Guide to the dashboard

Edition of 8 November 2026 · For the dashboard at https://www.carlohofer.com/rules-based-portfolio/

Contents

About the portfolio and this guide

Before you start

The page around the views

Overview

- The six tiles
- Figure 1 | Allocation now against the target and the band
- Figure 2 | Growth of 100: the portfolio, the reference portfolios and the index blend
- Figure 3 | Drawdown against the mandate's limit

Rebalancing

- Figure 4 | Equity weight in its band, with every order
- Figure 5 | Orders per month against the allowance of five

Costs

- Figure 6 | Cost per month by source
- Figure 7 | Implementation cost, cumulative
- Table 1 | What the portfolio costs a year

Contributions

- Figure 8 | Money paid in and the value of the portfolio
- Figure 9 | Market effect and contribution by month

Look-through

- Equity sleeve
- Figure 10 | Equity by region
- Figure 11 | Equity by country: the largest
- Figure 12 | Equity by sector
- Figure 13 | The ten largest companies
- Bond sleeve
- Figure 14 | Bonds by issuing country
- Figure 15 | Bonds by maturity, and duration
- Figure 16 | Credit quality of the bond sleeve
- Whole portfolio
- Figure 17 | Currency exposure, by sleeve
- Risk of the whole portfolio
- The risk tiles
- Figure 18 | Correlation of the components' monthly returns
- Figure 19 | Each component's share of the estimated variance
- Figure 20 | The efficient frontier of the components, and the portfolio's point
- Table 2 | Today's portfolio in five historical episodes

Factor exposures

Figure 21 | Factor loadings of the equity ETF
Figure 22 | Rolling loadings
Figure 23 | Where the difference to the index blend came from
Table 3 | The difference to the index blend in three parts

Method

Figure 24 | The allocation test: growth per year against the worst fall of each split
Figure 25 | Simulated worst fall of each split over resampled ten-year paths
Figure 26 | Rolling 36-month correlation of the two sleeves' monthly returns
What the portfolio does not do, and why
Figure 27 | The allocation test on the US stock market
Figure 28 | Growth of 100: reference portfolio A and the S&P 500 alternative
Figure 29 | The ETF selection: tracking difference by fund and year
Table 4 | Data sources
Table 5 | Controls
The rules
Other projects

Simulations

Mechanics simulation
Figure 30 | Simulation: the equity weight in its band under the rules
Historical simulation
Figure 31 | Simulation: growth of 100 under the rules
Figure 32 | Simulation: fall from the previous peak under the rules
Resampled paths
Figure 33 | Resampled paths of the portfolio's value over ten years

Log

Quarterly reports
Monthly notes
Cycle reports
Table 6 | Departures register
Table 7 | The monthly table

Terms

About the portfolio and this guide

The rules-based portfolio is a real-money portfolio run under written rules, fixed and tagged before its first order on 8 October 2026. It holds two exchange-traded funds, or ETFs, which are funds whose shares are bought and sold on a stock exchange like the shares of a company: a global equity ETF and a euro government bond ETF. A starting amount, paid in before the first purchase, and a fixed contribution every month fund it. The rules produce the orders the portfolio places, and every departure from them is kept in a register. A program in the public repository computes every figure on the dashboard from the private record of orders and contributions and from the data the ETFs' issuers publish.

The dashboard is where that record is published. It is rebuilt every Monday and on each day the rules are applied, and it is organised in nine views: the Overview, Rebalancing, Costs, Contributions, Look-through, Factor exposures, Method, Simulations and Log. The first five views describe the portfolio as it is, the next two explain what its holdings are exposed to and how it was designed, and the last two hold the simulations on history and the written record of each month.

This guide explains every view, section, tile and exhibit of the dashboard for readers with or without a background in finance: what each one shows, how to read it, where its numbers come from and what it cannot tell. The guide is built with the dashboard, so its figure and table numbers are the dashboard's. Each entry opens with a link to its place on the dashboard, the references in the text link to the entry they name, and the terms at the end define every word the dashboard uses in a sense of its own.

Before you start

The entries that follow assume twelve ideas, which recur across the views of the dashboard.

The portfolio, its sleeves and their weights. The portfolio is the two ETFs held under the rules, plus any contribution not yet invested, which is held as cash. A sleeve is the part of the portfolio held in one ETF: the equity sleeve holds the iShares MSCI ACWI UCITS ETF, whose index covers large and mid-sized companies in developed and emerging markets, and the bond sleeve holds the Vanguard EUR Eurozone Government Bond UCITS ETF, whose index covers euro area government bonds. A sleeve's weight is its value divided by the value of the portfolio, in per cent.

Target weights, the split and the band. A target weight is the weight the rules aim for: 70 per cent for equity and 30 per cent for bonds. A split is the pair of target weights, written equity/bonds, so the portfolio's split is 70/30. The band is the interval from 65 to 75 per cent for the equity weight. Inside it the rules sell nothing, and when the equity weight is outside it on the day the rules are applied, the portfolio is brought back to 70/30 that day.

The cycle day, the top-up and the orders. The cycle day is the day each month on which the rules are applied: the 5th, or the first trading day after it when the 5th is not one. The top-up is the fixed monthly contribution, 5 per cent of the starting amount, paid in by the 2nd and invested on the cycle day in the sleeve furthest below its target weight. The rules allow at most five orders a month, each placed between 15:45 and 17:00 Amsterdam time.

Net asset value. An ETF's net asset value is the value of one of its shares at the close of a day, as its issuer publishes it, usually on the next working day. The dashboard values the portfolio at these values, the equity ETF's converted from US dollars at the reference rate of the European Central Bank (ECB). Because the values arrive a day late, the record ends at the last day for which both ETFs have published one, and the dashboard's header gives that date.

Growth of 100, with contributions removed. Most charts of returns start at 100 and show what 100 of the starting amount would have become. Contributions are removed from these lines, so a top-up does not count as a gain: each day's change is the return on the money the portfolio held that day, and the days are chained together. This is the time-weighted return. The money-weighted return counts the timing of the contributions as well, since it is the rate at which the contributions, each from its own date, grow to the portfolio's value. The two differ when money arrives just before a rise or a fall.

Reference portfolios A and B. A reference portfolio is a portfolio computed from the same contributions and the same prices as the portfolio, without costs. Reference portfolio A follows the same rules and buys at net asset value, so it is the portfolio as the rules would have run it if placing an order cost nothing. Reference portfolio B receives the same contributions and goes back to 70/30 on every cycle day.

