THREE HORIZONS CAPITAL / FLOW RESEARCH

The Only Thing That Rotated Was the Narrative

Positioning for a longer conflict: what advisers said, what the money did, and why the standard models cannot price the difference.

Data through 10 August 2026 | All figures in US dollars | Written for professional investors

+$14.0bn
went *into* US government bond funds during the six weeks described as a flight from duration

6 of 6
weeks in the period we measured recorded a net inflow, the quietest still taking $1.17bn

1 of 7
major providers publish a gold assumption you could actually size a hedge with

Since the spring, one description of investor behaviour has become close to consensus. Geopolitical risk has moved from an episodic worry to a standing condition, so investors are buying gold and selling the government bonds that used to be ballast.

It is a coherent story and it explains several things at once. We had written parts of it ourselves. **Then we went back to the underlying records, and found that the advisers had not said it, the money had not done it, and the models most people use could not have priced it if they had.**

> **What this note is.** Not a view on the price of gold, the level of yields, or the course of the conflict. It is an account of what actually moved, what the people advising clients actually said, and what the standard assumption set is actually built to handle. The three agree with each other. The story is the outlier.

What is measured and what is inferred

Four tags run through this note. VERIFIED means we recomputed it ourselves from the underlying records and it is an observed quantity, with no model between the number and the money. MODEL means an attribution model produced it, it carries error, and we show every component including the part the model cannot explain. JUDGEMENT means it is our reading and reasonable people will differ. WORK means it is not a claim about the market at all. It describes something we would do, and it is tagged so that the commercial part of this note is held to the same labelling as the analysis.

01 Nobody sold the ballast

United States government bond funds took in money in every one of the six weeks to 10 August, between **+$1,173.7mm** in the quietest week, ending 20 July, and **+$3,617.5mm** in the largest, ending 6 July. VERIFIED Across the six weeks the total is **+$14.0bn**. All six weeks we measured recorded an inflow.

Three of those six weeks nonetheless carry a negative normalised score, and that is where the story came from. The score measures a week's flow against that category's own recent pattern. A negative reading means money arrived **more slowly than usual**. It does not mean money left. In the three negative weeks, between $1.17bn and $1.96bn arrived.

Every week was an inflow. Three of the six scored negative.

US government bond funds, weekly net flow (top) against the normalised score (bottom). The score measures pace against the category's own recent pattern. Negative means money arrived more slowly than usual, not that money left.

