A Funded-Path Random-Order Method for Portfolio Active Risk Attribution

Grant Holtes

July 20, 2026

Summary

The standard Euler decomposition of active risk evaluates every position at the final allocation, and so assigns large offsetting positive and negative contributions to positions that hedge one another, obscuring which decisions drive tracking error. This paper instead treats each allocation change as a discrete decision applied sequentially from the benchmark, defining its contribution as the expected change in active risk over a random ordering of the decisions. Each step is kept fully invested by one of two funding rules: offsetting pending trades, which treats decisions as relative trades, or the cash asset class, which treats them as standalone positions. Both attributions reconcile exactly to total active risk, remove the offsetting-sign artefact, and retain negative contributions only for positions that reduce risk on average, so a negative value identifies an allocation that is, on average, a hedge. The choice of funding rule lets practitioners align the attribution with how allocation decisions are framed.

1 Introduction

A portfolio is rarely held exactly at its benchmark, and the overweights and underweights that remain create *active risk*: the volatility of the portfolio's return relative to the benchmark, also called tracking error. Which asset-class decisions are responsible for that risk, and by how much? A useful answer should sum exactly to the total and give each asset class a share that matches its economic role.

The standard tool is the marginal, or Euler, decomposition, which splits total active risk into per-asset contributions using each position's sensitivity to a small change in its weight. It is exact, closed-form, and near instant, but it evaluates every position only at the final allocation. When two positions offset one another, for example an overweight funded by an underweight to a very similar asset class, it can assign one leg a large positive share and the other a large negative share, even though the pair adds almost no risk. These signs are mathematically correct but hard to interpret, and they obscure which decisions truly drive the risk. Table 1 shows the effect: the Euler contributions for listed equities and local bonds are large and negative.

Table 1: Example active risk attribution under Euler decomposition

	Benchmark	Portfolio	Tilt	Euler contribution
Local Equities	20%	18%	-2%	-0.1127%
Global Equities	40%	38%	-2%	-0.1140%
Private Equity	0%	4%	+4%	0.4823%
Local Bonds	15%	10%	-5%	-0.1099%
Global Bonds	15%	20%	+5%	0.1223%
Cash	10%	10%	+0%	0.0000%
TOTAL	100%	100%	+0%	0.2680%

Viewed differently, the gap between portfolio and benchmark is a set of trades: here, a trade from listed into private equities and a trade between the bond asset classes. On that view, the private equity allocation should carry the risk of the equity trade, and the global bond overweight the risk of the bond trade.

This paper formalises that intuition as a *funded-path random-order active risk attribution*, inspired by the Shapley value from cooperative game theory [3]. Each allocation change is treated as a step; starting from the benchmark, the steps are applied one at a time, the active risk added by each is recorded, and the results are averaged over a large number of random orderings [4].

The random-order construction can be viewed as a repair of the *incremental* (or leave-one-out) method developed by Balog et al.[1], which measures each position's contribution as the change in total risk when that position alone is added or removed. Evaluated at a single margin, these increments do not sum to total active risk and depend arbitrarily on whether the position is treated as first in or last in; averaging the increment over every possible position in the ordering removes both defects, restoring exact reconciliation while retaining the incremental method's intuitive question of what each decision adds.

To keep every step a fully invested trade, each change must be funded. The funding choice should align with how active decisions are framed in the portfolio, and we propose two rules:

Funding rule 1: Offsetting trades (FPRO). Each allocation change is offset by a change to an asset class with a pending move in the opposite direction; where several qualify, the pairing with the minimum active risk is chosen. This mirrors the example above, with equity trades funded from equities and bond trades from bonds, and suits processes that view decisions as relative trades (Figure 1).

