Reflex without Persistence: The Dollar's Multi-Day Safe Haven Reversal in 2025

Orhan Torul

October 2, 2026

Abstract

A safe haven currency appreciates when risk rises. Using sixteen years of high-frequency cross-asset futures data, I study the U.S. dollar on 465 risk-off days, defined as days on which S&P 500 futures fell by more than 1%. I separate the dollar's safe haven behavior into an intraday *reflex*, the five-minute comovement that lifts the dollar as equities fall within the trading session, and its *persistence* over the following week, measured by the cumulative dollar return. The two diverged in 2025. The reflex survived, but the dollar's five-day haven beta against the euro was positive in 2024–2026 ($\beta_{h=5} = +0.75$), with the reversal concentrated in 2025: deeper equity declines were followed by more dollar depreciation, while German Bunds and the rival reserve currencies rallied. Resampling whole stress episodes confirms the break from earlier years but cannot exclude zero for the 2024–2026 level, which rests on few episodes. The reversal appears mainly after unscheduled shocks, and the EUR/USD option-implied skew repriced dollar tail risk over the same window. A simple model of contestable reserve status organizes these facts. The effect fades in early 2026: it is a stress episode, not a proven regime change.

Keywords: Safe haven currencies; U.S. dollar; Flight to safety; Reserve currencies; Foreign exchange options; High-frequency data
JEL: F31, F33, G15, G12, E44

1 Introduction

In April 2025, as U.S. equities sold off on an escalating trade conflict, the two assets that investors hold to hedge equity risk, long-term Treasuries and the dollar, failed to rally over the days that followed. Figure 1 makes the contrast concrete. When Standard & Poor's downgraded the United States in August 2011, long Treasuries rallied as equities fell, the textbook flight to quality. In April 2025, Treasury prices drifted lower alongside a similar equity decline, while the dollar still firmed within the session before weakening over the following days. Some terminology clarifies what did and did not fail. A *safe asset* preserves value and supplies liquidity services; a *hedge* comoves negatively with risky assets on average; a *safe haven* appreciates conditionally, precisely when risk spikes; and a *reserve asset* is what official and institutional portfolios hold as their store of safety. These properties usually go together

for the dollar, but they are distinct, and in 2025 they diverged. Treasuries in Figure 1 stayed broadly stable while equities fell, which is consistent with intact *safe asset* status, yet they did not rally, which is a failure of the *safe haven* function; indeed, Du *et al.* (2026) document that the privileges of the dollar and of Treasuries can decouple. This paper measures the haven function, that is, the conditional appreciation. Episodes such as April 2025 reopened a pointed question for markets: is the dollar losing the safe asset status that underpins the "exorbitant privilege" of the United States (Gourinchas and Rey, 2007; Krishnamurthy and Vissing-Jorgensen, 2012; Jiang *et al.*, 2021)? The worry is not new. Even U.S. Treasuries saw their convenience value erode at the height of the "dash for cash" in March 2020 (He *et al.*, 2022), and a fast-growing literature now reads the market response of April 2025 as investors beginning to question the dollar's reserve role (Jiang *et al.*, 2025; Acharya and Laarits, 2026). That diagnosis, however, typically rests on a single episode, read off daily prices or an event window around one announcement. This paper brings sixteen years of high-frequency cross-asset futures data to the question, 465 risk-off days against which 2025 can be benchmarked, and the answer is more specific: the dollar's safe haven property did not fail as a whole. Its intraday *reflex* survived, while its multi-day *persistence* reversed in 2025.

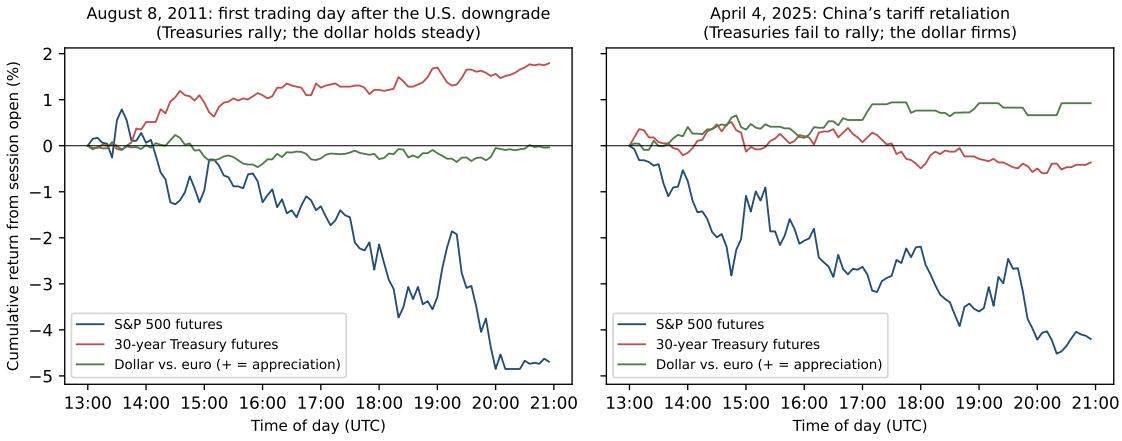


Figure 1: **The motivating contrast: the haven bid in 2011 and in 2025.** Cumulative intraday returns (percent from the U.S. session open) of S&P 500 futures, 30-year Treasury futures, and the dollar against the euro (positive values denote appreciation) on two stress days tied to U.S. policy: August 8, 2011, the first trading day after the S&P sovereign downgrade, and April 4, 2025, the day China announced its tariff retaliation. *Left:* equities fall sharply while Treasuries rally, the classic flight to quality, and the dollar holds steady. *Right:* equities fall by a similar amount while Treasuries drift lower rather than rallying. Treasuries nonetheless remain far more stable than equities, consistent with intact safe asset status even as the safe haven rally fails. The dollar still firms within the session, the intraday reflex documented in this paper, and its failure comes at the multi-day horizon: over April 1–21 the dollar lost 6.5% against the euro (Jiang *et al.*, 2025). The figure is an illustration, not an estimate. The systematic analysis below finds a large multi-day reversal for the dollar, whereas the multi-day response of Treasuries is statistically indeterminate.

The distinction matters for markets and for policy. The dollar's safe asset status lets the

United States borrow cheaply in its own currency and run persistent external deficits; its erosion would raise the cost of capital for the world's largest debtor and expose U.S. policy to a market discipline it has long been spared. That status shows up as a conditional comovement in which the dollar appreciates when global risk rises. The relationship has tightened since the global financial crisis (Lilley *et al.*, 2022) and reflects the dollar's dominant place in global safe asset portfolios (Maggiori *et al.*, 2020). Yet the debate on de-dollarization and the literature on safe haven currencies both treat that comovement as a single object. Most studies measure it at a daily or lower frequency (Habib and Stracca, 2012; De Bock and de Carvalho Filho, 2015), and even the high-frequency work that documents haven effects at horizons from hours to days (Ranaldo and Söderlind, 2010) treats them as a single phenomenon observed at different horizons. I show that the comovement has two components that can behave differently, and that in 2025 they took opposite signs within the same stress episodes.