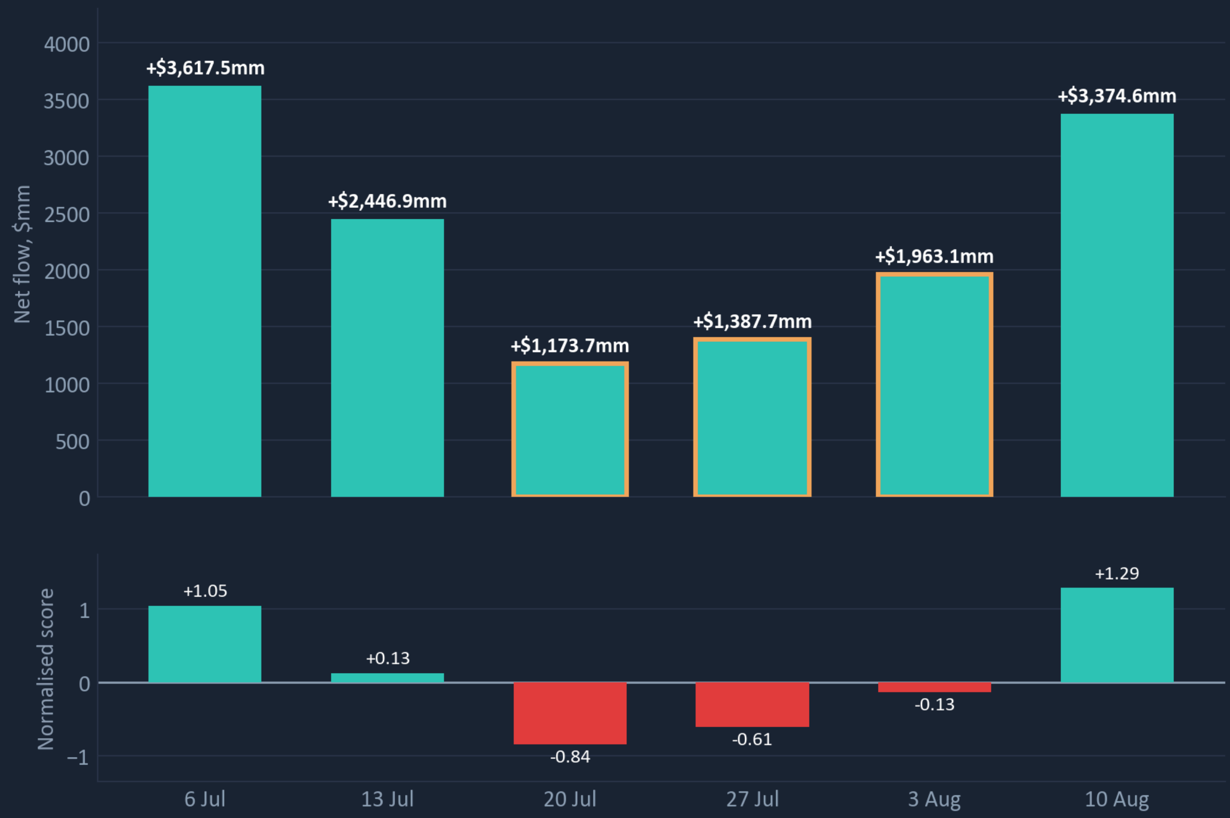

Three Horizons Capital. Weekly fund flow records, data through 10 August 2026. Figures in US dollars. Orange border marks the three weeks scoring negative.

The upper panel is money. The lower panel is pace. They disagree in three of the six weeks, and only one of them is a statement about direction.

A measure of pace was read as a measure of direction. The two point opposite ways in exactly the conditions this summer produced: a category taking large, steady inflows where the rate of arrival slows for a few weeks from an unusually high base.

> **We made this error ourselves.** Our own first pass described the United States as selling duration and reached a recommendation on that basis. It took a re-derivation from the underlying records to catch it. The normalised score is a good measure, used by careful people, and it answers a genuinely useful question. It does not answer this one.

01 Nobody sold the ballast CONTINUED

The composition defence, and why it does not hold

The reasonable objection is that money moved down the curve, out of long duration and into bills, and the category total conceals it. Splitting the same flows by maturity bucket, 15 June to 11 August. VERIFIED

MATURITY	NET FLOW	FUNDS
Ultra short	**+$4,236.7mm**	28
Short	**+$3,953.8mm**	23
Long	**+$3,835.5mm**	27
Broad	**+$2,827.6mm**	29
Intermediate	**+$2,117.6mm**	38
Maturity runoff	**+$542.3mm**	24

All six maturity buckets in our set recorded net inflows over that window. Long maturity is the third largest at $3.84bn across twenty-seven funds, and **+$3,613.1mm** restricting to Treasury exposure alone. **The pattern is the same at the front of the curve and at the back of it.**

SO WHAT

If a positioning case rests on the premise that the market has stopped treating government bonds as ballast, **that premise is not in the data.** It may still be right as a forecast. It is not currently a description, and the two should not be argued as though they were the same thing.

Data through 10 August 2026. Figures in US dollars.

02 What advisers said, and what the money did

If a rotation were underway, the people advising clients would be the first to say so. We hold a panel of **17,934 positioning statements from 2,499 advisers**. 7,858 of them carry a publication date, and the dated 2026 subset used below is **2,096 statements from 523 advisers**. It says something quieter than the coverage does. VERIFIED

Advisers tilted toward hedges and kept the ballast. They tilted rather than rotated.

Dated 2026 statements, by asset group. Counts are STATEMENTS, not advisers, and are not weighted by assets. The adviser count behind each group is shown on the left.

Three Horizons Capital. Adviser commentary panel, dated records only. The dated 2026 subset is 2,096 statements from 523 advisers. Hedges are favoured more emphatically than ballast, but government bonds remain net positive at 35 raising against 19 reducing.

Dated 2026 statements. Hedges are favoured more emphatically than ballast, but ballast remains net positive.