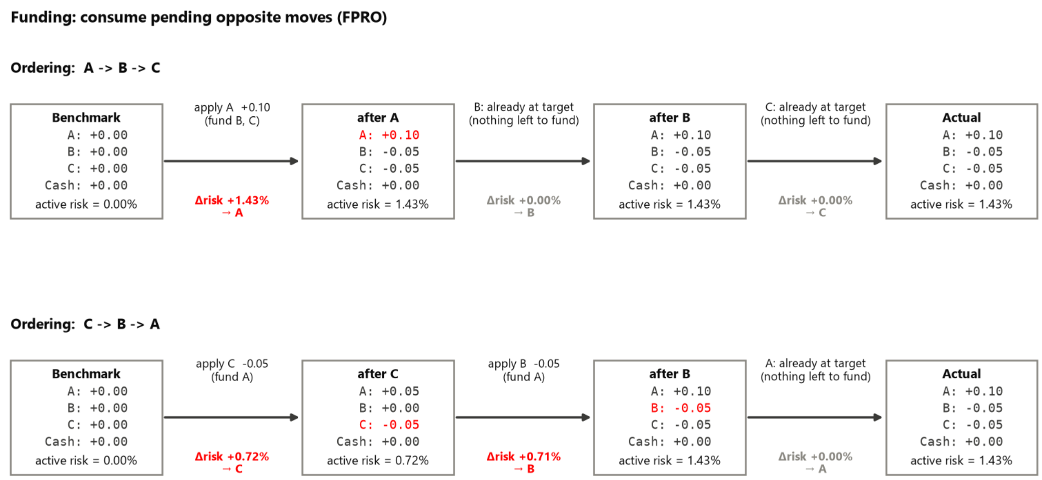

Figure 1: Simple example of two iterations of the FPRO algorithm. Note how ordering changes the active risk attributed to each asset class.

Funding rule 2: Cash (cFPRO). All changes are funded from the cash asset class, suiting processes that view decisions as standalone positions (Figure 2).

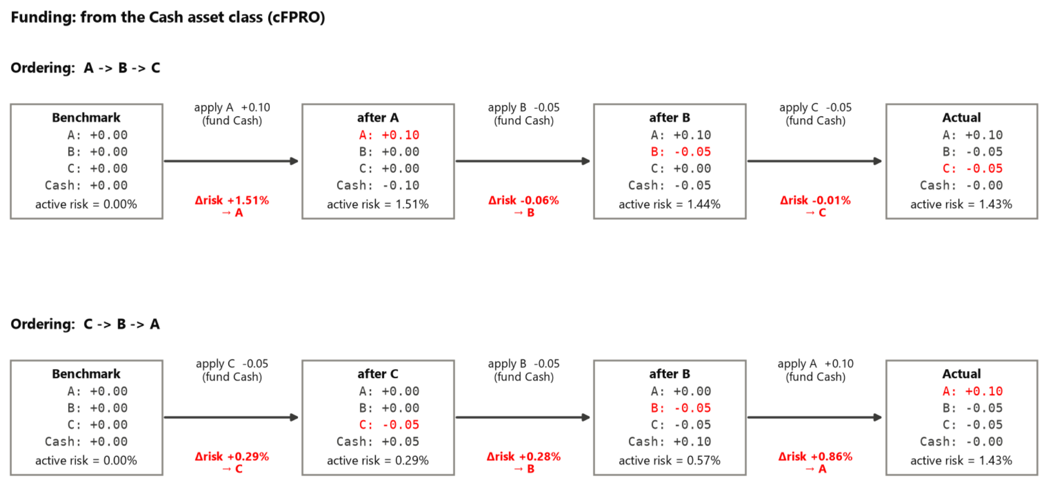

Figure 2: Simple example of two iterations of the cFPRO algorithm. Note how the cash value fluctuates through the steps, which may imply leverage in the process.

The two rules mirror a distinction Scherer [2] drew for marginal tracking error contributions, between the cash-funding convention of standard risk systems and the more intuitive framing of active positions as long/short trades.