The two objects are worth defining at the outset, because the paper's argument rests on the distinction. The *reflex* is the conditional comovement of five-minute dollar and equity returns within the U.S. trading session on stress days: whether the dollar ticks up in the same minutes in which equities tick down. I interpret it as reflecting mainly the demand for liquidity: in the first minutes of a stress episode, investors need to place cash immediately, and dollar instruments are the deepest and cheapest place to do so. The *persistence* is a different object at a different horizon: the local projection response of the dollar's cumulative return, from the close before the shock to five trading days after it, to the size of the day's equity decline. This longer response may also reflect portfolio reallocation: it shows whether the dollar keeps the safety bid once investors have had days rather than minutes to decide where to hold safe assets. The reflex turns out to be intermittent rather than constant (it lapsed in the mid-2010s); the persistence is the economically more important margin, and it is the one that reversed in 2025. Separating the two requires high-frequency measurement at explicit horizons, which is the paper's methodological contribution: a single daily beta blends the two components into an object that can look stable even as its economically relevant component reverses. Throughout, "risk-off day" is shorthand for the paper's selection rule based on equity returns, a day on which S&P 500 futures fall by more than 1%, and not an independently identified shock to global risk appetite; Section 5.6 shows that the reversal is present across thresholds and remains positive, though much smaller, under a definition of stress based on the VIX.

I study the 465 risk-off days at two frequencies, measuring the dollar against the euro in the headline results and against a basket and a six-currency cross-section alongside. *Intraday*, the dollar's safe haven reflex is intact in the recent windows: the dollar still rises in the same minutes in which equities fall, through 2026. *Over multi-day horizons*, by contrast, a local projection (Jordà, 2005) shows the dollar's haven beta turning positive in 2024–2026 (66 risk-off days), an effect carried by a pronounced stress episode in 2025: the deeper the day's equity decline, the more the dollar depreciated over the following week, in contrast with

the weakly hedging slope of earlier years. Over the same window, German Bunds rallied, consistent with haven returns shifting away from the dollar. The point estimate is specific to risk-off days, positive against every major currency but largest against the rival reserve havens, similar across thresholds and in non-overlapping windows, little changed by controls for high-frequency monetary policy surprises, and positive when any single episode is dropped or when the sample is restricted to risk-off days of U.S. origin. Its statistical precision depends on the unit of inference. Newey–West standard errors place the 2024–2026 slope far from zero, but a more conservative bootstrap that resamples whole stress episodes cannot exclude zero for the level of the slope, whether it is measured against the euro, against any other single currency, or against a basket, because the 66 days fall into only 35 episode blocks and one of them, the tariff turmoil of spring 2025, carries much of the estimate. The same bootstrap does exclude zero for the change in the dollar's response relative to 2010–2023, which is the paper's formal test, and for the reversal on unscheduled days; for the Bund rally it narrowly includes zero. When risk-off days are split by trigger, the reversal is concentrated on unscheduled days, and on the few days when the shock was a scheduled macroeconomic release (an FOMC decision, the payrolls report, or the CPI, on the actual release dates), the estimate is close to zero. The difference between the two groups is imprecisely estimated, however, and unscheduled days include foreign, geopolitical, and financial events as well as U.S. policy news, so the split is consistent with a credibility interpretation without establishing it. The qualification I stress throughout is that the effect is concentrated in 2025 and does not carry into the partial sample for early 2026: it is a stress episode, not a proven regime change.

This paper makes two contributions. First, *conceptually*, it separates the safe haven property into a high-frequency reflex and a multi-day persistence, two components that conventional daily measurement runs together, and shows that the two can diverge; the novelty is the decomposition and what it reveals, not the local projection itself, which is standard. Section 3 embeds the decomposition in a simple model of contestable reserve status that yields one testable prediction for each fact. Second, *empirically*, it documents a large reversal of the dollar's multi-day haven beta in 2025, together with a rotation of haven returns toward German Bunds and rival reserve currencies. The point estimates across currencies are ordered by reserve rival status, which distinguishes an erosion specific to the dollar from a generic breakdown of correlations, and the option-implied risk reversal skew repriced dollar tail risk over the same window, which corroborates the returns evidence with an independent, forward-looking instrument. The paper also separates what the evidence establishes from what it does not: the break in the dollar's response is statistically robust, but the level of the reversal is estimated from few independent episodes, and the reversal is a stress episode of 2025 rather than a demonstrated permanent erosion of exorbitant privilege.

The remainder of the paper proceeds as follows. Section 2 situates the contribution in four strands of the literature on international finance and market microstructure, and in the new work on the 2025 episode itself. Section 3 sets out a stylized framework of contestable reserve

status and its three predictions. Section 4 describes the high-frequency data, the risk-off sample, and the intraday and local projection estimators. Section 5 presents the results: the intraday reflex, the multi-day reversal, the cross-asset rotation, whether the reversal is specific to the dollar, the option-implied repricing of dollar tail risk, the robustness checks, and the mechanisms together with the caveat about 2025. Section 6 concludes.

2 Related literature

The paper speaks to four established literatures that have, for the most part, measured the dollar's safe asset status at a single frequency, and to a rapidly growing fifth literature on the 2025 episode itself.

A first strand explains why the United States earns a return differential on its external balance sheet: it issues the world's reserve safe asset and is paid a convenience premium for doing so (Gourinchas and Rey, 2007; Krishnamurthy and Vissing-Jorgensen, 2012). Jiang *et al.* (2021) formalize the dollar's safe asset status as a convenience yield that moves with global risk, and Maggiori *et al.* (2020) document the dollar's dominant place in global safe asset portfolios. The status is not exogenous. He *et al.* (2019) model which sovereign's debt becomes the safe asset as an equilibrium outcome of fundamentals and investor coordination, Caballero *et al.* (2017) stress the chronic global shortage of such assets, and Farhi and Maggiori (2018) and Gopinath and Stein (2021) show that reserve currency status is contestable across issuers. In these accounts the privilege is a slow-moving but not immutable structural feature. The question this paper asks is whether its market footprint, the dollar's conditional bid in stress, can move at a much higher frequency, and whether the intraday and multi-day components of that bid can part ways.

A second strand studies what happens to safe assets precisely when safety is most in demand. He *et al.* (2022) show that even U.S. Treasuries can lose their convenience premium in acute stress, as in the dash for cash of March 2020, and Du *et al.* (2018) document violations of covered interest parity that price a dollar funding premium. These papers motivate the central object here, the behavior of the haven bid *within* stress episodes, and the distinction I draw between a liquidity reflex that survives the first minutes and a portfolio reallocation that plays out over days.

A third strand documents the conditional comovement that operationalizes the dollar's haven status: the dollar appreciates when global risk rises, a relationship that Lilley *et al.* (2022) show has tightened since the global financial crisis. This comovement is the empirical handle of the present paper, but I argue that it is not a single object. The conditional response measured at the daily horizon mixes an intraday reflex with a multi-day persistence, and the paper shows that the two can be measured separately and that they behaved differently in 2025.

A fourth strand identifies which currencies hedge global equity risk and how. Because it is

the strand closest to this paper, the relation between the two deserves precision. Ranaldo and Söderlind (2010) study high-frequency exchange rates over 1993–2008 and show that the safe haven appreciation of the franc and the yen materializes at horizons from a few hours to several days; the safe haven fact at multiple horizons is theirs. Kaul and Sapp (2006) document safe haven trading of the dollar itself in intraday and order flow data around Y2K. At daily and lower frequencies, Habib and Stracca (2012) ask which fundamentals make a safe haven currency in a broad panel; De Bock and de Carvalho Filho (2015) show that in risk-off episodes the yen, the franc, and, of particular relevance here, the U.S. dollar appreciate systematically; Grisse and Nitschka (2015) document time variation in the safe haven behavior of the franc across currency pairs; Fatum and Yamamoto (2016) show that the ranking *within* the group of haven currencies is itself unstable, with the franc the safest haven before the global financial crisis, the yen during and after it, and the dollar only third; and Sakurai (2021) documents that the comovement between safe haven assets and U.S. equities in bear markets shifted after that crisis. Relative to this strand, the contribution of the present paper is not high-frequency measurement per se, which Ranaldo and Söderlind (2010) pioneered, but the demonstration that the intraday and multi-day components of one currency's haven response can take *opposite signs* within the same stress episodes. This divergence, dated to 2025 against a sixteen-year benchmark, is what measurement horizon by horizon is designed to detect and what a single daily beta conceals. The cross-section of reserve rivals in Section 5.4 extends the question of the ranking within the haven group, studied by Fatum and Yamamoto (2016) and De Bock and de Carvalho Filho (2015), to the stress of 2025.

