THREE HORIZONS CAPITAL · PORTFOLIO ANALYTICS

A 60/40 Portfolio Is Not One Strategy

The label names a weight. It does not name a portfolio, and it does not name a return.

14–66%
Net equity across five live funds all sold as balanced, moderate or tactical

146bp
Return spread across five portfolios held at exactly 60/40 by weight

99bp
Spread on one unchanged portfolio, depending on whose assumptions were used

Fund allocations read from our own data 31 August 2026, at each fund's own portfolio date, spanning 31 March to 31 July 2026 · Capital market assumptions read 31 August 2026, 2026 vintage, ten-year horizon

US dollars throughout · Professional investors · No fund, house, provider or issuer is named

Executive summary

FINDING 1 Five funds currently marketed as balanced, moderate or tactical hold between **13.8% and 66.3%** in equities. Two of them carry the **identical category tag** and sit 18.5 points apart, with unrelated strategies underneath.

FINDING 2 Hold the split at exactly 60/40 and the portfolios still diverge by **146 basis points of return and 280 of volatility**, with no leverage and nothing exotic. Diversification ranges from a two-asset bet to four effective positions.

FINDING 3 Growth beta still dominates every one of them. That is the honest caveat, and it belongs at the front: the differentiation is real but it lives at the margins, in tail loss, inflation protection and whether illiquidity was taken deliberately.

FINDING 4 Change nothing in the portfolio and its expected return still moves **99 basis points** across six houses' published assumptions. The houses differ by nearly three times on what the equity leg earns over the bond leg.

FINDING 5 One house publishes two correct expected returns for the same asset, **124 basis points apart**. That single construct choice is larger than the disagreement between all six houses.

> **The label hides the portfolio, and the model hides the label.** A 60/40 allocation is not a decision you have made. It is a decision you have deferred, twice.

1 Five funds sold as balanced hold between 14% and 66% in equities

Two portfolios arrive on the same page of the same board pack. Both report sixty per cent equities and forty per cent fixed income. The committee compares them on fee, on manager tenure, and on three-year performance, and moves on.

They are not the same portfolio. They may not even be the same kind of portfolio.

Start with what is actually on sale. These are five live funds, all currently marketed under balanced, moderate or tactical allocation labels, and between them they hold $545bn.

CHART 1 · WHAT THE LABEL DOES NOT SAY

Five funds all marketed as balanced, moderate or tactical hold anywhere from 14% to 66% in equities.

Net equity exposure, disclosed, each at its own portfolio date. The two funds outlined in white carry the *identical* category tag.

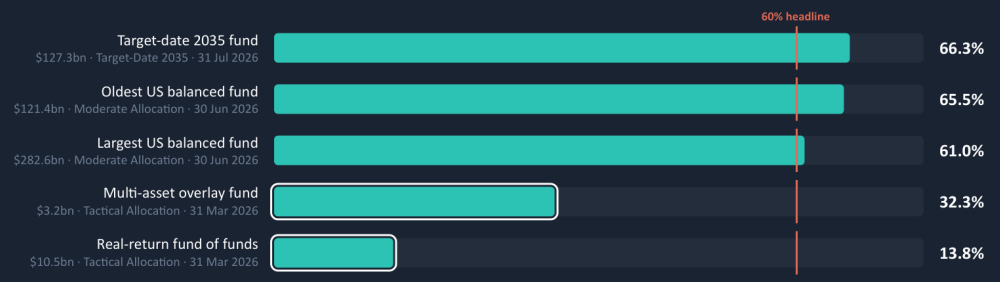

Disclosed net equity, read from our own fund data 31 August 2026. **The five funds report to three different portfolio dates, from 31 March to 31 July 2026, shown per fund above.** A spread of 52.5 points is far too wide to be an artefact of four months. Funds are not named: every finding holds without them.

Disclosed net equity, each fund at its own portfolio date. The spread is 52.5 points.