The implementation cost and the basis point. The implementation cost is reference portfolio A's value relative to the portfolio's, with contributions removed: what running the portfolio has cost against the rules followed at no cost. It is given in basis points, and a basis point is one hundredth of a percentage point, so 10 basis points are 0.1 per cent. The implementation cost is positive when the portfolio is worth less than reference portfolio A and negative when it is worth more, as when a purchase was filled below the day's net asset value.

Drawdown, worst fall and the risk limit. A drawdown is the fall in the portfolio's value below its previous peak on a given day, with contributions removed, in per cent. The worst fall is the largest drawdown over a period. The risk limit is the worst fall the mandate, the document that sets the portfolio's objective, aims to stay within: about one third, tested as 35 per cent, with 40 per cent as an outer bound.

Look-through. The look-through is the portfolio's exposure to the securities inside the two ETFs, each ETF's holdings weighted by the ETF's weight in the portfolio. It turns two funds into the companies, countries, sectors, currencies and bonds that the portfolio holds through them.

The record, history and simulations. The dashboard shows three kinds of numbers, and each exhibit says which kind it shows. The record is what the program computes from the ledger, the private record of every contribution and order, and from the prices, since the first purchase. The allocation test and the simulations use index returns from February 1999 to December 2025, which run back before the two ETFs existed, to show how the rules behave over a longer history, and they are never part of the record. A record of five to ten years cannot show that one set of rules is better than another, so the dashboard does not offer its returns as evidence for the rules.

Growth per year, volatility and correlation. Growth per year is the compound annual return, with contributions removed. Volatility is the standard deviation of returns, a measure of how widely they spread around their average, and the dashboard gives it for a year. A correlation runs from −1, when two returns always move in opposite directions, through 0, when they are unrelated, to 1, when they always move together.

Exhibits and their parts. Every chart is a figure, and every table that stands on its own is a table, each numbered in page order across the views. Under the title, an "as of" line gives the dates and the basis of the numbers. A note explains what the exhibit shows, a source line names the file in the repository that holds its numbers, and a Table toggle opens the numbers behind the chart. Nothing depends on hovering over a chart, since every value is in its table.

The page around the views

The header, the bar of views and the footer stay the same on every view.

The header. Under the title, the dashboard gives the date of the record and the day of the first purchase, followed by links to the repository, the rules, the mandate and this guide. When the prices and the holdings have different dates, as after a cycle day whose orders are not yet valued, both dates are written.

The bar of views. The nine views are tabs in one bar, which stays at the top of the screen as the page scrolls. On a phone the bar scrolls sideways, and a fade marks each edge beyond which more tabs lie. A link to a part of a view opens that view at that part.

The footer. The footer states the basis of the returns: in euro, at the issuers' net asset values, time-weighted except where the money-weighted return is named, and not offered as evidence for the rules. It then lists the terms the dashboard uses, which are the terms at the end of this guide, and closes with the build date and a link to the code.

Overview

On the dashboard: the Overview view

The Overview answers where the portfolio stands on the record's date and how it has done since the first purchase, with contributions removed. It holds six tiles and three figures.

The six tiles

The six tiles give the figures a reader checks first, each with a line under it that says what the figure is measured against.

Equity weight. The equity sleeve's weight on the record's date. The line under it gives the change since the last complete month end, in percentage points, so a top-up, an order or a market move shows here.

Drift from 70%. The equity weight minus its target weight of 70 per cent, in percentage points. The band allows a drift from −5 to +5 points before the rules sell anything.

Return since the first purchase. The time-weighted return since the first purchase, with contributions removed: the growth of 100 of Figure 2, less 100, in per cent.

Money-weighted return. The money-weighted return from the first purchase to the record's date. It is cumulative in the first year and given per year after it. The sentence under the tiles explains why it can differ from the time-weighted return: it counts the timing of the contributions, which the time-weighted return removes.

Drawdown now. The fall of the growth of 100 below its previous peak on the record's date, in per cent. Zero means that the portfolio is at a new peak.

Orders this month. The orders placed in the current month, against the allowance of five set by rule 6. Unlike the record, this tile also counts orders placed after the last net asset values, since the allowance applies to orders as they are placed.

Figure 1 | Allocation now against the target and the band

On the dashboard: Figure 1, in the Overview view

Figure 1 shows the weights of the two sleeves on the record's date against their target weights. Each bar is a sleeve's weight, blue for equity and orange for bonds, and the black tick on each bar is its target weight. The shaded interval behind the equity bar is the band, from 65 to 75 per cent. The bond sleeve has no band of its own, since its weight is what the equity weight and the cash leave. Cash, the contributions not yet invested, is given in the note, and the table lists all three weights.

The figure tells a reader what the next cycle day would do at these weights. When the equity bar ends inside the shaded interval, the rules sell nothing on that day and the top-up alone moves the weights towards their targets. When it ends outside, the next cycle day brings the portfolio back to 70/30 (rule 5). An order placed after the last net asset values enters the figure with the next ones, and the note names it.

Figure 2 | Growth of 100: the portfolio, the reference portfolios and the index blend

On the dashboard: Figure 2, in the Overview view

Figure 2 shows what 100 of the starting amount has become since it was invested, beside the two reference portfolios and the index blend, all with contributions removed. The navy line is the portfolio, labelled with its last value. The dark grey line is reference portfolio A, the same rules without costs, and the mid grey line is reference portfolio B, back to 70/30 on every cycle day. The light grey line is the index blend, 70 per cent MSCI ACWI net total return and 30 per cent the Bloomberg Euro-Aggregate Treasury index, in euro, rebalanced at each month end. The blend is monthly, starts at the portfolio's level on the first month end and is drawn once a second month's returns are published.

Three comparisons sit on this chart, and each answers its own question. The gap between the portfolio and reference portfolio A is the implementation cost, which Figure 7 follows. The gap between reference portfolio A and reference portfolio B is the drift effect, the effect on returns of letting the weights drift inside the band, before costs. The gap between the portfolio and the index blend adds the tracking difference of the two ETFs, each fund's calendar-year return minus its index's, and from the twelfth month end Table 3 splits it into these three parts. Over a record this short, none of the gaps says whether the rules are good ones.

The line under the figure links to the part of the Method view that explains what the portfolio does not do: why it has no optimiser, no index benchmark (an index against which a manager's work is measured), no views and no signal.

Figure 3 | Drawdown against the mandate's limit

On the dashboard: Figure 3, in the Overview view

Figure 3 shows how far the portfolio has been below its previous peak on each day since the first purchase. The filled navy area is the drawdown, drawn downwards, so the deeper the area, the larger the fall. The axis is scaled to the deepest fall so far and reaches at least 10 per cent, and the mandate's two lines, the test limit of 35 per cent and the outer bound of 40 per cent, appear once the axis reaches them. The grey line, from the second month end, is the index blend's drawdown at month ends, and the "as of" line gives the worst fall to date.