The results of these methods can be compared with the Euler results in Table 2, which shows that the large offsetting negatives have been removed and the results align more closely with the decision-based intuition of where risk is concentrated.

Table 2: Active-risk contribution by asset class: Covariance, FPRO, and cFPRO.

	Covariance	FPRO	cFPRO
Local Equities	-0.11%	0.06%	0.02%
Global Equities	-0.11%	0.06%	0.02%
Private Equity	0.48%	0.12%	0.17%
Local Bonds	-0.11%	0.02%	-0.00%
Global Bonds	0.12%	0.02%	0.06%
TOTAL	0.27%	0.27%	0.27%

2 Properties and comparison of the estimators

One of the main motivations of a non-Euler approach is to avoid the large offsetting +/- contributions that it produces for highly correlated pair trades. This can be seen in Figure 3, where the Euler attribution of a 100% Bonds portfolio against a 100% Equities benchmark has a large positive and a negative component that emerges as correlations increase towards 1.

In contrast, the FPRO attributions provide positive, non-offsetting results for both asset classes. The offsetting trades version allocates the active risk equally between the two assets as the trade impact is averaged over the possible orderings. This equal split is a property of the two-move case rather than of FPRO in general; in larger portfolios the ordering average does not reduce to a coin flip and paired legs need not share risk equally.

In contrast, the cash funded cFPRO version allocates more of the risk to the higher volatility asset, with the higher vol asset class absorbing a larger portion of the risk as the correlation increases.

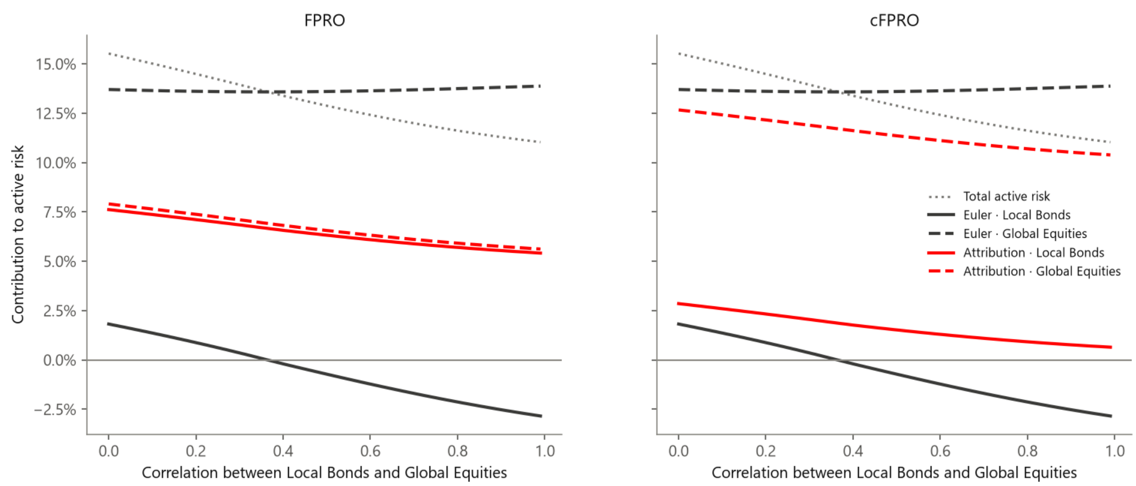

Figure 3: Per-asset contribution to active risk for a spread pair (one asset long, one short) as the correlation between the two sleeves is varied from 0 to 1. Each panel plots the Euler decomposition (grey) against the random-order attribution (red) under one funding method: FPRO (left, fund from pending opposite moves) and cFPRO (right, fund from the Cash asset class).

Negative attributions can still arise under either method, but only for moves whose expected marginal contribution across orderings is negative, that is, for actual hedges: positions that reduce total active risk given the rest of the active portfolio, such as a tactical unwind of a strategic tilt. Under Euler, a negative contribution is ambiguous between a hedge and a spread artefact; under the random-order methods it is unambiguous.

