16.7108%
Total	-16.0000%	-4.0000%	23.0000%	16.0000%	15.0572%

Table 13: Geometrically smoothed return contributions for portfolio

c_B	Q1	Q2	Q3	Q4	Total
Sector 1	0.0000%	3.0000%	-6.4245%	7.7410%	3.8438%
Sector 2	3.8906%	0.0000%	-4.2892%	5.8093%	5.2110%
Sector 3	13.5810%	0.0000%	-10.6763%	0.0000%	1.4547%
Total	18.0000%	3.0000%	-20.0000%	14.0000%	10.8445%

Table 14: Geometrically smoothed return contributions for benchmark

Note that the overall return in the bottom right of each table can be calculated by

- combining the aggregated returns over each quarter;
- combining the aggregated returns over each sector;
- combining individual performance contributions over sector *and* time

We have now adjusted the performance contributions so that they aggregate geometrically over market and over time to the known overall portfolio return. The aggregation is path-independent, as required.⁶

For relative, or active, returns, the geometric approach uses

$$r_{active} = \frac{1 + r_P}{1 + r_B} - 1 \quad (9)$$

so in this case the active return is given by the geometric difference

$$r_{active} = \frac{1 + 15.0572\%}{1 + 10.8445\%} - 1 = 3.8005\%$$

rather than the arithmetic difference

⁶ To combine many return contributions c_i geometrically, it is often easier to calculate $S = \sum_i \log_e(1 + c_i)$; then the summed aggregated return will be $e^S - 1$. This applies equally to spreadsheets and databases, where summation functions are usually built-in.

$$r_{active} = 15.0572\% - 10.8445\% = 4.2127\%$$

Because return contributions now combine path-independently, active contributions can also be calculated in this way. For instance, the active return r^{S1} of Sector 1 over the four reporting quarters is given by

$$r^{S1} = \frac{1 + 5.5425\%}{1 + 3.8438\%} - 1 = 1.6358\%$$

The active returns for each sector may then be combined to give the total active return.

Unlike Carino smoothing, geometrical smoothing does not require that the original unsmoothed data be stored. Once security-level geometrically smoothed return contributions are calculated, they may be combined with previously calculated returns over any interval.

Foreign exchange return and smoothing

The situation becomes more complex still when multi-currency portfolios are considered.

Recall that local currency return and currency return combine geometrically:

$$1 + r_{base} = (1 + r_{local})(1 + r_{FX}) \quad (10)$$

or equivalently

$$1 + r_{FX} = \frac{(1 + r_{base})}{(1 + r_{local})} \quad (11)$$

In other words, the decomposition of base currency returns into local currency returns and foreign exchange returns is subject to the same difficulties as is compounding over time.

While it is possible to cast Brinson attribution in geometric terms to remove this problem (see, Chapter 5) the mathematics is complex, especially for cases where each security generates multiple sources of return. Given that the results will be adjusted anyway, it is often simpler

to treat FX return additively, and set the local currency of a security as its base currency return, minus its local currency return:

$$r_{FX} = r_{base} - r_{local} \quad (12)$$

To apply smoothing, FX return may then be treated in the same way as any other source of return. The overall return of the portfolio will be unaffected, and it is arguable whether any significant detail will be lost after smoothing is applied. The benefit is a much simpler calculation.

Summary

Pragmatically, the user should be able to present returns in a form that is acceptable by clients. Inevitably, some will comment that ‘the numbers don’t add up’ when presented with an aggregated geometric report. Rather than trying to convey the concepts outlined in this chapter, it is often easier to present an arithmetically smoothed report.

For the same reason, you may also find different areas of your organization requesting different smoothing algorithms for results on the same portfolio.

Unless your returns are particularly extreme, the choice of smoothing method will seldom make much difference to the overall conclusions of your attribution analysis. There is no ‘best’ way to perform smoothing, and a good software system will offer a range of options to allow the user to select and experiment with various smoothing algorithms.

The material in this paper has been adapted from Chapter 5 of ‘Mastering Attribution in Finance’, Colin, A., 2016, FT Press/Pearsons.