The figure and the mandate measure falls differently. The mandate's test used month-end values, which miss falls that reverse within a month, while the figure uses daily values, so the worst daily fall is at least as deep as the worst month-end fall. A drawdown is never a reason to act: under rule 8 a fall in prices neither pauses the cycle nor adds an order.

Rebalancing

On the dashboard: the Rebalancing view

The Rebalancing view shows the equity weight against its band and every order against the rules. Rebalancing means moving the weights back towards the target weights, through the top-ups (rule 4) or through the band (rule 5).

Figure 4 | Equity weight in its band, with every order

On the dashboard: Figure 4, in the Rebalancing view

Figure 4 follows the equity weight day by day since the first purchase and marks each order on the day it was placed. The blue line is the equity weight, the shaded interval is the band from 65 to 75 per cent, and the black line is the target of 70 per cent. A blue dot marks a purchase of the equity ETF and a smaller orange dot a purchase of the bond ETF, drawn smaller so that a day with both purchases shows both. A red diamond labelled "Band order" marks a day on which the band brought the portfolio back to 70/30, and an open orange circle marks an order that the rules called for and that was not placed. The table gives each day's weights of equity, bonds and cash.

Between cycle days the line moves with markets alone, since no order is placed. On a cycle day the top-up moves it towards 70 per cent, and only when the weight has left the band does the cycle day bring it back to 70 exactly. An order placed after the last net asset values enters the figure with the next ones, and the note names it.

Figure 5 | Orders per month against the allowance of five

On the dashboard: Figure 5, in the Rebalancing view

Figure 5 counts the orders placed in each month. Each bar is a month's orders, and the line at five is the allowance of rule 6, the number of orders a month that the account's plan carries without commission. The table lists every order with its date, sleeve, side and the rule that called for it, whether it was placed inside rule 7's window of 15:45 to 17:00 Amsterdam time, and whether the bid and ask were recorded when it was placed.

Apart from the first purchases, the rules call for orders on the cycle day only, at most two, so any other order, or an order outside the window, departs from the rules. The note names each departure, and the departures register of the Log view, Table 6, keeps it with its consequence.

Costs

On the dashboard: the Costs view

The Costs view measures what running the portfolio costs. It separates two kinds of cost: the cost of placing the portfolio's orders, measured against reference portfolio A, and the ETFs' own yearly costs, which are taken inside their net asset values.

Figure 6 | Cost per month by source

On the dashboard: Figure 6, in the Costs view

Figure 6 splits each month's implementation cost into four sources. Each bar is a month, in basis points of the portfolio's value on the day of each order, and its four parts are shades of navy, from the darkest for the price paid against net asset value, through the currency conversion and the half-spread at execution, to the lightest for the commissions. A part below zero is a saving. The table adds the orders without a recorded bid and ask and the month's total.

The four sources have exact meanings. The half-spread is half the gap between the ask and the bid quoted when an order is placed, which is what buying or selling at once costs at that moment. The commissions are the broker's fees, and the currency conversion is the cost of changing euro into another currency, which neither ETF needs, since both are bought in euro, so this part is zero. The price paid against net asset value is the gap between each order's price and the ETF's net asset value of that day, beyond the half-spread. It is computed as the month's implementation cost minus the other three, so the four parts of a month add up to that month's implementation cost in Figure 7. For an order without a recorded bid and ask the half-spread cannot be measured, and it stays inside the price paid against net asset value.

One month's sign says little. An order fills at a price during the day while the net asset value is set at the close, so the price paid against net asset value can be negative as well as positive.

Figure 7 | Implementation cost, cumulative

On the dashboard: Figure 7, in the Costs view

Figure 7 follows the implementation cost since the first purchase, cumulative, in basis points. The line is reference portfolio A's value relative to the portfolio's, with contributions removed, and it is positive when reference portfolio A is worth more, that is, when running the portfolio has cost money. The line at zero marks no cost at all, and the end of the line carries its last value.

The line and Figure 6 are two views of the same cost. A cost of 10 basis points in one month of Figure 6 moves this line up by 10 basis points in that month, so the slope of the line is the cost of each month. The ETFs' yearly costs are not in it, since reference portfolio A pays them too, inside the same net asset values.

Table 1 | What the portfolio costs a year

On the dashboard: Table 1, in the Costs view

Table 1 sets the ETFs' yearly costs beside the implementation cost. For each ETF it gives the ongoing costs, its TER, or total expense ratio, plus its transaction costs, as the ETF's key information document states them, with the document's date and the ETF's weight. The portfolio's row weights the two ETFs by their sleeve weights, and a cash row completes the weights.

The note sets the implementation cost of Figure 7 beside the ongoing costs and explains its sign order by order. The two numbers measure different things: the ongoing costs are what holding the ETFs costs every year, taken inside their values, and the implementation cost is what placing the portfolio's own orders has cost since the first purchase.

Contributions

On the dashboard: the Contributions view

The Contributions view separates the money paid in from the gains and losses on it. Its figures are in units of the starting amount, which counts as 100, so the view shows the portfolio's money without publishing a euro amount.

Figure 8 | Money paid in and the value of the portfolio

On the dashboard: Figure 8, in the Contributions view

Figure 8 compares what has been paid into the portfolio with what it is worth. The starting amount counts as 100 units and each top-up adds 5. The grey line is the money paid in, which steps up with each top-up, the navy line is the portfolio's value, and the shaded area between them is the gain or loss on the money paid in. Unlike Figure 2, this figure keeps the contributions in, since it shows money and not returns.

Figure 9 | Market effect and contribution by month

On the dashboard: Figure 9, in the Contributions view

Figure 9 shows, for each month, how much the portfolio's value changed through markets and how much through new money. The upper panel is the market effect, the month's change in value that did not come from contributions, in per cent of the value at the start of the month. The lower panel is the contribution, the money paid in during the month, in the same unit, and the table gives both by month. Together the two make up the month's change in value. While the portfolio is young, a top-up of 5 per cent of the starting amount can be large against a month's market effect.

Look-through

On the dashboard: the Look-through view

The Look-through view shows what the portfolio owns through its two ETFs, in four parts: the equity sleeve, the bond sleeve, the whole portfolio and the risk of the whole portfolio. Each holding is weighted by its weight in its ETF and by the ETF's weight in the portfolio, so the shares add up across the two sleeves. The "as of" line of each exhibit gives two dates, that of the issuer's holdings file and that of the sleeve weights.

Equity sleeve

On the dashboard: Equity sleeve, in the Look-through view

The first part covers the equity sleeve, the iShares MSCI ACWI UCITS ETF, and its heading gives the sleeve's share of the portfolio on the record's date. iShares publishes the fund's holdings every day, so the holdings and the weights share a date. Every figure in this part gives shares of the whole portfolio, and its table adds the shares of the sleeve.