The two funding rules repair the Euler artefact in different ways, and Figure 4 makes the contrast visible across 250 randomly generated active portfolios. FPRO pairs each position with an offsetting counterparty, so the spurious negative leg of a spread is absorbed into the trade and its sign flips positive: it disagrees with Euler on sign frequently, and by design. cFPRO does not pair positions; it shrinks both legs of an offsetting spread towards zero without necessarily crossing the sign boundary, so its attributions track Euler closely and sign disagreements are rarer, concentrated near the origin where the contributions are economically small.

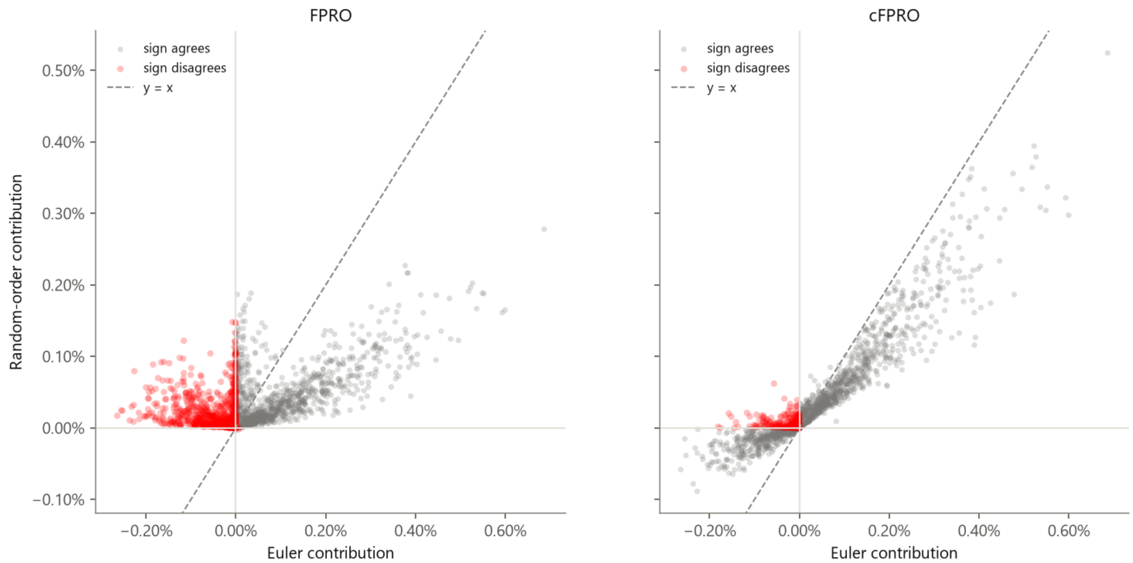

Figure 4: Per-asset active-risk contributions across a large sample of randomly generated active portfolios: the Euler decomposition (horizontal axis) against the random-order attribution (vertical axis), for FPRO (left) and cFPRO (right). Points where the method disagrees with Euler on the sign of the contribution are highlighted in red; the dashed line marks $y = x$.

One major difference between the funding methods is the level of "concentration" in the attributions, with FPRO resulting in attributions being more spread out over the asset classes (less concentrated), while cFPRO tends to increase concentration (Figure 5).