Finally, a contemporaneous literature studies the 2025 episode directly. Jiang *et al.* (2025) read the response of bond and currency markets to the April tariff announcement (a 6.5% depreciation of the dollar against the euro between April 1 and 21, despite a *rise* of 48 basis points in the 10-year U.S.–German yield spread) as investors questioning the dollar's reserve status: the Treasury basis turned negative over the same window and, scaled by their earlier estimates, they infer a decline in the convenience yield on dollar safe assets large enough to explain most of the depreciation. Acharya and Laarits (2026) decompose Treasury yields around the same shock and find that the hedging property of the convenience yield eroded at long maturities while short-dated instruments kept theirs. Du *et al.* (2026) argue that the privileges of the dollar and of Treasuries have decoupled, with the convenience yield on Treasuries and the dollar's exchange rate premium no longer moving together, a decoupling consistent with this paper's finding that the large multi-day reversal is the dollar's while the response of Treasuries is statistically indeterminate. Boissay and Huang (2025) find no wholesale exodus from U.S. assets in the portfolio flows of global investors, but diversification at the margin toward euro area sovereign bonds; and the ECB's financial stability assessment shows that in April and May 2025 the dollar's comovement with a broad safe haven factor switched sign and that of U.S. Treasury yields weakened, while the franc

and the euro comoved with the factor more strongly than before (European Central Bank, 2025). These studies share a focus on the episode itself. This paper adds a benchmark and a decomposition: sixteen years of measurement within stress episodes, against which 2025 can be quantified as a break; a split into reflex and persistence that locates *which* component of the haven property failed; a split by shock type that shows *when* it failed; corroboration from an independent instrument in the options market; and a partial window for 2026 in which the effect fades, which is the evidence that separates a stress episode from a regime change.

Methodologically, the paper applies the local projection of Jordà (2005) to a cross-asset setting conditioned on risk-off days, tracing the haven beta as an explicit function of the horizon rather than reading it off a single daily regression. The horizon profile is itself the object of interest: it reveals that the dollar's intraday haven bid is given back by the same day's close and then reverses over the following week.

3 A stylized framework: contestable reserve status

The three facts this paper documents, namely an intraday reflex that survives, a multi-day persistence that reverses, and a reversal that is largest against reserve rivals and after unscheduled shocks, can be organized by a simple model in which the dollar's reserve status is *contestable* and safety is sought in two stages. The model centers on one candidate mechanism, a loss of credibility, but it does not identify that mechanism. Its purpose is to show that the three facts *would be* the joint signature of that mechanism and to make explicit the prediction that the split by shock type in Section 5.7 examines.

A unit mass of global investors holds a budget B of safe assets to carry into a risk-off state. There are N candidate reserve assets, $i = 1, \ldots, N$; asset 1 is the U.S. dollar (safe claims denominated in dollars), and the others are rival reserve assets (the Bund, the franc, and so on). Each asset has a *perceived safety* $s_i \in (0,1)$: the probability, in investors' eyes, that it holds its value in the bad state, that is, its credibility or convenience as a safe asset (Krishnamurthy and Vissing-Jorgensen, 2012; Jiang *et al.*, 2021; He *et al.*, 2019). Safety is contestable in the sense of Farhi and Maggiori (2018): the s_i are not fixed, and a shock can move them.

Safety is sought in two stages. At *impact* ($t = 0$, the first minutes), investors must park cash in the deepest market, and a fraction $\lambda > 0$ of the flight is absorbed mechanically by the most liquid asset, the dollar. At the *reallocation* horizon ($t = 1$, the following days), investors rebalance the budget B across reserve assets according to perceived safety, with portfolio shares

$$w_i = \frac{s_i^\theta}{\sum_{j=1}^{N} s_j^\theta}, \qquad \theta > 0, \tag{1}$$

where θ governs how substitutable the reserve assets are. The dollar's exchange rate e_1 (an increase denotes appreciation) clears the market for dollar safe assets. With imperfectly elastic financiers (Gabaix and Maggiori, 2015), e_1 moves with net demand for dollar claims,

$de_1 = \Gamma^{-1}\, d(\text{dollar demand})$, where $\Gamma > 0$ is the financiers' risk-bearing capacity.

A risk-off episode is one of two kinds. A *global risk* shock (a data surprise or a foreign crisis) raises the demand for safety, $dB > 0$, but leaves relative credibility $\{s_i\}$ unchanged. A *credibility* shock, meaning bad news about the United States itself, lowers the dollar's perceived safety, $ds_1 < 0$, with B unchanged. The first leaves the portfolio shares in (1) fixed; the second shifts them away from the dollar.

The comparative statics of (1) are standard: $\partial w_1/\partial s_1 > 0$, and $\partial w_j/\partial s_1 = -(\theta s_1^{\theta-1}/S)\, w_j < 0$ with $S \equiv \sum_k s_k^\theta$, so the outflow to rival j is proportional to its own share w_j. They deliver three predictions.

P1 *(The reflex does not depend on the type of shock.)* At impact, dollar demand rises by λ times the flight, so $de_1|_{t=0} > 0$ after a risk-off shock of *either* type: the intraday bid is a liquidity reflex and does not depend on s_1.

P2 *(Persistence depends on the type of shock.)* At the reallocation horizon, dollar demand is $B\, w_1$, so $de_1|_{t=1} \propto\ w_1\, dB + B\,(\partial w_1/\partial s_1)\, ds_1$. A global risk shock ($ds_1 = 0,\ dB > 0$) *appreciates* the dollar, and persistence holds. A credibility shock ($ds_1 < 0$) makes the second term negative, and for a large enough loss of credibility the dollar *depreciates*: persistence reverses.

P3 *(Rotation runs to the closest rivals.)* The outflow from the dollar, $-B\, dw_1$, is absorbed by rival j in proportion to $|\partial w_j/\partial s_1| \propto\ w_j$, which is increasing in s_j. Every rival with a positive share gains, but the gains are largest for the *safest* rivals (the Bund, the franc, the euro) and small for risk currencies with low s. The haven loss is therefore *ordered by reserve rival status*. The model's assets are safe claims denominated in national currencies; gold, although central banks hold it as an official reserve asset, lies outside this set by assumption, and the model makes no prediction for it.

The three predictions correspond to the three facts. P1 is the surviving intraday reflex of Section 5.1, with one clarification: the model treats λ as a parameter of market structure, not a constant of nature, so the content of P1 is that the impact bid does not depend on the dollar's perceived safety s_1, not that its strength never varies. The reflex can lapse for reasons the model does not address, and it did in the mid-2010s (Figure 2); what P1 rules out is a reflex that fails *because* s_1 falls, and the reflex is in fact present throughout 2024–2026, the very window in which the persistence margin suggests that the dollar's safety came into question. P2 is the conditional reversal of Sections 5.2 and 5.7: the multi-day persistence reverses, and the reversal is concentrated on unscheduled shocks, the group that contains the documented credibility events, while the point estimate on days of scheduled macroeconomic releases is close to zero (Table 9). This is the pattern that the distinction between $ds_1 < 0$ and $dB > 0$ predicts, although with few scheduled release days the difference between the two groups is not precisely estimated. P3 is the gradient across reserve rivals and the Bund rotation

of Section 5.4; gold lies outside the model, so its small and insignificant response in euro terms (Section 5.3) is a separate fact rather than a test of P3. The paper does not measure s_1 directly, and the 2025 episode identifies a fall in s_1 only if one reads the tariff turmoil as a credibility shock rather than a one-off event. The Treasury basis evidence of Jiang *et al.* (2025), a measured decline over the April window in the convenience yield that foreign investors attach to dollar safe assets, is an independent market gauge of precisely this object, and it moved in the direction that the credibility reading requires.