Figure 10 | Equity by region

On the dashboard: Figure 10, in the Look-through view

Figure 10 groups the equity holdings into MSCI's regions under its market classification: North America, Europe and the Middle East, the Pacific and emerging markets, with the fund's cash and derivatives on their own. Each bar is a region's share of the portfolio. The four regions are also the equity components of the risk estimates in Figure 18 to Figure 20.

Figure 11 | Equity by country: the largest

On the dashboard: Figure 11, in the Look-through view

Figure 11 gives the countries with the largest shares of the portfolio, by the location of each company as iShares reports it. Where the fund holds another iShares ETF that itself holds one market, that holding counts under its market, and the note names those markets.

Figure 12 | Equity by sector

On the dashboard: Figure 12, in the Look-through view

Figure 12 groups the equity holdings by sector under GICS, the industry classification that MSCI and S&P publish, as iShares reports it. The single-market ETFs that the fund holds have no single sector, so they form a bar of their own.

Figure 13 | The ten largest companies

On the dashboard: Figure 13, in the Look-through view

Figure 13 lists the ten companies with the largest shares of the portfolio, with the share classes of one company summed, so that a company with two listed classes of shares counts once. The title gives their combined share of the portfolio, which shows how much of it depends on ten companies.

Bond sleeve

On the dashboard: Bond sleeve, in the Look-through view

The second part covers the bond sleeve, the Vanguard EUR Eurozone Government Bond UCITS ETF, and its heading gives the sleeve's share of the portfolio. Vanguard publishes the fund's holdings and its factsheet once a month, so these exhibits are as of the last month end the issuer has published, while the sleeve weight is the record's.

Figure 14 | Bonds by issuing country

On the dashboard: Figure 14, in the Look-through view

Figure 14 groups the bonds by the government that issued them. Each bar is a country's share of the portfolio, and the table adds its share of the sleeve. The bars follow the size of each government's bond market in the fund's index, so they show which governments the portfolio lends to and in what proportion.

Figure 15 | Bonds by maturity, and duration

On the dashboard: Figure 15, in the Look-through view

Figure 15 groups the bonds by years to maturity, the time until each bond repays, measured from the date of the holdings file. Each bar is a group's share of the bond sleeve, and the grey bar is the cash the fund holds. A longer bond moves more when interest rates change, which the two tiles under the figure measure.

Duration of the bond sleeve. The duration is a bond's average time to its payments, weighted by their value, in years. It is also, approximately, the percentage fall in the bond's price when its yield rises by one percentage point. This tile gives the bond sleeve's average duration from Vanguard's factsheet, so a duration of 7 years means that a rise of one percentage point in yields would lower the sleeve's value by about 7 per cent.

Duration contribution to the portfolio. The sleeve weight times its duration: the same sensitivity to yields, measured for the whole portfolio.

Figure 16 | Credit quality of the bond sleeve

On the dashboard: Figure 16, in the Look-through view

Figure 16 gives the credit ratings of the bond sleeve as Vanguard's factsheet reports them, with each issue rated by the median of its Moody's, Fitch and S&P ratings. A rating grades the issuer's ability to repay, from AAA, the highest, down to BBB, the lowest grade still counted as investment grade. Each bar is a rating's share of the bond sleeve. The fund's index holds investment-grade bonds only, so no bar falls below BBB.

Whole portfolio

On the dashboard: Whole portfolio, in the Look-through view

The third part puts both sleeves and cash together, for the one exposure that both sleeves carry: currency.

Figure 17 | Currency exposure, by sleeve

On the dashboard: Figure 17, in the Look-through view

Figure 17 gives the portfolio's exposure to each currency, each bar split by where the exposure comes from: blue from the equity sleeve, orange from the bond sleeve and grey from cash. Each holding counts in the currency of its market, as the issuers report it, so a US company's shares are an exposure to the US dollar whichever currency they were bought in. Buying the equity ETF in euro on Xetra therefore does not change the exposure. Currencies are named until the rest is below 10 per cent of the portfolio, and the rest is grouped as Other. The share outside the euro is the part of the portfolio whose value in euro moves with exchange rates.

Risk of the whole portfolio

On the dashboard: Risk of the whole portfolio, in the Look-through view

The fourth part estimates the risk of the whole portfolio at its look-through weights. The record is too short to measure risk, so the estimates use index returns over the last 120 months. These index returns price the components, the parts into which the look-through divides the portfolio: MSCI's regional indices in euro for the four equity regions, and three maturity groups of the bond sleeve, each priced as a zero-coupon euro area government bond from the European Central Bank's yield curve. Cash is a component with no return. None of the estimates is part of the record, and none is a limit, since the mandate binds on the worst fall over years.

The risk tiles

Seven tiles give the risk estimates, each with its basis in the line under it.

Estimated volatility. The standard deviation of the portfolio's return over a year, estimated from the components' monthly returns over the last 120 months, at the latest look-through weights.

Estimated worst fall over ten years. The median worst fall over ten years in the bootstrap of the allocation test (Figure 25), at the split nearest the equity weight on the record's date. The line under it adds the 95th percentile, the fall exceeded in one path in twenty.

The forecast's calibration. The standard deviation, over the months that have a forecast, of each month's return divided by the volatility estimated for it at the end of the month before. A value of 1 means that the estimates matched the returns, and a value above 1 that they were too low. It needs twelve months of returns, so it starts at the October 2027 month end. Every estimate is kept as first written, in `outputs/ex_ante_risk.csv`, so it cannot be revised after the event.

Value at risk, historical. The loss over one month that the portfolio, at its weights on the record's date, exceeded in one month in twenty of the 120-month window: the sixth worst of the 120 months.

Expected shortfall, historical. The mean loss of those six worst months, which says how large a loss is once the value at risk is passed.

Value at risk, normal. The same loss read from the normal distribution with the estimated volatility and a mean of zero: 1.645 times the monthly volatility.

Expected shortfall, normal. The normal distribution's mean loss beyond that value at risk.

The historical figures depend on which crises the 120-month window holds, and the note under the tiles names them. Table 2 adds the older crises. Each month end's value at risk is kept as first written, so that the months in which the loss passed it can be counted against the one in twenty expected. Twelve months give that count almost no power, and it starts to say something at about sixty.

Figure 18 | Correlation of the components' monthly returns

On the dashboard: Figure 18, in the Look-through view

Figure 18 gives the correlation of the monthly returns of each pair of components over the 120-month window. Each cell of the grid prints the correlation of its row and its column, coloured from blue for negative through white to red for positive. The table gives each component's weight, its volatility a year and the series that prices it. MSCI publishes no series for Europe and the Middle East together, so that region is priced by MSCI Europe, and the note says so.