The two funds circled are the sharper finding. They do not merely sit at different points on a spectrum. They carry the **identical category tag**, and they sit 18.5 points apart on equity with unrelated strategies underneath. One runs trend and macro overlays. The other is a levered fund of funds built around inflation-linked bonds and real assets.

That second fund reports a cash position of **minus 21.5 per cent**. It is borrowing, through futures and repo, and the borrowing is disclosed in the filings. It is not disclosed by the words on the front of the fund. A plain-language allocation name gives no indication that leverage is present at all.

> So the category tag is not a backstop for the allocation label. It fails in the same direction, and it fails for the same reason.

2 Hold the split at exactly 60/40 and the portfolios still diverge

The obvious objection is that those five funds are not really 60/40 vehicles, so of course they differ. Fair. So we removed the variable.

We built five portfolios that are precisely sixty per cent return-seeking and forty per cent risk-reducing, to the decimal, and changed only what sits inside each sleeve. Any divergence that survives is composition and nothing else.

It survives. Across the five, expected return spans **146 basis points** and volatility spans **280 basis points**, with no leverage and no exotic instruments anywhere in them. The concentration measure runs from 0.24 to 0.85 on a scale where 1.0 is a single holding, and the effective number of independent bets runs from **1.18 to 4.17**.

The clearest way to see it is the effective number of positions, which asks how many genuinely independent bets a portfolio is running rather than how many lines its allocation table happens to carry.

CHART 2 · WHAT THE LABEL HIDES

Hold the 60/40 split exactly constant and diversification still ranges from a two-asset bet to six effective positions.

Effective number of positions. Five portfolios built to be precisely 60/40 by weight, and five live funds.

Effective number of positions. A value near 1.0 means the portfolio behaves as a single bet however many lines its allocation table carries. Constructed scores are ours. Live-fund scores are provisional, pending independent re-derivation.

The classic domestic construction scores **1.18**. It has a four-line allocation schedule and it behaves as a two-asset bet. The pension-style construction, built from the same sixty and forty, scores **4.17**. Both are ours and both are re-derived.

The live funds sit across the same range and reach further at the top. Those scores are the analyst's computation and are **provisional pending independent re-derivation**, so they are reported here as direction rather than as settled figures. On that basis two things stand out. The longest-running balanced fund in the set sits near the bottom, running one dominant exposure with a bond position attached. And two funds less than a point apart on the headline split, at 65.5 and 66.3 per cent equity, are materially different underneath, because one puts roughly a third of its equity sleeve outside its home market and the other does not.

A board that reviews the equity and bond percentages, and stops there, is holding two structurally different risk books and calling them the same thing.

3 The honest caveat, stated up front

Here is the part that a paper making this argument is tempted to bury.

Every one of these portfolios, constructed and live, is still dominated by a single factor. Growth beta explains the large majority of the variance in all ten. The creative ones are less dominated than the plain ones, but none of them escapes. If you were hoping that a cleverly built 60/40 becomes a genuinely balanced set of risks, it does not.

That does not weaken the argument. It sharpens it. The differentiation between these portfolios is real but it lives at the margins: in tail loss, in how much inflation protection is actually present, and in whether illiquidity has been taken deliberately or by accident. Those margins are exactly what the headline label cannot show you, and exactly what determines behaviour in the year the portfolio is tested.

The one portfolio in the set that scores best on risk efficiency is the one deliberately built against something other than growth beta. That is the tell. A book holding that profile under a 60/40 label made a choice. A book holding the concentrated profile under the same label may simply have inherited a default that nobody has interrogated since it was set.

> The label cannot distinguish between those two boards. It reports the same number for both.

4 Now change nothing in the portfolio and watch the answer move anyway

Everything above was computed on one house's capital market assumptions. That is standard practice, and it is also a choice, and the choice is not neutral.

We hold the published long-term assumptions of eight houses in one place, normalised to a common structure. So we took the simplest possible portfolio, sixty per cent US equity against forty per cent US core bonds, held it completely fixed, and priced it on each house's own numbers. Same weights, same currency, same ten-year horizon, same vintage.