4 Data and empirical strategy

I use continuous CME futures from Databento at one-minute and daily frequency from June 7, 2010, to June 26, 2026, which follow the contract nearest to expiry and are not adjusted at rolls: the E-mini S&P 500 (equities); the euro, Swiss franc, yen, sterling, Canadian dollar, and Australian dollar futures (the dollar against the euro, against an equally weighted basket of the first five, and against each currency); Treasury futures from the 2-year note to the Ultra T-bond, with the 30-year contract as the benchmark; and gold. Futures are the right instrument for this question. They trade nearly around the clock, they are the venue in which the flight to quality in the first minutes actually clears, and they allow the same cross-asset response to be measured at the five-minute and multi-day horizons within one consistent set of instruments. The continuous series have two known drawbacks. First, the contract nearest to expiry is not always the most traded one: for the products with monthly serial contracts (the currencies other than the franc, and gold) it is at times thinly traded, and for every product the series holds the expiring contract until the last trading day, after most trading has moved to the next one. Second, each roll enters one daily return as the price gap between contracts. Appendix C therefore re-estimates the main results on series that follow the most traded contract and are adjusted at every roll, and on Federal Reserve spot exchange rates. The conclusions are unchanged; the one estimate that does not survive is noted where it appears. Ten-year sovereign yields for the United States, Germany, Canada, and Japan come from FRED, the ECB's AAA euro area yield curve (a proxy for the Bund in which German bonds carry the largest weight), the Bank of Canada, and Japan's Ministry of Finance.

A risk-off day is a trading day on which S&P 500 futures fall by more than 1% ($N = 465$). Trading days are the daily observations of the futures calendar, which includes the Sunday evening session (Appendix A). Of the 465 risk-off days, 429 fall on weekdays and have intraday minute data (the other 36 are Sunday evening sessions, which contain no U.S. cash session), 428 have intraday dollar quotes, and 455 have a complete window from the prior close to $t+5$ for the multi-day projection. The 1% threshold is a convention, not a knife edge: Section 5.6 re-estimates the main result for thresholds from -0.5% to -2% and, as an entirely different definition, on days of VIX spikes; the reversal is present throughout, although it is much

smaller under the VIX definition.

The dollar return is appreciation, $-\Delta \ln(\text{EUR/USD})$, and the basket return is the equally weighted average of the corresponding dollar appreciations, each $-\Delta \ln(\text{foreign/USD})$, over days with all five quotes. Two conventions deserve emphasis so that every later result can be read without ambiguity. First, unless a table or figure states otherwise, "the dollar" means the dollar against the *euro*, the headline pair, and each results section, table, and figure below states its measure explicitly; the basket and the full six-currency cross-section appear in Sections 5.4 and 5.6. Second, the basket spans the euro, franc, yen, sterling, and Canadian dollar, the five constituents of the standard dollar indices for major currencies that have liquid CME futures (the ICE dollar index adds only the Swedish krona, whose futures are thinly traded). The Australian dollar is deliberately excluded from this benchmark basket. It is a commodity-linked risk currency and serves as the counterexample in the cross-section of Section 5.4, but nothing rides on the exclusion: Section 5.6 re-estimates the reversal on a six-currency basket that adds it. In the results on cross-asset rotation, I express each asset as a comparable percentage return so that magnitudes can be read against one another: the dollar as appreciation, gold as a log return, and bonds as duration-adjusted price returns (modified duration ≈ 8.5 applied to $-\Delta y$ for the U.S. and German 10-year bonds; for the United States I also use 30-year Treasury futures returns directly). Appendix A collects the exact variable definitions, the construction of episode blocks, and the classification of risk-off days by origin.

The object of interest is the haven beta β in the conditional comovement

$$r_t^j = \alpha + \beta\, r_t^{\text{eq}} + \varepsilon_t, \tag{2}$$

estimated on risk-off days, where r^j is the return on candidate haven asset j and r^{eq} is the equity shock. A negative β means that the asset hedges: it rises when equities fall. Throughout, β measures a conditional comovement on risk-off days, not a causally identified shock to haven demand; the design isolates a *change* in the dollar's conditional response across periods, which is the object of interest.

For the high-frequency reflex, I estimate (2) on five-minute returns within a fixed window from 13:30 to 21:00 UTC, which contains the U.S. cash session under both daylight saving and standard time, with standard errors clustered by day. Because the fixed UTC window picks up an extra hour before the open or after the close depending on daylight saving time, and because the cash open itself moves between 13:30 and 14:30 UTC across the year, Section 5.1 re-estimates every reflex regression on the exact New York cash session, 09:30 to 16:00 Eastern time; the estimates are nearly identical. Returns are winsorized at the 0.1st and 99.9th percentiles to limit the influence of erroneous ticks. This recovers the automatic bid for dollars within the session: whether, in the minutes in which equities fall, the dollar rises.

For the multi-day margin, I run a local projection (Jordà, 2005) of the cumulative return

from the prior close to $t+h$ on the equity shock of day t,

$$r^{j}_{t-1\to t+h} = \alpha_h + \beta_h\, r^{\mathrm{eq}}_{t} + \varepsilon_{t,h}, \qquad h = 0, 1, \ldots, 10, \tag{3}$$

a specification with a single shock and no added controls. The horizon profile $\{\beta_h\}$ is the central object: it shows whether the haven bid present intraday survives, decays, or reverses as portfolios are reallocated over the following days.

Two features of the design complicate inference: the projection is estimated on a subsample of risk-off days, and the windows of nearby risk-off days overlap. At $h = 5$, the dependent variable runs from the close before day t to the close of day $t+5$, so it spans six daily returns, those of the shock day and of the five trading days that follow; two risk-off days that lie five or fewer trading days apart therefore share at least one daily return. I use two devices to guard against the resulting dependence. The first is an episode block bootstrap. Risk-off days are grouped into episode blocks, with a new block starting whenever a risk-off day falls more than five trading days after the previous one, so that days with overlapping windows always share a block and windows in different blocks share no daily return. The bootstrap resamples whole blocks (50,000 draws, percentile intervals); Appendix B reports the resulting intervals for each principal estimate, the distribution of block sizes, the most influential episodes, and a comparison with blocks defined by calendar days. The second device re-estimates the projection on non-overlapping windows, selecting risk-off days in date order so that consecutive selected days are more than five trading days apart. Newey–West (HAC) standard errors, with ten lags at $h = 5$, are reported alongside. Local projection residuals are serially correlated at horizons beyond zero (Montiel Olea and Plagborg-Møller, 2021); here the dependence arises both from overlapping windows and from the clustering of risk-off days into stress episodes, which HAC standard errors capture only imperfectly.

The two approaches do not always agree. For the levels of the 2024–2026 slopes, the HAC t-statistics are large, whereas the episode bootstrap, which treats each stress episode rather than each day as the unit of independent information, yields 95% intervals that include zero for the dollar against every individual currency and against both baskets. The bootstrap intervals exclude zero for the change in the dollar's response between 2010–2023 and 2024–2026 (the formal break test of Section 5.6) and for the reversal on unscheduled days; for the Bund rally the interval narrowly includes zero. I regard the bootstrap as the more reliable guide for the 2024–2026 levels and state throughout which criterion a claim rests on. Boldface in the tables marks significance at the 5% level by the reported HAC t-statistic (clustered by day for the intraday estimates), and the text reports the bootstrap interval wherever the two criteria diverge. I report the formal structural break test, the range of estimates when one episode is left out, and placebo break cutoffs alongside the point estimates, so that the strength of the evidence and its fragility are visible together.