Counts here are statements, not advisers, because an adviser may comment more than once. On gold and precious metals, **43 statements raised against 12 that reduced**, across 30 advisers. On commodities, 18 against 2 across 18 advisers. On defense and aerospace, 9 statements across 8 advisers, all of them raising. The tilt toward hedges is real and it is one-sided.

But on government bonds the panel reads **35 raising against 19 reducing** across 51 advisers. Net positive. Nobody was telling clients to abandon ballast, and the flows show clients not abandoning it. **Said and done agree.** The advisers tilted toward hedges without funding the tilt by selling the ballast leg, and that is precisely what the flow data shows happening.

02 What advisers said, and what the money did CONTINUED

The gold buying is real, but younger and more regional than the story

Over the eight weeks to 3 August, United States precious metals funds averaged **-$17.1mm per week**. VERIFIED Approximately flat. The buying arrives at the very end of the window, **+$1,158.9mm** on 3 August and **+$1,587.9mm** on 10 August. A fortnight, not a quarter.

Europe is the genuine case. European precious metals funds averaged **+$224.3mm per week** of net buying across the same eight weeks, steadily rather than at the end. Decomposed, the two regions are close to mirror images. MODEL

	UNITED STATES	EUROPE
Observed net flow	-$17.1mm/wk	+$224.3mm/wk
Persistent component	-$107.3mm/wk	+$810.7mm/wk
News-driven component	+$1,222.9mm/wk	-$1,379.1mm/wk
Unexplained residual	-$1,132.7mm/wk	+$792.6mm/wk
Residual share of attribution	49.7%	24.8%

The three components sum exactly to observed flow. Read the American column carefully: the news-driven component is large and positive, and almost exactly cancelled by a residual of similar size and opposite sign. Quote the first two rows without the third and you produce a picture of heavy sentiment-led buying. Include the third and observed flow is roughly zero, which is what it was. **We show the residual because the alternative is to publish a number the next line down withdraws.**

Europe's buying persists while its news-driven component is negative, which is the signature of an allocator adding to a strategic position on weakness rather than reacting to a headline. Europe also carries the lowest residual share of the five markets we run, so it is where the attribution guesses least.

The slowest money has not moved at all. Across 1,067 disclosed United States public pension commitments in 2023 to 2026 vintages, a screen for gold, precious metals, mining, defense or aerospace strategies returns 13 matches, all of which resolve to manager names rather than to strategies. VERIFIED The caveat is large and we state it rather than lean on the result: 972 of those 1,067 records, or 91%, carry no asset-class tag, and a fund whose name does not announce its strategy would not be caught. The honest claim is that **no commitment named for these strategies appears in four vintage years**, which is weaker than saying pensions hold none.

SO WHAT

Advisers tilted toward hedges and kept the ballast. The money did the same. Every source we hold agrees with every other source we hold, and **the only thing that rotated was the way the position was described.**

Data through 10 August 2026. Figures in US dollars.

03 What the assumptions are actually built for

If a persistent conflict is the risk, the natural next question is what a standard toolkit does with it. The answer is uncomfortable, and it is a matter of record rather than opinion.

The standard stress library holds **three scenarios**: a normal recession, a hard recession of the 2009 kind, and a bond shock of the 2022 kind. VERIFIED **All three are growth shocks or rate shocks.** The library was built from the last three crises, and each of those was growth-negative or rate-driven. A supply-side shock is a different shape, and the shelf as it stands does not carry one.

The shelf carries three shapes. The one under discussion is a fourth.

The three stress scenarios in the standard set. In both recessions gold and duration rise together. In the bond shock they diverge, but gold barely moves.