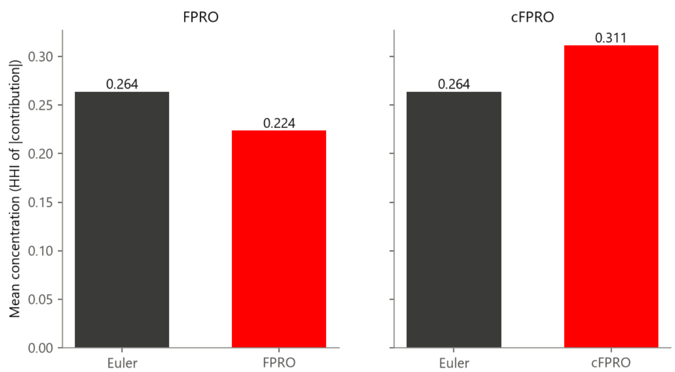

Figure 5: Concentration of the attributed active risk, measured by the Herfindahl–Hirschman index of the shares of absolute per-asset contributions (HHI $= \sum_i s_i^2$, $s_i = |c_i| / \sum_j |c_j|$) and averaged across the random portfolios. Each panel compares Euler with one funding method: FPRO (left) and cFPRO (right).

As the algorithm is a Monte Carlo estimate, the sample size used determines the level of error in the estimate. This error decays with $1/\sqrt{n}$, but practically the error is very small once multiple thousand samples are used. For the purposes of the other figures in this paper 2000 samples are used.

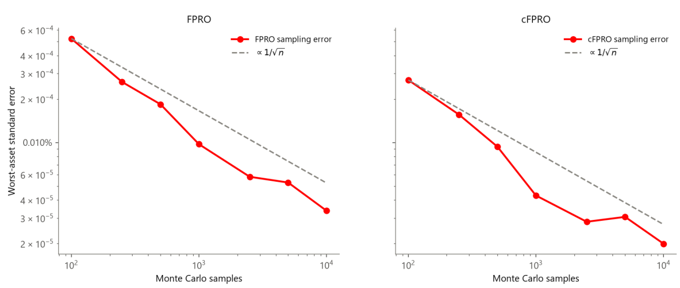

Figure 6: Monte Carlo sampling error of the random-order attribution (worst asset) against the number of sampled orderings, on logarithmic axes with a $1/\sqrt{n}$ reference line, for FPRO (left) and cFPRO (right).

The run-time of the algorithm scales with both the number of samples and the number of asset classes with active allocations, as this determines the length of the trade paths simulated. The more complex funding mechanism of the FPRO version results in a materially longer run time than cFPRO as asset class count increases.

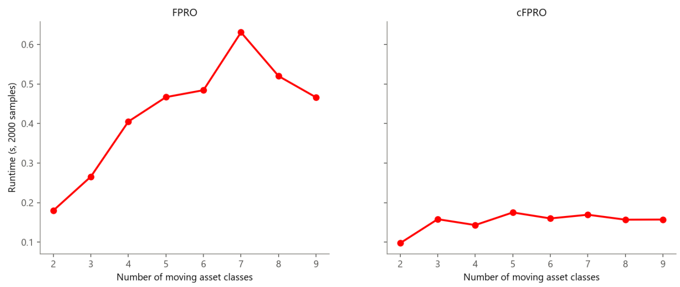

Figure 7: Run time of the random-order attribution, at a fixed number of sampled orderings, as a function of the number of moving asset classes, for FPRO (left) and cFPRO (right).

3 Detailed Methodology

3.1 Euler (covariance) contribution

Let $\mathbf{w} = \mathbf{w}_{\text{actual}} - \mathbf{w}_{\text{bench}} \in R^N$ be the vector of active weights and Σ the asset covariance matrix. Total active risk (tracking error) is

$$\sigma \;=\; \sqrt{\mathbf{w}^\top \Sigma\, \mathbf{w}}. \tag{1}$$

Because σ is homogeneous of degree one in $\mathbf{w}$, Euler's theorem yields an exact additive split into per-asset *component contributions to risk*,

$$\sigma \;=\; \sum_{i=1}^{N} c_i, \qquad c_i \;=\; w_i\,\frac{(\Sigma\mathbf{w})_i}{\sigma}, \tag{2}$$

where $(\Sigma\mathbf{w})_i/\sigma = \partial\sigma/\partial w_i$ is asset i's marginal contribution to risk. The decomposition is closed-form and exact, but each c_i is evaluated at the single, final active position and may be negative for a position that diversifies against the rest of the active bet.