I do not claim a causal shock to haven demand. Equation (2) is a conditional comovement,

and the rate channel controls in Section 5.6 are themselves endogenous to the shock, so they *condition on* rather than *rule out* a monetary channel. What the design does isolate is a *change* in the dollar's conditional multi-day response between the pre-2024 sample and 2024–2026, and the robustness checks are built to show that this change is not an artifact of thresholds, windows, single episodes, or pooling.

5 Results

5.1 The intraday reflex survives

Table 1 reports the dollar's five-minute haven beta against the euro on risk-off days, by period. The beta is negative and significant in the early 2010s and in every window since 2021 (-0.052, $t = -4.3$, in 2024–2026), and controlling for the contemporaneous Treasury move changes little, so the reflex is not driven by a mechanical rate channel. It was absent in the mid-2010s, so the reflex is intermittent: reliably present in the recent stress windows but not an unconditional constant. Excluding March 2020 leaves the beta for 2016–2020 at +0.003, so the lapse is not an artifact of the dash for cash. The dollar still rises as equities fall intraday in the recent windows, even as its multi-day response reverses.

Two measurement checks discipline the window and the timing. First, because the fixed UTC window straddles the shifts to and from daylight saving time, Table 1 also reports every estimate on the exact New York cash session (09:30 to 16:00 Eastern time). The betas are nearly identical (2024–2026: -0.041, $t = -4.6$; orthogonalized, -0.046, $t = -5.0$), so the reflex is not an artifact of the clock convention. Second, since the baseline regression is contemporaneous at the five-minute frequency, I also add three five-minute lags of the equity return. In 2024–2026 the coefficient on the current bar is -0.051, while the three lag coefficients sum to +0.008. The lags are jointly different from zero at the 5% level ($p = 0.047$), but their combined effect is about a sixth of the contemporaneous coefficient and of the opposite sign. The accurate statement is therefore that the dollar rises *in the same minutes* in which equities fall, not with a measurable delay, and I use that language throughout. A third check concerns data quality. The continuous euro series at times follows a thinly traded contract (Section 4), and, on 158 of the 429 days, more than half of its five-minute returns are zero; on the days with full euro coverage, the reflex is stronger, not weaker (-0.070, $t = -3.4$, in 2024–2026; Appendix C), so stale quotes attenuate the estimate rather than create it.

The intermittency is worth examining year by year rather than only in period averages, and the period buckets need justification, since any four-way split of sixteen years is somewhat arbitrary. Figure 2 plots the intraday haven beta for each calendar year with 95% confidence intervals clustered by day; the yearly estimates, not the buckets, are the primary evidence, and the buckets are descriptive summaries whose boundaries carry no inferential weight. The beta is deeply negative in the early 2010s, positive in a mid-2010s patch (2014–2017),

Table 1: Intraday (five-minute) dollar haven beta on risk-off days

	2010–2015	2016–2020	2021–2023	2024–2026
Dollar five-minute β on equity	**−0.117**	+0.003	**−0.093**	**−0.052**
t (clustered by day)	(−4.7)	(0.3)	(−7.4)	(−4.3)
Orthogonalized to the 30-year Treasury move	**−0.124**	−0.011	**−0.092**	**−0.057**
t (clustered by day)	(−5.2)	(−1.5)	(−7.6)	(−4.7)
New York cash session (09:30–16:00 ET)	**−0.116**	+0.001	**−0.094**	**−0.041**
t (clustered by day)	(−4.5)	(0.1)	(−7.2)	(−4.6)

Notes: Five-minute returns of the dollar against the euro regressed on the equity return across risk-off days, winsorized at the 0.1st and 99.9th percentiles to limit the influence of erroneous ticks; $\beta < 0$ means the dollar rises when equities fall. Boldface marks coefficients significant at the 5% level by the reported t-statistic, here and in all subsequent tables; Appendix B reports the more conservative episode block bootstrap intervals for the principal multi-day estimates. The first two panels use the fixed 13:30–21:00 UTC window, which contains the U.S. cash session under both daylight saving and standard time; the second controls for the contemporaneous 30-year Treasury return, isolating the dollar bid from the rate channel; the third re-estimates the simple specification on the exact New York cash session, which removes the drift of the fixed UTC window with daylight saving time (orthogonalized, the 2024–2026 New York beta is −0.046, $t = -5.0$). The panel is unbalanced because some days have gaps in the five-minute quotes; the dollar regressions use the 426 of the 429 weekday risk-off days that have both dollar and 30-year quotes (37,167 five-minute observations in the UTC window).

and reliably negative again from 2021 through 2026. The confidence intervals show that the sign flip of the mid-2010s is not sampling noise: the betas for 2014–2016 are individually positive and significant, and pooled over 2014–2016 the beta is +0.131 ($t = 7.9$). A sup-Wald test for a single unknown break in the intraday beta, run over every candidate break date in the central 70% of the sample, selects December 2013, the onset of that patch, with a statistic far beyond the 5% critical value, so the lapse of the mid-2010s is a regime dated by the data rather than an artifact of where the table draws its period lines. The sign flip also has economic content. Controlling for the contemporaneous Treasury move shrinks the pooled 2014–2016 beta by roughly a quarter, to +0.094 ($t = 5.0$), so part of the period in which the dollar stopped hedging reflects a rate channel. In those years of tapering and liftoff, punctuated by the devaluation of the renminbi in 2015 and stress in emerging markets, equity declines tended to push back the expected path of U.S. policy, weakening the dollar through rates even as risk rose. The remainder is consistent with the dominant risk-off shocks of that period being news about global growth, for which the dollar was not the natural hedge. The qualitative message stands: the reflex is a property the dollar *usually but not always* displays in stress, and it is present, not absent, in exactly the recent windows over which its multi-day persistence reverses.

5.2 The multi-day persistence reverses

Figure 3 plots the dollar's haven beta against the euro by horizon. In each pre-2024 subperiod, it is weakly negative at every horizon: the safety bid holds, if loosely, for days. In 2024–2026 it is positive and *grows*: $\beta_h = +0.10, +0.21, +0.27, +0.36, +0.37, +0.75$ for $h = 0, \ldots, 5$ (HAC $t = 5.5$ at $h = 5$). The intraday haven bid is already given back by the same day's close (the daily response at $h = 0$ is an insignificant +0.10, which reconciles the negative intraday beta with

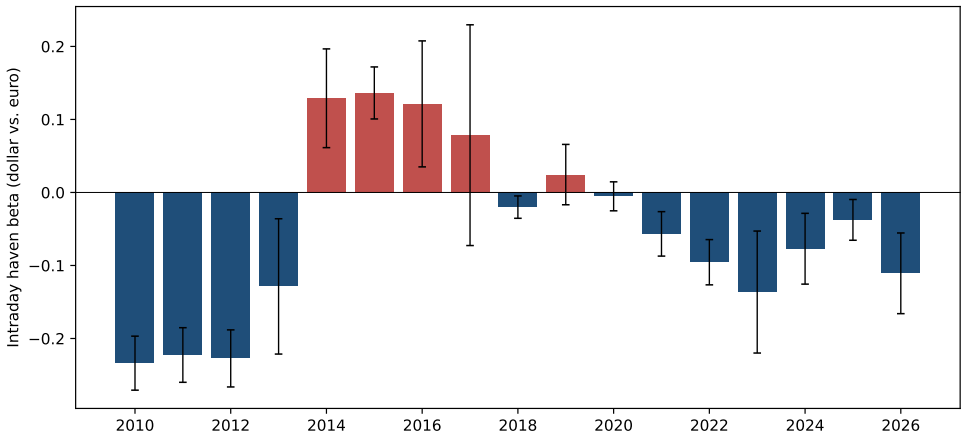


Figure 2: **The intraday reflex is intermittent, and present in the recent windows.** Intraday (five-minute) haven beta of the dollar against the euro on risk-off days, by calendar year (negative values mean that the dollar hedges, rising when equities fall); whiskers are 95% confidence intervals clustered by day. The reflex is deeply negative in the early 2010s, turns positive in a mid-2010s patch (significantly so in 2014–2016), and is reliably negative again from 2021 through 2026. This heterogeneity across years underlies the period averages in Table 1, whose bucket boundaries are descriptive only; a sup-Wald test dates the single largest break in the intraday beta to December 2013 (see text).