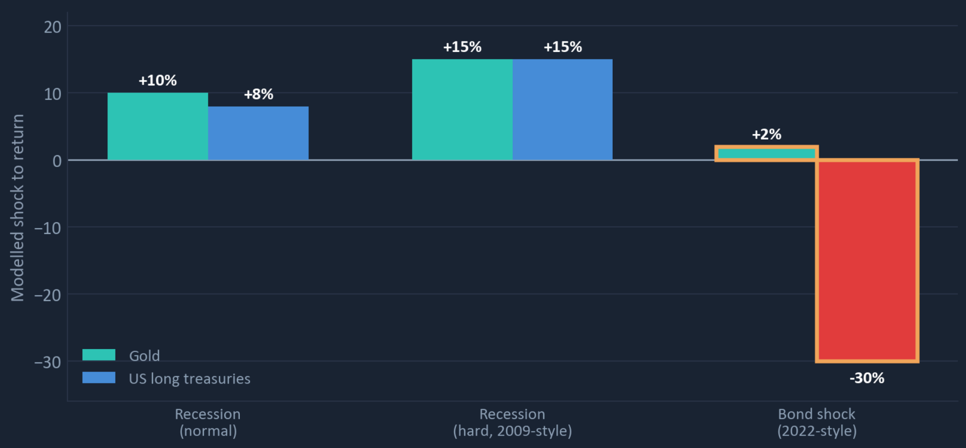

Three Horizons Capital. Standard scenario library, 2026 vintage. Orange border marks the only case where the two diverge. The combination under discussion, gold sharply up and duration sharply down, is a fourth shape.

The library was built from the last three crises, and every one of them was a growth shock or a rate shock. Nothing on the shelf is a supply shock.

Look at what that implies. In the normal recession gold is modelled at **+10%** and long Treasuries at **+8%**. In the hard recession both are at **+15%**. Only the bond shock separates them, and there gold is nearly flat at **+2%** against long Treasuries at **-30%**. **In all three, gold and duration either rise together or gold barely moves.** The combination under discussion, gold sharply higher while duration sells off, is a fourth shape, and the toolkit has not been asked to carry it.

03 What the assumptions are actually built for CONTINUED

And most providers cannot help you size the hedge

Seven providers publish 2026 long-term assumptions in our set. **Three publish a commodities line, a fourth publishes a listed natural resources line, and one publishes a gold assumption.** VERIFIED

PROVIDER COVERAGE, 2026 VINTAGE	COMPOUND RETURN	VOLATILITY
Commodities, provider A	2.3%	14.3%
Commodities, provider B	3.5%	18.0%
Commodities, provider C	3.8%	18.3%
Natural resources, provider D	6.4%	not published
Gold, the only provider publishing one	4.0%	16.7%

The three commodities lines span **2.3% to 3.8%** on return while agreeing on volatility within four points. The providers differ by roughly 150 basis points a year on the asset and barely at all on its risk. Providers are not named here, consistent with the rest of this note.

The fourth line looked like a commodities assumption and is not. It is a listed natural resources line, and our own mapping layer had merged it into the commodities bucket, which is where the wider spread we first computed came from. Separating the two took a re-derivation against raw published names. We mention it because it is the same error as the one on page two, committed by us, one layer further down.

SO WHAT

You cannot stress this risk with the standard scenario set, and for six providers out of seven you cannot price the asset most people are using to hedge it. **The gap is not in anyone's portfolio. It is in the toolkit everyone is using to think about the portfolio.**

Data through 10 August 2026. Figures in US dollars.

04 What would change our reading

We hold no data on the geopolitical calendar and we are not going to speculate about it. What we can do is publish the thresholds at which this reading would change, so that a reader can check them independently and hold us to them. JUDGEMENT

The persistent component of US gold flow turns positive and stays positive for four consecutive weeks

Would mean: the American pattern has become accumulation rather than reaction, and the fortnight-old turn has become a position. This is the single thing we watch most closely.

Government bond flow turns negative in dollars, not merely in the normalised score

Would mean: an actual rotation has begun. Every week in the period examined here recorded an inflow instead.

A defense or aerospace category appears in the fund taxonomy at all

Would mean: the industry has recognised a structural theme rather than a handful of funds gathering assets. Today the closest proxy in the taxonomies we run is the thematic equity bucket.

A second provider publishes a gold assumption, or any provider publishes a conflict scenario

Would mean: the assumption set is catching up with the conversation, and the risk becomes modellable rather than merely discussable.

05 How we read it

Taking the three together, and stating this as a reading of the evidence rather than as advice: JUDGEMENT

The flows say this is not yet a repositioning. One region is accumulating steadily, one turned a fortnight ago, one contradicts the thesis outright, and the slowest money has not moved. That is an early-stage picture, not a consensus.

The assumption set says that if it becomes one, most standard toolkits are not built to price it. Three scenarios, all of them growth or rate shocks, and gold coverage in one provider out of seven.