The grid shows where diversification comes from. The equity regions tend to move together, and the correlation between equity and bonds decides how much the bond sleeve offsets a fall in equity.

Figure 19 | Each component's share of the estimated variance

On the dashboard: Figure 19, in the Look-through view

Figure 19 sets each component's share of the portfolio's estimated risk beside its weight. The variance is the volatility squared, and a component's risk contribution is its weight times the covariance of its return with the portfolio's return (how far the two move together, in the units of variance), divided by the portfolio's variance, so the contributions sum to 100 per cent. The pale bars are the weights and the navy bars the risk contributions. Equity is far more volatile than government bonds, so its share of the risk is larger than its share of the weight, and the note gives the two shares for the equity regions and for the maturity groups.

Figure 20 | The efficient frontier of the components, and the portfolio's point

On the dashboard: Figure 20, in the Look-through view

Figure 20 places the portfolio among the portfolios that could be built from the components. The efficient frontier is the set of long-only, fully invested portfolios of the components with the lowest estimated volatility for each estimated return. The horizontal axis is the estimated volatility and the vertical axis the estimated return, both in per cent a year. The dark grey curve is the frontier of the 120-month window and the light grey curve that of the 120 months before it. The open circle is the minimum-variance portfolio, the portfolio on the frontier with the lowest estimated volatility, and the navy dot is the portfolio at its look-through weights.

The portfolio's point lies away from the curve because its two ETFs are cap-weighted, that is, they hold each security in proportion to its market value, and no optimiser chose the weights. The distance from the curve is how far the cap weights sat from the frontier estimated on that window. It is no cost that the portfolio paid, since the frontier rests on the window's own estimated mean returns, the least reliable input, and the two curves show how far it moves from one window to the next. The Method view's section on what the portfolio does not do explains why the rules do not follow the curve.

Table 2 | Today's portfolio in five historical episodes

On the dashboard: Table 2, in the Look-through view

Table 2 estimates what the portfolio, at its weights on the record's date, would have lost in five historical crises: 2000 to 2003, 2007 to 2009, the euro crisis of 2011, February to March 2020, and 2022. Each episode holds today's look-through weights from the end of its first month to the end of its last, without rebalancing, so the two sleeves' contributions add up to the total. The table gives each sleeve's part in per cent of the portfolio, the episode's month ends and the series used. The components' series start in 2001 and 2004, so the fall of 2000 to 2003 uses the allocation test's two series at the sleeve weights. These are estimates on index returns at today's weights and form no part of the record, and month-end values miss falls that reverse within a month.

Factor exposures

On the dashboard: the Factor exposures view

The Factor exposures view describes the equity ETF by its factors. A factor is one of the return series in Kenneth French's data library for developed markets: the market's return minus the risk-free rate, the one-month US Treasury bill return in French's files, and the returns of long-short portfolios for size, value, profitability, investment and momentum. A long-short portfolio holds one group of stocks and sells another short, so its return is the difference between the two groups: small companies against large ones for size, companies whose shares are cheap relative to their accounting value against expensive ones for value, profitable against unprofitable ones for profitability, companies that invest little against those that invest heavily for investment, and recent winners against recent losers for momentum.

A loading is the coefficient of the ETF's return minus the risk-free rate on a factor's return, which says how far the ETF moves when the factor moves by one. A fund that holds the market at market weights should load close to 1 on the market and close to 0 on the other five factors.

Figure 21 | Factor loadings of the equity ETF

On the dashboard: Figure 21, in the Factor exposures view

Figure 21 gives the ETF's loading on each factor, estimated on weekly returns over the last three years. A standard error is the estimated standard deviation of an estimate, and each loading carries a horizontal line of two standard errors on either side, so a loading whose line crosses zero cannot be told apart from zero. Weeks run from Wednesday to Wednesday, to avoid effects at the start and the end of a week. The regression uses the ETF's net asset value in US dollars and French's factors, also in US dollars, so no exchange rate enters it, and its standard errors are Newey-West, which allow for returns that are correlated from one week to the next.

The table adds a daily estimate as a check, with each factor's return of the same day and of the day before (Dimson's correction), since the fund's valuation point and French's daily returns may not be timed alike. The bond sleeve is not regressed on equity factors. A market loading close to 1 and small loadings on the rest describe a fund that holds the market at its weights, which is what the equity sleeve is meant to be.

Figure 22 | Rolling loadings

On the dashboard: Figure 22, in the Factor exposures view

Figure 22 repeats the estimate over rolling windows of 36 months of weekly returns, stepped at each month end, in one panel per factor, so that a reader can see whether the loadings have stayed stable. The five panels after the market share one scale, so their lines can be compared directly, and each has a line at zero. Small movements are expected, since 36 months of returns estimate a loading with error. A lasting move away from zero would indicate that the index itself had shifted towards a factor.

Figure 23 | Where the difference to the index blend came from

On the dashboard: Figure 23, in the Factor exposures view

Figure 23 will split the portfolio's monthly return against the index blend into its factor exposures times the factors' returns, and a residual. It starts at the October 2027 month end, when twelve months of returns exist, and until then it holds a line saying so.

Table 3 | The difference to the index blend in three parts

On the dashboard: Table 3, in the Factor exposures view

Table 3 will split the difference between the portfolio and the index blend into three parts that sum to it, over the period of Figure 23: the portfolio against reference portfolio A, which is the implementation cost, the drift effect, which is reference portfolio A against reference portfolio B, and the tracking difference of the two ETFs. It starts at the October 2027 month end as well.

Method

On the dashboard: the Method view

The Method view explains how the split and the two ETFs were chosen and what the portfolio does not do, and it holds the data sources, the controls and the rules. Its exhibits use history, and each says which history.

Figure 24 | The allocation test: growth per year against the worst fall of each split

On the dashboard: Figure 24, in the Method view

Figure 24 shows how the split was chosen. The allocation test, in `mandate/allocation_check.py`, runs every split from 40/60 to 100/0, in steps of five points, through monthly euro returns from February 1999 to December 2025. The equity series combines Kenneth French's developed-market and emerging-market returns, 90 and 10 per cent, converted to euro, and the bond series is a seven-year euro area government bond priced from the European Central Bank's yields. Each split is rebalanced by its band, and for each the test measures the growth per year and the worst fall.