Layer attribution. The active bet can be split along the path Benchmark $\rightarrow$ SAA $\rightarrow$ Actual into a strategic and a tactical layer,

$$\mathbf{w} = \underbrace{(\mathbf{w}_{\text{SAA}} - \mathbf{w}_{\text{bench}})}_{\mathbf{w}^S} + \underbrace{(\mathbf{w}_{\text{actual}} - \mathbf{w}_{\text{SAA}})}_{\mathbf{w}^T}. \tag{3}$$

Because the marginal contribution to risk $(\Sigma\mathbf{w})_i/\sigma$ is a *fixed* vector (evaluated once at the total active position) and c_i is linear in the weight it multiplies, each asset's contribution splits additively across layers,

$$c_i = w_i\,\frac{(\Sigma\mathbf{w})_i}{\sigma} = c_i^S + c_i^T, \qquad c_i^\ell = w_i^\ell\,\frac{(\Sigma\mathbf{w})_i}{\sigma}, \tag{4}$$

and the layer totals $\sum_i c_i^\ell$ in turn sum to σ. Adding layers is trivial: for any chain of intermediate portfolios $\mathbf{w}_{\text{bench}} = \mathbf{p}_0, \mathbf{p}_1, \ldots, \mathbf{p}_L = \mathbf{w}_{\text{actual}}$ the active weight decomposes as $\mathbf{w} = \sum_{\ell=1}^{L}(\mathbf{p}_\ell - \mathbf{p}_{\ell-1})$, and scoring each increment against the same fixed marginal-risk vector introduces no further computation.

3.2 Random-order attribution

We decompose the same total active risk $\sigma = \sqrt{\mathbf{w}^\top \Sigma \mathbf{w}}$, with $\mathbf{w} = \mathbf{w}_{\text{actual}} - \mathbf{w}_{\text{bench}}$, by treating each individual allocation change as a player, applying the changes sequentially from the benchmark, and averaging each change's marginal risk contribution over the order of application. Write $\sigma(\mathbf{x}) = \sqrt{\mathbf{x}^\top \Sigma \mathbf{x}}$ for the active risk of any active-weight vector $\mathbf{x}$.

This is a single framework admitting two alternative *funding rules*, which differ in how each allocation change is financed and, in general, produce different attributions. We describe the shared framework first, then the two rules.

Players. As the allocation travels Benchmark $\to$ SAA $\to$ Actual, each asset i makes up to two *moves*: a strategic move $\delta_i^S = w_i^{\text{SAA}} - w_i^{\text{bench}}$ and a tactical move $\delta_i^T = w_i^{\text{actual}} - w_i^{\text{SAA}}$. Discarding zero moves leaves a set $\mathcal{M}$ of signed weight changes; move m acts on asset $\alpha(m)$ with size δ_m, and within each layer the moves are collectively self-funding, $\sum_m \delta_m = 0$. Adding layers is trivial: the path may pass through any chain of intermediate portfolios, each additional waypoint contributing one further move per asset, with the game and its estimator otherwise unchanged.

Random-order walk. Given an ordering π of $\mathcal{M}$, we walk from the benchmark portfolio $\mathbf{p} = \mathbf{w}_{\text{bench}}$, applying one move at a time. Applying a move changes one asset by its size and, to keep the weights summing to one, is *funded* by an offsetting change elsewhere (the funding rule fixes where). After each funded move the active risk is re-evaluated, and the move's marginal contribution in this ordering is the induced change,

$$\Delta_m(\pi) \;=\; \sigma\big(\mathbf{p}^{\text{after}} - \mathbf{w}_{\text{bench}}\big) - \sigma\big(\mathbf{p}^{\text{before}} - \mathbf{w}_{\text{bench}}\big)\,, \tag{5}$$

credited entirely to the move being applied; the funding counterparties are passive.