THE READING

Those two facts together are the reason to look now rather than later. **Not because the trade is obvious, but because the measurement is not.** Anyone acting on this is acting ahead of both the institutional channel and the modelling toolkit, and that is a defensible thing to do deliberately and an uncomfortable thing to discover afterwards.

06 What this looks like against a book

Everything above is a reading of the market. The same three tests run against a single portfolio answer a narrower and more useful question, and each one falls directly out of a finding in this note. WORK

Separate the dollars from the pace, in your own reporting

The first finding here is that a widely used measure answers a different question from the one being asked of it. We would run the same decomposition across the categories a portfolio actually holds, net flow and normalised score side by side, and mark every position where the two point in opposite directions. Those are the positions where a reader of your own reporting is most likely to reach the conclusion we corrected on page two. It took a re-derivation to catch it here, and nothing about that is specific to this market.

Price the shock the shelf does not carry

The standard library holds three scenarios and all three are growth shocks or rate shocks. We would build the fourth shape, a supply-side shock, and run it against an allocation across all eight assumption sets we hold, seven of them on the same 2026 vintage used above, rather than the single set a toolkit ships with. The disagreement between providers is the point rather than a caveat: the three providers publishing a commodities line span 2.3% to 3.8% on return while agreeing on volatility within four points, so a portfolio priced off any one of them is carrying roughly 150 basis points a year of assumption risk it cannot currently see.

Find the product before the taxonomy names it

The closest thing to a defense or aerospace bucket in the taxonomies we run is the thematic equity bucket. For an allocator that is a measurement problem. For a manager it is a product question, because a real bid with no category name is where a shelf gap sits. We can show which adjacent buckets are taking the money, at what pace, how that differs across the five markets in this note, and what the adviser panel is saying about each of them.

The last time you were shown evidence that the market had moved, was it a measure of how much money moved, or a measure of how quickly it arrived relative to normal?

Those two answers point in opposite directions more often than is comfortable. Any of the three above is the same work as this note, pointed at your book instead of the market. We will run the first of them on a single portfolio at no cost, because the fastest way to show what the distinction is worth is to find it in something you already hold.

07 Limits and method

The figures cover exchange-traded fund flows in five markets, a dated adviser commentary panel, disclosed United States public pension commitments, and published long-term assumptions. They exclude separately managed accounts, direct holdings, unlisted vehicles and all private wealth structures. We hold no dataset that tags anything as a family office, so nothing here should be read as a family-office survey.

The provider comparison in section 03 is drawn from a multi-provider assumption store. The portfolio-analytics library our own tools run on is a separate, single-provider table, so a reader should not infer that any optimiser sees all seven sets. The comparison is a comparison, not the input to a model.

Counts in the adviser section are statements, not advisers, because an adviser may comment more than once, and adviser totals are taken on a resolved identifier rather than on the raw name, which carries spelling variants. Asset groups come from a priority-ordered crosswalk over free-text labels: gold or precious, then defense or aerospace, then commodities, then government bonds, treasuries, sovereigns or gilts, one group per statement. Two of the four groups did not reproduce against the figures first drafted and have been restated to what the crosswalk returns.

The adviser panel carries a publication date on 43.8% of records, so the dated 2026 subset used above is a fraction of the whole. Counts are of statements, not weighted by assets, and the per-group adviser counts range from 11 to 56, which is small.

We examined the news-theme attribution behind precious metals flow, hoping to say something about why the money moved. The records duplicate, the robustness flag is unpopulated, and the ordering is dominated by unrelated themes. Earlier work of ours characterised those drivers more confidently, in both directions. We are withdrawing that characterisation rather than picking whichever version suits the argument. We can see what money did. We cannot currently say why.

Canadian government bond and thematic series have not updated since 23 February 2026. Australian data has not updated since 29 June 2026. Neither supports a current read and none is drawn.

The attribution model is fitted jointly across five geographies. Across 55 evaluation windows its out-of-sample explanatory power averages 0.54, ranging from 0.18 to 0.86. It is a short-horizon cross-sectional signal, not a forecast, and at individual fund level its reliability is far lower and is not used. It was last refitted on 11 May 2026 and applied forward from there.

Every figure was recomputed from source for this publication. Where a figure originated in earlier internal analysis it was re-derived rather than carried forward, and two claims from that analysis did not survive and have been corrected or withdrawn here.