Each grey dot is a split, placed by its worst fall on the horizontal axis and its growth per year on the vertical axis, and the curve joins the splits from 40/60 to 100/0. The navy dot is 70/30, the split held, and the red lines are the risk limit of 35 per cent and the outer bound of 40 per cent. The mandate's rule is the largest equity target weight whose worst fall stayed within 35 per cent, and on this history that is 70/30, with a worst fall of 34.7 per cent from August 2000 to March 2003. With two assets every split lies on this one curve, so the limit picks the point and no estimate of future returns enters. The test is in-sample, that is, computed on the data on which the choice was made: the split was chosen on the returns the chart shows.

Figure 25 | Simulated worst fall of each split over resampled ten-year paths

On the dashboard: Figure 25, in the Method view

Figure 25 asks how often each split could fall more than the limit over ten years if the test's history were reshuffled. The bootstrap, a resampling of the allocation test's returns, builds 10,000 paths of ten years, each joining blocks of twelve consecutive months that start at random months, and applies the band at each month end. A block keeps its twelve months in their order, so a fall and the correlation of the two sleeves within a year stay as they happened.

For each split, a line runs from the median worst fall over the paths, a filled dot, to the 95th percentile, an open circle, which is the fall exceeded in one path in twenty. The black diamond is the allocation test's single historical path, 70/30 is in navy, and the red lines are 35 and 40 per cent. The note gives how often 70/30 and 60/40 fall more than 35 per cent, and the table adds paths of five years and blocks of 6, 24 and 36 months. The paths reuse the same returns, so they say nothing about years unlike those, and the split was chosen on this sample. The results are an input to the review of the rules in October 2027 and change no rule.

Figure 26 | Rolling 36-month correlation of the two sleeves' monthly returns

On the dashboard: Figure 26, in the Method view

Figure 26 follows the correlation between the returns of equity and of bonds over rolling windows of 36 months. The grey line uses the allocation test's index series from 1999, and the navy line the two ETFs' net asset values from March 2019, the first full month of the bond ETF's share class. A

negative correlation means that bonds tended to rise when equity fell, which keeps the worst fall of a 70/30 portfolio below that of equity alone. The years in which the line turned positive are the ones in which bonds did not offset the falls of equity. The test's worst fall and the bootstrap's bands rest on this correlation, which the block resampling keeps within each block.

What the portfolio does not do, and why

On the dashboard: What the portfolio does not do, and why, in the Method view

This section answers the questions readers ask of a portfolio like this one: why it has no optimiser, why it has no index benchmark and why it does not hold the S&P 500 alone, why it takes no views and why it follows no signal. It changes no rule. Each answer is a paragraph with a run-in label, the values in the text are computed from the repository's files, and two figures sit inside the section.

No optimiser. An optimiser chooses weights from estimates of each asset's expected return and risk. The section compares the standard error of the equity series' mean yearly return with the difference in growth per year between the 40/60 and the 100/0 split, and the two are of the same size, so the data cannot tell the splits apart by their returns. It cites the study of DeMiguel, Garlappi and Uppal (2009) on how much history an optimiser needs, and a replication of it.

No index benchmark, and why not the S&P 500. A benchmark is the portfolio a client would hold without a manager, against which the manager's contribution is measured. This portfolio has no client and no manager apart from its rules, so its reference portfolios are the rules themselves without costs. The S&P 500 holds US companies only, which the global fund already holds at about two thirds of its weight, so holding the index alone puts the whole portfolio into one country and one currency.

No views and no signal. A view is a statement that an asset's expected return differs from what its price implies, and a signal is a view produced by a rule and tested on history. The rules take neither, and a signal could enter only through the yearly review of the rules, with its test fixed in advance.

The section's notes, numbered from 1, give the full reference for each source with a link.

Figure 27 | The allocation test on the US stock market

On the dashboard: Figure 27, in the Method view

Figure 27 repeats the allocation test with the US stock market in euro in place of the global equity series, with the same bonds, the same band rule, the same splits and the same months. For each split, the upper bar is the global series, as in Figure 24, with 70/30 in navy and the other splits in grey, and the lower bar, in teal, is the US series. The red lines are 35 and 40 per cent, and the table gives both series' worst falls, their growth per year and whether each split stayed within 35 per cent.

The split was set on the global series alone, so this figure chooses nothing. It shows how a single-country equity sleeve would have fared under the same limit, and the section's second paragraph gives the largest US split within 35 per cent and the US worst fall at 70/30.

Figure 28 | Growth of 100: reference portfolio A and the S&P 500 alternative

On the dashboard: Figure 28, in the Method view

Figure 28 shows the S&P 500 alternative beside reference portfolio A, each as a growth of 100 since the first purchase. The S&P 500 alternative is the same contributions on the same days in one S&P 500 ETF at its net asset value in euro, in fractional units and without costs, starting as reference portfolio A starts. Its fund is the S&P 500 ETF that the rules' own selection criteria pick, named in the caption and ranked in `method/selection_sp500.csv`. The dark grey line is reference portfolio A, the teal line is the alternative, and the table gives the gap between them in points of the growth of 100, with its sign.

Both lines are rule sets without costs, so the comparison is like for like, and the portfolio's own line is left off because reference portfolio A stands for it. The difference between the portfolio and reference portfolio A is the implementation cost, which the Costs view shows. A record of months is far too short for the gap to say anything about either rule set, which is the same limit that applies to every return on the dashboard.

Figure 29 | The ETF selection: tracking difference by fund and year

On the dashboard: Figure 29, in the Method view

Figure 29 shows how the two ETFs were chosen. For each sleeve the funds that meet the criteria of the rules, such as the index tracked, the costs and the fund size, are ranked in an order fixed before they were compared. Equity ETFs rank first by Irish domicile, which lowers the tax the fund pays on US dividends, then by the highest average tracking difference over 2023 to 2025, then by the lowest TER. Bond ETFs rank by the highest average tracking difference, then by the lowest TER. The tracking difference is a fund's calendar-year return minus its index's return, from the issuers' factsheets, so a higher figure means a smaller shortfall against the index.

Each row is a fund, numbered in rank order, and each mark is one calendar year: a circle for 2023, a square for 2024 and a diamond for 2025. The short vertical bar is the three-year average. The held funds are in navy and the others in grey, and hollow marks are funds the account does not offer. The held fund of each sleeve is the first in rank order that the account offers, and the table gives every value, the factsheet dates and which funds are offered and held.

Table 4 | Data sources

On the dashboard: Table 4, in the Method view

Table 4 lists every source the program reads, with the last date each one covers and the day it was downloaded: the issuers' net asset values, holdings, factsheets and key information documents, the European Central Bank's exchange rate and yield curve, MSCI's index levels and Kenneth French's factors. Every source is downloaded when the dashboard is built, cached locally and not redistributed, and only the figures computed from it are published.