Attribution. A move's attribution is *defined* as its expected marginal contribution under a uniformly random ordering of the moves,

$$\phi_m \;=\; E_\pi[\Delta_m(\pi)]\,, \tag{6}$$

and an asset's contribution sums the values of its moves across the layers in which it moves,

$$\Phi_i \;=\; \sum_{m:\,\alpha(m)=i} \phi_m, \tag{7}$$

which here means the strategic and tactical layers (at most two terms per asset, and one term per layer in general). How the expectation in (6) is computed depends on the funding rule, as set out below.

Funding rule 1: offsetting trades (FPRO). The offsetting $-\delta_m$ is drawn from moves that still have an unexecuted amount in the opposite direction, so the funding advances those assets toward their own targets. Among eligible funders we select the counterparty (a different asset class wherever possible) that minimises the active risk of the resulting trade[1], spilling to the next-best only when one lacks capacity. Because each funded trade consumes counterparties' pending

[1] The active risk measured is the total change in active risk for the trade, not the traded pair alone

remainders, the portfolio reached after a given set of moves depends on the order in which they arrived. Every move's delta nonetheless lands on its own asset exactly once in every ordering. The expectation (6) has no closed form and is estimated by Monte Carlo over K independently sampled uniform permutations $\pi_1, \ldots, \pi_K$,

$$\hat{\phi}_m \;=\; \frac{1}{K}\sum_{k=1}^{K} \Delta_m(\pi_k), \tag{8}$$

converging to ϕ_m as K grows large. We call this funded-path random-order attribution (FPRO).

Funding rule 2: cash (cFPRO). The offsetting $-\delta_m$ is held in a designated cash asset class, so move m contributes the fixed active-weight vector $\delta_m(\mathbf{e}_{\alpha(m)} - \mathbf{e}_{\text{cash}})$, and cash carries its own volatility and correlations through Σ. Each move is thereby measured as a position in the asset funded from cash, that is, as an excess-over-cash exposure. Cash serves as the numeraire rather than a player: its own moves are dropped from $\mathcal{M}$, and the funding legs of the remaining moves accumulate, by the self-funding identity $\sum_m \delta_m = 0$ within each layer, to exactly cash's own active weight, so the grand coalition still lands precisely on $\mathbf{w}$. A consequence of this convention is that a deliberate active cash position is not reported as its own line item; its risk effect is absorbed into the funding legs of the other moves.
We call this cash-funded random-order attribution (cFPRO), noting that the random ordering here matters only to the sampling estimator, not to the value being estimated.

Reconciliation. Under either funding rule every ordering terminates exactly at the Actual portfolio (the moves consume one another under rule 1, and sum to $\mathbf{w}$ under rule 2), so the per-move marginals telescope: $\sum_{m\in\mathcal{M}} \Delta_m(\pi) = \sigma(\mathbf{w})$ for every π. Hence the attributions sum to total active risk exactly, with no rescaling.

3.3 Algorithm

Algorithm 1 Funded-path random-order attribution (FPRO)

Require: moves $\mathcal{M}$ with signed sizes δ_m on assets; $\alpha(m)$; covariance Σ; number of samples K
Ensure: attribution ϕ_m for every move

$\phi \leftarrow \mathbf{0}$
for $k = 1, \dots, K$ **do**
  $x \leftarrow \mathbf{0}, \quad r \leftarrow \delta, \quad v_{\text{prev}} \leftarrow 0$ ▷ x: active weights; r: unexecuted amounts
  for m in a uniformly random permutation of $\mathcal{M}$ **do**
    $x_{\alpha(m)} \leftarrow x_{\alpha(m)} + r_m$ ▷ apply the move
    need $\leftarrow |r_m|$
    $C \leftarrow \{\, j \neq m : \operatorname{sign}(r_j) = -\operatorname{sign}(r_m) \,\}$ ▷ pending opposite moves
    sort C: a different asset class first, then the smallest change in active risk
    for $j \in C$ **while** need > 0 **do**
      $a \leftarrow \min(\text{need}, |r_j|)$
      $x_{\alpha(j)} \leftarrow x_{\alpha(j)} - \operatorname{sign}(r_m)\, a, \quad r_j \leftarrow r_j + \operatorname{sign}(r_m)\, a$ ▷ fund from j
      need $\leftarrow$ need $- a$
    end for
    $r_m \leftarrow 0$
    $v \leftarrow \sigma(x)$
    $\phi_m \leftarrow \phi_m + (v - v_{\text{prev}}), \quad v_{\text{prev}} \leftarrow v$ ▷ credit the marginal to m
  end for
end for
return ϕ / K