the positive daily one), and the dollar then weakens further: each additional percentage point of equity decline is followed by about 0.75 percentage points of cumulative dollar *depreciation* over the next five trading days, against a pre-2024 slope near zero. The selloff shows no sign of mean reversion by $h = 10$ (+0.91), so it is not a transient overshoot that unwinds within two weeks. Inference that treats whole stress episodes as the unit of observation should be read alongside the point estimate. The episode block bootstrap 95% interval for the level of the beta at $h = 5$ is $[-0.01, +0.91]$, with 97.4% of draws positive, so at the episode level the 2024–2026 slope is not distinguishable from zero at the 5% level, if only narrowly. The *change* relative to the pre-2024 slope, by contrast, is significant: the formal break test of Section 5.6, which uses all 455 risk-off days rather than the subsample of 66 days, delivers a bootstrap interval that excludes zero. Appendix B collects the intervals for all principal estimates.

The magnitudes are economically meaningful but not extreme, which is consistent with the reading of a stress episode that I defend below. The reversal shows up in average returns as well as in the slope: the average risk-off day was followed by a dollar *appreciation* of 0.28% over the next five trading days before 2024 and by a *depreciation* of 0.31% in 2024–2026, with deeper shocks bringing proportionally more. A benchmark puts these numbers in perspective. Over the full sample, the dollar's mean absolute daily move against the euro is 0.35% and its mean absolute five-day move is 0.88% (0.29% and 0.71% in 2024–2026). The average drift after a risk-off day in 2024–2026 is thus a little under half of a typical five-day move in that period, while the slope is large relative to it: an equity decline of 2% implies about 1.5 percentage points of conditional five-day depreciation, roughly twice the typical

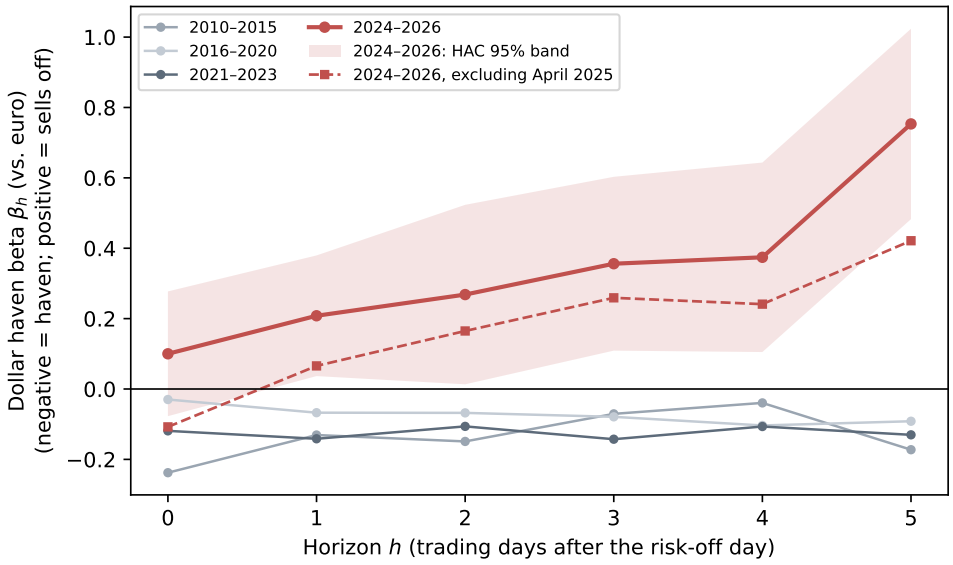


Figure 3: **The dollar's safe haven bid reversed over multi-day horizons in 2024–2026.** Haven beta β_h of the dollar against the euro (local projection of the cumulative dollar return on the equity shock of day t, risk-off days) by horizon h and period; the shaded region is the HAC 95% confidence band (ten lags) for the 2024–2026 path, which is narrower than the episode block bootstrap interval reported in Appendix B. Negative values denote a haven; positive values mean that the dollar sells off over the following days. Only the 2024–2026 profile turns positive; the dashed line excludes the April 2025 tariff window, is positive from $h = 1$ onward, and at $h = 5$ is about half the full estimate. Figure 4 reports the same exercise for the dollar against the franc.

absolute five-day move. These are the moves of a haven that temporarily stopped working at the multi-day horizon, not of a currency in disorderly retreat.

The intraday reflex and the multi-day reversal seem to pull in opposite directions; they are reconciled once one asks *when* the multi-day move accrues. I decompose the dollar's cumulative response, from the prior close to the close at t+5, into three non-overlapping clock segments and project each on the *same* equity shock of day t, so that the three betas sum exactly to the total: (1) the session window of the shock day (the fixed 13:30–21:00 UTC window in which the reflex is measured, which contains the U.S. cash session); (2) the rest of that day outside the window (the overnight hours and those after settlement); and (3) the following five trading days, taken as whole days, since the decomposition does not split them by time of day. One caution before the numbers: this is an accounting of timing, which locates when the move is booked, and not an identification of separate causal mechanisms operating in each segment; the discussion of mechanisms is deferred to Section 5.7. Unlike the rest of the paper, this exercise uses the dollar against the Swiss *franc* rather than the euro, for a data reason. Computing a return for the session of each day requires a continuous series of one-minute prices, and, on the risk-off days of 2024–2026, the continuous euro series is too sparse for that purpose (a median of 92 of the 480 one-minute bars in the 13:00–21:00 UTC

window from which each risk-off day is downloaded, because the series then often follows a thinly traded serial contract), while the franc, which has no serial contracts, is well covered (a median of 464). The franc is also the reserve rival with the largest multi-day haven loss in the cross-section of Section 5.4, which makes it the conservative choice: if even the gains of the strongest rival fall outside the session window of the shock day, the point holds a fortiori. To keep the franc results comparable with the headline euro results, Figure 4 reproduces the full horizon profile of Figure 3 for the dollar against the franc; against the franc the daily $\beta_{h=5}$ is +1.11, of which the subsample with intraday data recovers +1.01.

Table 2 reports the split. In 2024–2026 the estimated contribution of the shock day's session window is small and statistically insignificant (+0.04, $t = 0.5$), and most of the estimated response accrues outside it: in the rest of the shock day (+0.21, $t = 3.1$) and over the following five trading days (+0.76, $t = 4.1$). The last component includes the U.S. cash sessions of those days as well as their other hours, which the decomposition does not separate, so the evidence places the multi-day loss outside the session in which the shock occurred, not outside U.S. trading hours in general. Before 2024 the dollar was a mild multi-day haven (−0.23), and the change is concentrated in the two segments outside the shock day's session window, which move from near zero or negative values to positive ones while the session segment stays small throughout. This is the temporal anatomy of "reflex without persistence": the window of the shock day in which the high-frequency bid operates contributes little to the multi-day loss, most of which accrues outside it. Consistent with this, the high-frequency reflex of Section 5.1 remains intact against the euro in 2024–2026 (the −0.052 of Table 1) even though the cumulative move within the session window is small. Against the franc, the strongest rival, even the five-minute reflex has faded to a slightly positive +0.029 ($t = 1.8$), so the gradient of Section 5.4 reaches into the intraday frequency. The euro decomposition, restricted to the days on which it is well covered, is underpowered ($n = 18$) but points the same way: the session segment carries no dollar weakness (a residual bid of −0.08), and the positive sign sits in the following days.