Table 5 | Controls

On the dashboard: Table 5, in the Method view

Table 5 lists the checks on the record, each with how often it runs and where its evidence is: the ledger against the broker's statements each month, the ledger against rules 7 and 10 at each build, the cycle report written before the orders, the departures register, the test that no private value

appears in a public file, the fingerprint of the private values, the checks on the public text, and the rebuild of the page from the committed tables.

The rules

On the dashboard: The rules, in the Method view

The rules card shows `rules/RULES.md`, the ten rules as fixed before the first order and tagged `rules-v1.0`, with its table of amendments. Rule 1 sets the target weights of 70 and 30 per cent and rule 2 the starting amount and the top-up. Rule 3 sets the cycle day, rule 4 how the top-up rebalances and rule 5 the band. Rule 6 sets the allowance of five orders a month and rule 7 the window for orders and what the ledger records for each. Rule 8 forbids any order outside the rules, so news adds no purchase and a fall in prices does not pause the cycle. Rule 9 sets how a rule changes, by a dated amendment that applies from the first cycle day at least 30 days after its commit, with the first review in October 2027, and rule 10 how a sale is recorded. The card also explains how the two ETFs were chosen and how the private euro amounts are fixed by a fingerprint without being published.

Other projects

On the dashboard: Other projects, in the Method view

The last card of the Method view links to the author's other investment research projects, each with its page, its dashboard or note, and its code.

Simulations

On the dashboard: the Simulations view

The Simulations view applies the rules to the allocation test's returns and to paths resampled from them, to show how the rules behave over a longer history than the record. Every chart in it is a simulation on index returns, gross of costs, and none is the portfolio's record.

Mechanics simulation

On the dashboard: Mechanics simulation, in the Simulations view

The first part shows how often the rules act. It runs them on the allocation test's returns with a starting amount of 100 and a top-up of 5 a month, with the cycle at each month end in place of the 5th.

Figure 30 | Simulation: the equity weight in its band under the rules

On the dashboard: Figure 30, in the Simulations view

Figure 30 follows the equity weight at each month end from February 1999 to December 2025, before the top-up, with the band and the target. Each top-up buys the sleeve furthest below its target weight, and the band brings the portfolio back to 70/30 when the weight leaves it. The red diamonds mark the months with a band order. The table counts the cycle days with one and with two orders and the band orders, and gives the most orders on one cycle day against the allowance of five. Every purchase counts as an order, whatever its size.

Historical simulation

On the dashboard: Historical simulation, in the Simulations view

The second part runs the whole history as one path, from February 1999 to December 2025.

Figure 31 | Simulation: growth of 100 under the rules

On the dashboard: Figure 31, in the Simulations view

Figure 31 compares the growth of 100 under the rules, the navy line, with the same portfolio taken back to 70/30 on every cycle day, the grey line, with contributions removed. The gap between them is the drift effect over 27 years of history. The simulation is in-sample, since the split was chosen on these returns, and its index series run back before the two ETFs existed. The table gives the year ends, and the months are in the source file.

Figure 32 | Simulation: fall from the previous peak under the rules

On the dashboard: Figure 32, in the Simulations view

Figure 32 shows the same two simulations as falls below the previous peak, month by month, with the test limit of 35 per cent and the outer bound of 40 per cent.

Resampled paths

On the dashboard: Resampled paths, in the Simulations view

The third part carries the bootstrap's paths forward from the first purchase.

Figure 33 | Resampled paths of the portfolio's value over ten years

On the dashboard: Figure 33, in the Simulations view

Figure 33 shows the range of values the portfolio could reach over ten years from the first purchase if the future repeated the test's history, reshuffled in twelve-month blocks. Values are in units of the starting amount, with the top-up of 5 a month and the rules applied at each month end. The pale band holds the middle 90 per cent of the 10,000 paths, from the 5th to the 95th percentile, the darker band the middle half, and the grey line is the median. The dashed line is the money paid in, and the navy line, from the first month end, is the record itself.

The bands say what the test's sample implies and assume nothing beyond it. The record's place inside the range says where markets have taken it and is no evidence about the rules, which a record of this length cannot give. A record below the lower band is a question for the review of the rules in October 2027.

Log

On the dashboard: the Log view

The Log view keeps the written record of each quarter and month in five parts, newest first: the quarterly reports, the monthly notes, the cycle reports, the departures register and the monthly table.

Quarterly reports

On the dashboard: Quarterly reports, in the Log view

One report covers each quarter. Each report is a page, with a PDF printed from it, built once from the tables of the commit it names and never rebuilt, so it stays as issued while the dashboard moves on. A correction goes into the next issue as an erratum. The first report covers 8 October to 31 December 2026 and is published after the cycle day of 5 January 2027.

Monthly notes

On the dashboard: Monthly notes, in the Log view

After each cycle day a note of 120 to 200 words, written by hand, says what happened and why. Its first line gives the cycle day, and a line of figures written by the program gives the day the top-up was recorded, the equity weight before and after the cycle day, whether the band was triggered, the orders placed against those proposed, the implementation cost of the month before and any departure from the rules.

Cycle reports

On the dashboard: Cycle reports, in the Log view

On each cycle day, before any order, the program writes the cycle report: each rule that bears on the proposed orders, its figure and a pass or a fail, and the orders proposed. After the orders, one line compares what was placed with what was proposed, and under each report the dashboard lists that day's orders as the record holds them.

Table 6 | Departures register

On the dashboard: Table 6, in the Log view

Table 6 lists every departure from the rules since the first purchase, that is, every order placed, or left unplaced, other than as the rules require. Each row gives the date, the rule, what happened and its consequence. A departure stays in the record as it happened and is never corrected after the event.

Table 7 | The monthly table

On the dashboard: Table 7, in the Log view

Table 7 holds every figure behind the tiles and charts, by month, in three blocks. The first block gives returns and risk: the month's return, the time-weighted return since the first purchase, the money-weighted return, cumulative and per year, the drawdown, the worst fall to date and, from the twelfth month, the volatility of the last twelve monthly returns. The second gives the weights at the month end, the drift, the orders, the money paid in and the value in units, and the market effect and the contribution. The third gives the costs in basis points: commissions, half-spread, the price

paid against net asset value, and the implementation cost of the month and cumulative. Every tile is a cell of these tables or of `outputs/portfolio_daily.csv`, except that the orders tile also counts orders placed after the last net asset values.