CHART 3 · WHAT THE MODEL HIDES

Change nothing in the portfolio and its expected return still moves 99 basis points, depending only on whose assumptions you use.

The identical 60% US equity / 40% US core bonds, priced on each house's own 2026 ten-year assumptions. Circled: the two ends of the range.

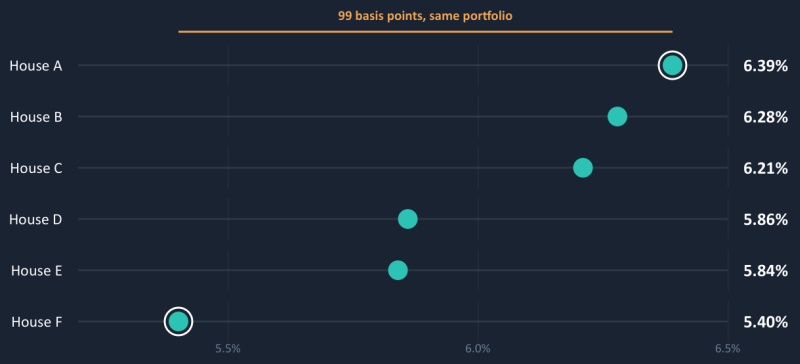

Six houses' published long-term capital market assumptions, 2026 vintage, US dollars, ten-year horizon, compound returns, read from our own store on 31 August 2026. Of the eight houses we hold, one publishes no single core-bond line and is excluded rather than estimated, and one is a 2025 vintage and is excluded to hold the vintage constant. Houses are not named.

Ninety-nine basis points, on a portfolio that never changed. Set against the 146 basis points of spread across five deliberately different portfolios, the lens is worth roughly two thirds of the thing it was brought in to measure.

And the disagreement is not a simple optimism dial, which would at least be easy to adjust for. It is a disagreement about shape.

CHART 4 · THE DISAGREEMENT IS ABOUT SHAPE, NOT LEVEL

The houses differ by nearly three times on what the equity leg is paid over the bond leg: 140 to 383 basis points.

Each house's own US equity and US core bond assumption, and the gap between them.

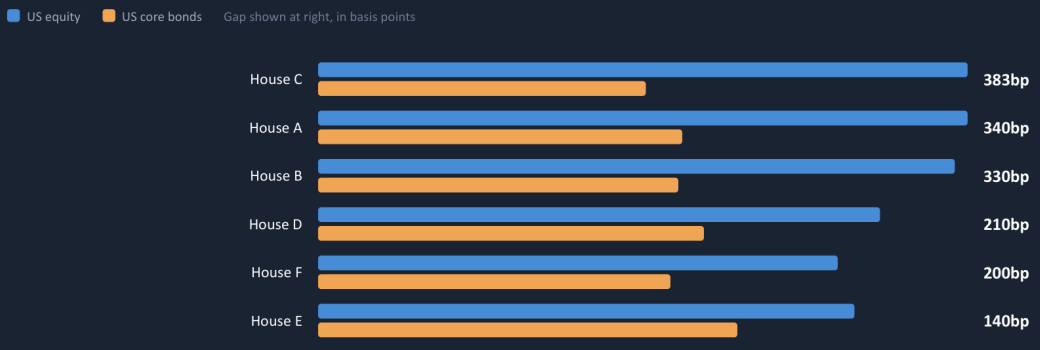

Bars are scaled to the highest equity assumption in the set. The widest gap comes from a house pairing a near-top equity number with the lowest bond number. **These gaps compound in the premium but partly offset in a 60/40 blend**, because the houses expecting most from equities also expect least from bonds, which is why a 243 basis point spread in the premium produces only 99 in the portfolio. Same source and date as Chart 3.

The premium these houses expect for holding the equity leg rather than the bond leg runs from 140 to 383 basis points.