5.3 Where the haven returns go: the cross-asset rotation

If the dollar ceded its multi-day haven role, the natural question is where the haven returns went. Table 3 and Figure 5 report the five-day response of each asset to a U.S. risk-off shock *in comparable return units*, contrasting the pooled pre-2024 sample with 2024–2026. These are return responses, not measured capital flows; I read them as the price footprint of where safety was sought, not as an accounting of flows. The flow evidence that exists points the same way: in the portfolio flows of global investors over the same window, Boissay and Huang (2025) find no exodus from U.S. assets (the April sales largely reversed within two months) but a tilt toward euro area sovereign bonds.

The largest signal specific to the United States is the dollar. Its response moves from a

Table 2: When the multi-day dollar move accrues (dollar against the franc, β on the equity shock of day t)

Clock segment	2010–2023	2024–2026
(1) Shock day, session window (13:30–21:00 UTC)	+0.01 (0.3)	+0.04 (0.5)
(2) Same day, outside the window (overnight and after settlement)	−0.02 (−0.7)	**+0.21** (3.1)
(3) Following five trading days, all hours (close $t \to t$+5)	**−0.22** (−2.5)	**+0.76** (4.1)
= Total ($\beta_{h=5}$)	**−0.23** (−2.3)	**+1.01** (4.3)

Notes: Additive decomposition of the dollar's cumulative five-day appreciation against the Swiss franc (prior close to the close at t+5) into three non-overlapping clock segments, each a local projection on the same U.S. equity shock of day t over risk-off days; $\beta > 0$ means that the dollar *depreciates* (haven lost). Segment (1) is the move within the fixed 13:30–21:00 UTC window of the shock day, from the one-minute futures; the window contains the U.S. cash session under both daylight saving and standard time but is not identical to it. Segments (2) and (3) are residuals against the daily projections at h=0 and h=5, so the three sum to the total by construction, and the informative content is the split; segment (3) covers all hours of the following five trading days. HAC t-statistics (ten lags) in parentheses; boldface marks coefficients significant at the 5% level by that criterion; n = 365 (2010–2023) and 62 (2024–2026) risk-off days with intraday coverage of the franc. As for the other level estimates of 2024–2026, the episode block bootstrap interval for the total includes zero: on the 62 days of this table (33 blocks, 50,000 draws, as in Appendix B) it is [−0.34, +1.28], against [−0.38, +1.38] for the franc row of Table B.1, which uses all 66 days. The decomposition informs the timing of the move rather than the significance of its level. The euro is the paper's headline numeraire, but in 2024–2026 its continuous intraday series is too sparse for a session return on most days; the franc, the strongest reserve rival (Table 4), is the conservative substitute and is well covered. Measuring segment (1) on the exact New York cash session (09:30–16:00 ET) leaves the split unchanged (2024–2026: within the session +0.01, t = 0.1; outside the session +0.24, t = 3.7; following days +0.76, t = 4.1).

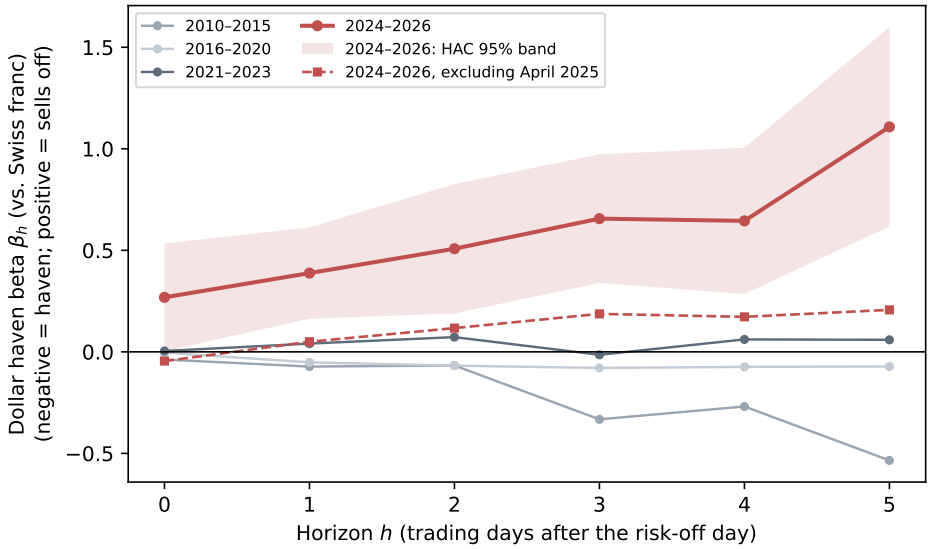


Figure 4: **The multi-day reversal against the Swiss franc.** Haven beta β_h of the dollar against the franc by horizon and period, constructed exactly as in Figure 3; the shaded region is the HAC 95% band (ten lags) for 2024–2026, and the dashed line excludes April 2025. The franc profile mirrors the euro profile with larger magnitudes ($\beta_{h=5}$ = +1.11 against +0.75), which is why the clock decomposition of Table 2, which requires the franc's intraday coverage, can be read against the headline euro results.

statistically insignificant −0.09 (HAC $t = -1.3$) before 2024 to a selloff of +0.75 (HAC $t = 5.5$) in 2024–2026. German Bunds *rally* (−0.23, $t = -5.4$), having shown no significant haven response before.[1] At the episode level, both estimates are borderline: the bootstrap intervals for the Bund, [−0.49, +0.01], and for the level of the dollar response, [−0.01, +0.91], each narrowly include zero, whereas the change in the dollar's response from the earlier period is significant under either criterion (Section 5.6 and Appendix B). The U.S. Treasury response is positive but insignificant whether measured by 30-year futures (+0.58, $t = 1.4$) or in duration-adjusted 10-year terms (+0.26, $t = 1.0$); the data are uninformative about whether long Treasuries kept their multi-day haven role. This point deserves emphasis, because the motivating Figure 1 shows Treasuries failing in 2025, but a single episode is not a systematic estimate. Gold rallies sharply in dollar terms (−0.92, $t = -2.0$), but this is largely the mechanical mirror image of dollar depreciation: in euro terms the response is small and insignificant (−0.15, $t = -0.4$), as it is in yen and Swiss franc terms (−0.29 and +0.22, both insignificant). The cross-asset evidence is thus a shift in haven *returns* away from the dollar and toward German Bunds. It is a statement about prices, consistent with portfolio reallocation, and it does not establish that the same investors moved between the two assets. Once the translation effect in gold is set aside, the dollar response is the largest in the table.

Table 3: Five-day response to a U.S. risk-off shock, in comparable return units

	Dollar (appreciation)	U.S. Treasury (30-year futures)	German Bund (duration-adjusted)	Gold (USD)	Gold (EUR)
2010–2023	−0.09	−0.07	+0.09	+0.35	+0.13
HAC t	(−1.3)	(−0.5)	(1.0)	(1.1)	(0.5)
2024–2026	**+0.75**	+0.58	**−0.23**	**−0.92**	−0.15
HAC t	(5.5)	(1.4)	(−5.4)	(−2.0)	(−0.4)

Notes: Local projection of each asset's cumulative five-day return on the equity shock of day t, risk-off days; the earlier period pools 2010–2023 to mirror the formal break test. All responses are in comparable percentage return units: the dollar as appreciation against the euro, gold as a log return (in USD and in EUR), and bonds as duration-adjusted price returns (the United States via 30-year Treasury futures). Positive values mean that the asset sells off when equities fall (haven loss); negative values mean that it rallies (haven gain). Boldface marks coefficients significant at the 5% level by the HAC t-statistic; under the episode block bootstrap of Appendix B, the 2024–2026 interval narrowly includes zero for both the Bund and the dollar. The responses of the Treasury and of gold in euro terms are insignificant in both periods; the rally of gold in dollar terms is largely the mirror image of dollar depreciation.