An asset's contribution is the sum of its moves' values, $\Phi_i = \sum_{m:\, \alpha(m)=i} \phi_m$, and the layer totals follow by summing within each layer. The cash-funded variant (cFPRO) is identical except for the funding step: instead of consuming pending opposite moves, each move is funded from a single cash asset, i.e. the funding block is replaced by the single update $x_{\text{cash}} \leftarrow x_{\text{cash}} - r_m$.

Note:

Views expressed are the author's, and may differ from those of JANA investments. This material does not constitute investment advice and should not be relied upon as such. Investors should seek independent advice before making investment decisions. Past performance cannot guarantee future results. The charts and tables are shown for illustrative purposes only.

References

[1] Dóra Balog et al. "Properties and Comparison of Risk Capital Allocation Methods". In: *European Journal of Operational Research* 259.2 (2017), pp. 614–625. DOI: 10.1016/j.ejor.2016.10.052.

[2] Bernd Scherer. "A Note on Tracking Error Funding Assumptions". In: *Journal of Asset Management* 2.3 (2001), pp. 235–240. DOI: 10.1057/palgrave.jam.2240048.

[3] Lloyd S. Shapley. "A Value for n-Person Games". In: *Contributions to the Theory of Games, Volume II*. Ed. by Harold W. Kuhn and Albert W. Tucker. Vol. 2. Princeton: Princeton University Press, 1953, pp. 307–317.

[4] Robert J. Weber. "Probabilistic Values for Games". In: *The Shapley Value: Essays in Honor of Lloyd S. Shapley*. Ed. by Alvin E. Roth. Cambridge: Cambridge University Press, 1988, pp. 101–119.

4 Appendix

Table 3: Asset-class volatility assumptions (annualised). All values are examples only and are not based on actual index data.

Asset class	Volatility
Local Equities	16.1%
Global Equities	16.2%
Private Equity	20.7%
Unlisted Infra	12.0%
Private Debt	11.5%
High Yield	9.5%
Local Bonds	4.9%
Global Bonds	4.9%
Cash	0.5%

Table 4: Asset-class correlation assumptions. Columns are numbered as the rows. All values are examples only and are not based on actual index data.

	1	2	3	4	5	6	7	8	9
1. Local Equities	1.00	0.98	0.95	0.85	0.94	0.80	-0.11	-0.11	-0.13
2. Global Equities	0.98	1.00	0.94	0.85	0.94	0.80	-0.12	-0.11	-0.13
3. Private Equity	0.95	0.94	1.00	0.84	0.94	0.83	0.01	0.01	0.01
4. Unlisted Infra	0.85	0.85	0.84	1.00	0.85	0.78	0.09	0.09	0.06
5. Private Debt	0.94	0.94	0.94	0.85	1.00	0.89	0.13	0.13	0.11
6. High Yield	0.80	0.80	0.83	0.78	0.89	1.00	0.39	0.38	0.35
7. Local Bonds	-0.11	-0.12	0.01	0.09	0.13	0.39	1.00	0.96	0.94
8. Global Bonds	-0.11	-0.11	0.01	0.09	0.13	0.38	0.96	1.00	0.93
9. Cash	-0.13	-0.13	0.01	0.06	0.11	0.35	0.94	0.93	1.00