Terms

Portfolio	the two ETFs held under these rules, plus any contribution not yet invested.
Sleeve	the part of the portfolio held in one ETF.
Weight	a sleeve's value divided by the value of the portfolio.
Target weight	the weight the rules aim for, 70 per cent equity and 30 per cent bonds.
Band	the interval from 65 to 75 per cent for the equity weight. When the equity weight is outside it on a cycle day, after the top-up, the portfolio is brought back to its target weights.
Rebalancing	moving the weights back towards the target weights, through top-ups (rule 4) or through the band (rule 5).
Contribution	money paid into the portfolio, meaning the starting amount and each top-up.
Starting amount	the first contribution, fixed before the first order.
Top-up	the fixed monthly contribution, 5 per cent of the starting amount.
Cycle day	the day each month on which the rules are applied.
Trading day	a day on which Xetra, the Frankfurt exchange where both ETFs are listed in euro, is open.
Ledger	the private record of every contribution and every order, kept outside the repository.
Record	what the program computes from the ledger and the prices: holdings, weights, orders and their costs, returns, the look-through and the reference portfolios.
Net asset value	the value of one share of an ETF at the close of a day, as its issuer publishes it.
Look-through	the portfolio's exposure to the securities inside the two ETFs, each ETF's holdings weighted by the ETF's weight in the portfolio.
Reference portfolio	a portfolio computed from the same contributions and prices as the portfolio, without costs. Reference portfolio A follows the same rules, and reference portfolio B goes back to the target weights on every cycle day.
S&P 500 alternative	the same contributions on the same days into one S&P 500 ETF at its net asset value in euro, in fractional units and without costs, starting with the starting amount at the net asset value of the day of the first purchase, exactly as reference portfolio A starts. It stands for the rule "hold the S&P 500 and nothing else" with no costs, so that it is compared with reference portfolio A like for like, two rule sets on the same money.
Index blend	70 per cent MSCI ACWI net total return and 30 per cent the Bloomberg Euro-Aggregate Treasury index, in euro, rebalanced at each month end.
Implementation cost	reference portfolio A's value relative to the portfolio's, with contributions removed, in basis points, positive when the portfolio trails.
Half-spread	half the gap between the ask and the bid quoted when an order is placed.
Basis point	one hundredth of a percentage point.
Split	the pair of target weights, written equity/bonds, for example 70/30.
Risk limit	the worst fall the mandate aims to stay within, about one third, tested as 35 per cent.
Growth per year	the compound annual return, with contributions removed.
Drawdown	the fall in the portfolio's value below its previous peak on a given day, with contributions removed, in per cent.
Worst fall	the largest drawdown, meaning the largest fall in the portfolio's value from a previous peak to a later low, with contributions removed.
Tracking difference	an ETF's return minus its index's return over a calendar year, in percentage points.
Duration	a bond's average time to its payments, weighted by their value, in years. It is also, approximately, the percentage fall in the bond's price when its yield rises by one

	percentage point. For the bond sleeve it is the average duration in the issuer's factsheet.
Factor	one of the return series in Kenneth French's data library for developed markets: the market's return minus the risk-free rate, and the returns of long-short portfolios for size, value, profitability, investment and momentum.
Loading	the coefficient of the ETF's return minus the risk-free rate on a factor's return.
Risk-free rate	the one-month US Treasury bill return in French's files.
Standard error	the estimated standard deviation of an estimate.
Allocation test	the historical test in mandate/allocation_check.py that derived the split from the risk limit, on monthly euro returns from February 1999 to December 2025.
Bootstrap	the resampling of the allocation test's returns in blocks of twelve consecutive months, each starting at a random month, into paths of five and ten years, with the worst fall measured on each path.
Mechanics simulation	the rules applied to the allocation test's returns, with a starting amount of 100 and a top-up of 5 a month, for the counts of their orders.
Resampled paths	the bootstrap's ten-year paths carried forward from the first purchase with the top-up and the rules, in units of the starting amount.
In-sample	computed on the data on which a choice was made, here the returns from 1999 to 2025 on which the split was chosen.
Time-weighted return	the return with contributions removed, chained from day to day, so the timing and the size of the contributions do not affect it.
Market effect	a month's change in the portfolio's value that did not come from contributions, in per cent of the value at the start of the month.
Money-weighted return	the rate at which the contributions, each from its date, grow to the portfolio's value, so the timing of the contributions counts. Cumulative in the first year and per year after it.
Cycle report	the file cycles/YYYY-MM.md written on a cycle day before any order, with each rule that bears on the proposed orders, its figure and a pass or a fail, and one line after the orders that compares them with the proposal.
Monthly note	the note of 120 to 200 words, in notes/YYYY-MM.md, written by hand after each cycle day.
Departure	an order placed, or left unplaced, other than as the rules require. The departures register lists each one with its rule and its consequence.
Rolling correlation	the correlation of two series of monthly returns over the 36 months to each month end.
Component	one of the parts into which the look-through divides the portfolio to estimate its risk: the four equity regions, the bond sleeve's three maturity groups, and cash, which has no return.
Maturity group	the bonds of the bond sleeve with 1 to 5, 5 to 10 or over 10 years to maturity, each priced as a zero-coupon euro area government bond at 3, 7.5 or 20 years from the European Central Bank's yield curve.
Estimated volatility	the standard deviation of the portfolio's return over a year, estimated from the components' monthly returns over the last 120 months at the latest look-through weights.
Realised volatility	the standard deviation of the portfolio's last 12 monthly time-weighted returns, times the square root of 12, so a figure for a year.
Value at risk	the loss over one month that the portfolio, at today's weights, exceeded in one month in twenty of the 120-month window, read from those months or from the normal distribution with the estimated volatility.
Expected shortfall	the mean loss over one month in the months beyond the value at risk, from the window or from the normal distribution.

Forecast's calibration	the standard deviation, over the months that have a forecast, of each month's return divided by the monthly volatility estimated for it at the end of the month before. 1 means the estimates matched the returns, and above 1 that they were too low.
Covariance	how far two returns move together, in the units of variance, the square of volatility: their correlation times the product of their standard deviations.
Risk contribution	a component's weight times the covariance of its return with the portfolio's return, divided by the portfolio's variance. The contributions sum to 100 per cent.
Efficient frontier	the long-only, fully invested portfolios of the components with the lowest estimated volatility for each estimated return, on the means and covariances of one window of 120 months.
Minimum-variance portfolio	the portfolio on the efficient frontier with the lowest estimated volatility.
Cap-weighted	holding each security in proportion to its market value, as the indices of both ETFs do.
TER	total expense ratio, the management and other operating costs an ETF takes from its assets each year, as a percentage of them.
Ongoing costs	what an ETF takes from its assets each year, its TER plus its transaction costs, as its key information document states them. They are inside its net asset value.
Drift effect	reference portfolio A's return minus reference portfolio B's, the effect of letting the weights drift inside the band, before costs.

Carlo Hofer · Data: the issuers' published data · Code on GitHub · Edition of 8 November 2026