They differ by nearly three times on whether the sixty is worth having at all. Run an optimiser on the low number and it will tell you to hold less equity. Run it on the high number and it will tell you to hold more.

Both are defensible published views from serious institutions.

Notice what the blend then does to that disagreement. A 243 basis point spread in the premium, applied at a sixty per cent weight, would arithmetically produce about 146 basis points of portfolio spread. It produces 99. The reason is that the houses expecting the most from equities are also the ones expecting the least from bonds, so at the portfolio level the two disagreements partly cancel. **The 60/40 blend does not just hide what a portfolio holds. It hides how far apart the forecasters are.** A committee looking only at the blended expected return sees a fraction of the disagreement that actually exists underneath it.

One further detail, because it is the sharpest of the lot. One of these houses publishes both an arithmetic and a compound expected return for the same US equity line, and the two differ by **124 basis points**. That is a choice about which of two correct numbers to use, and on its own it is larger than the 99 basis points of disagreement between all six houses, and comparable to the 146 basis points of difference between five genuinely different portfolios.

> The measurement choice outweighs the thing being measured. That should stop a committee in its tracks.

5 What to do about it

Not "60/40 is dead". That argument has been made many times and it is not what this shows.

What this shows is that the label is a classification, and a classification is not a risk statement. Three questions follow, and they are answerable in an afternoon.

- **What is actually inside the sleeves?** Not the asset-class percentages. The geography, the credit quality, the duration, and whether anything in the risk-reducing forty is behaving like equity when it matters.
- **How many independent bets is the book really running?** If the answer is close to one, the allocation schedule is decoration. That is a governance finding, not a performance finding, and it does not appear anywhere in a standard report.
- **Whose assumptions produced the expected return you are planning against, and what happens under someone else's?** If the number moves by a hundred basis points and nobody in the room knew it could, the planning conversation was never really about the portfolio.

None of these require a new mandate, a new manager or a new system. They require decomposing what is already there.

WHAT WE CHECKED, AND WHAT WE DID NOT

The five live funds' disclosed allocations and assets were read from our own fund data and re-derived there on 31 August 2026; four of the five reproduce exactly and the fifth differs by 0.13 of a point because our data has moved to a more recent portfolio date. **The five funds report to three different portfolio dates, from 31 March to 31 July 2026.** Each is stated per fund on the first chart rather than averaged into one "as at". A 52.5-point spread in net equity is far too wide to be an artefact of four months, but the dates belong on the page.

The five live funds' computed risk measures, including the effective-position scores, are the analyst's computation and are **pending independent re-derivation**. The five constructed portfolios' figures are ours. The two sets are kept apart in the text for that reason.

The six-house comparison was derived from the published assumptions we hold, on 31 August 2026, at a single vintage, currency and horizon. **The mapping from each house's own asset naming to a common pair of legs was declared by hand rather than matched on names**, because a name match returned only four of the houses: the others write the same asset as "US equity", "US stocks" and "US large cap equities". That under-count would have halved the comparison without announcing itself.

Two of the eight houses are excluded rather than estimated: one publishes no single core-bond line, and one is a different vintage. We would rather drop a data point than manufacture one, since the whole argument here is about what happens when choices are made invisibly.

Two of the five live funds hold a material portion of assets that their filings describe only in aggregate. Those portions were mapped to the nearest comparable asset class using each fund's stated strategy. That is an estimate, it is the largest source of uncertainty in this work, and the funds concerned are the two most extreme in the set. The direction of the finding does not depend on it. The precise figures for those two funds do.

The constructed portfolios are illustrations built to isolate one variable. They are not model portfolios and nobody is invested in them. Funds and houses are not named: every finding holds without the names, and two of them are adverse.

The offer

If you run or oversee a balanced book, send us the allocation line. We will come back with what it actually holds, how many independent bets it is really running, and how much of its expected return is a modelling choice rather than a portfolio choice.

It takes us an afternoon. It is the question the label was never able to answer.