5.4 Is the reversal specific to the dollar?

A natural concern is that the reversal of 2024–2026 reflects not an erosion specific to the dollar but a generic breakdown of diversification, in which a U.S. equity selloff drags *everything* down together. Two cuts of the data speak against that reading: which currencies the dollar loses ground to, and whether foreign sovereign bonds behave differently from Treasuries.

[1] The ECB AAA curve is a proxy for the Bund. On the risk-off days of 2025–2026 for which Eurex Euro-Bund futures are available in the Databento feed (from March 2025), the futures return corroborates the rally ($\beta = -0.22$, $t = -2.4$, against −0.17, $t = -2.5$, for the proxy on the same 36 days).

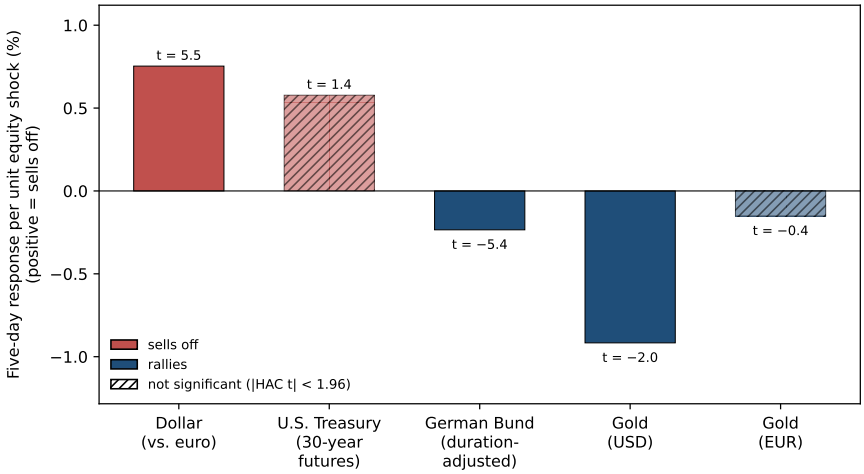


Figure 5: **The rotation in 2024–2026, in comparable return units.** Five-day response to a U.S. risk-off shock, as a percentage return per unit of equity shock (negative values denote rallies; the dollar is measured against the euro). The dollar sells off (red) and German Bunds rally (blue); the responses of the Treasury and of gold in euro terms (hatched) are insignificant by the HAC criterion. The rally of gold in dollar terms is largely the mirror image of dollar depreciation.

Table 4 estimates the dollar's haven beta against each of six currencies, separating rival reserve and haven currencies (the Swiss franc, euro, and yen) from risk and other currencies (sterling and the Canadian and Australian dollars). The grouping follows the literature on safe haven currencies rather than any construction of mine. The franc and the yen are the two currencies that this literature consistently identifies as safe havens (Ranaldo and Söderlind, 2010; De Bock and de Carvalho Filho, 2015; Grisse and Nitschka, 2015; Fatum and Yamamoto, 2016), and the euro is the principal reserve alternative to the dollar in official portfolios; the Canadian and Australian dollars are the commodity-linked, cyclical currencies that the same literature finds depreciate in risk-off states, and sterling, while a secondary reserve currency by official holdings, behaves cyclically rather than as a haven in these classifications. The results do not hinge on the borderline assignment of sterling, since Table 4 reports every currency separately. The pattern is not a uniform breakdown. In 2024–2026 the point estimate of the dollar's multi-day haven loss is largest against the rival havens ($\beta_{h=5} = +1.11$ against the Swiss franc, +0.75 against the euro, and +0.62 against the yen, with HAC t-statistics of 4.4, 5.5, and 4.1) and smallest against the pure risk currency, the Australian dollar (+0.21, $t = 0.9$), with sterling and the Canadian dollar in between (both near +0.3, with HAC t-statistics of 2.2 and 2.9). Figure 6 shows the ordering: the gradient runs with reserve rival status. Two qualifications keep the claim in proportion. First, by the HAC criterion, the dollar weakened significantly against five of the six currencies at the five-day horizon, so the accurate claim is not that the loss is *confined* to reserve rivals but that it is *ordered* by reserve rival status. Second, the ordering is a property of the point estimates. Under the episode block bootstrap (Table B.1), the 95% interval for every individual currency pair includes zero and the intervals

overlap, so the cross-section establishes a pattern rather than a set of individually significant level effects.

The contemporaneous (same-day) betas in Panel A show where that pattern is most pronounced. Before 2024 the dollar's same-day beta was negative against the risk currencies and close to zero against the franc and the yen: pooled over 2010–2023, it averages −0.16 across sterling and the Canadian and Australian dollars and −0.02 across the three rivals. By 2024–2026 the same-day beta has *flipped sign* against the rival havens (the dollar now depreciates against the Swiss franc, +0.27, $t = 2.0$, and the yen, +0.16, $t = 1.8$, the latter significant only at the 10% level), while against the risk currencies it is close to zero for sterling (+0.00) and the Canadian dollar (−0.02) and negative for the Australian dollar (−0.34, $t = -2.4$). Averaged within groups, the same-day beta moved from −0.02 to +0.18 against reserve rivals but only from −0.16 to −0.12 against risk currencies: the loss of the same-day bid is concentrated where the dollar competes with other reserve havens rather than with cyclical, commodity-linked currencies. The ECB's assessment of the April episode records a similar asymmetry from a different vantage point: the dollar's comovement with a broad safe haven factor switched sign, while the franc and the euro comoved with it more strongly than before (European Central Bank, 2025).

Table 4: Dollar haven beta by currency and period: reserve rivals versus risk currencies

	Reserve rival or haven			Risk or other		
	CHF	EUR	JPY	GBP	CAD	AUD
Panel A: contemporaneous (same-day) $\beta_{h=0}$						
2010–2015	−0.04	−0.24	−0.04	−0.17	−0.33	−0.52
2016–2020	−0.00	−0.03	+0.01	−0.08	−0.09	−0.12
2021–2023	+0.00	−0.12	−0.05	−0.30	−0.25	−0.41
2024–2026	**+0.27**	+0.10	+0.16	+0.00	−0.02	**−0.34**
HAC t	(2.0)	(1.1)	(1.8)	(0.1)	(−0.4)	(−2.4)
Panel B: five-day $\beta_{h=5}$						
2010–2015	−0.53	−0.17	+0.06	−0.27	−0.27	−0.39
2016–2020	−0.07	−0.09	−0.17	−0.27	−0.19	−0.37
2021–2023	+0.06	−0.13	+0.11	−0.53	−0.48	−0.71
2024–2026	**+1.11**	**+0.75**	**+0.62**	**+0.32**	**+0.34**	+0.21
HAC t	(4.4)	(5.5)	(4.1)	(2.2)	(2.9)	(0.9)

Notes: Dollar haven beta against each currency on risk-off days (S&P 500 futures below −1%), by period. $\beta > 0$ means that the dollar *depreciates* against that currency when equities fall (haven bid lost); $\beta < 0$ means that it appreciates (haven). Panel A is the contemporaneous same-day response; Panel B is the cumulative five-day local projection. Reserve rival or haven = CHF, EUR, JPY; risk or other = GBP, CAD, AUD (see the text for the basis of the grouping in the literature on safe haven currencies). Newey–West (HAC) t-statistics (ten lags) are shown for the 2024–2026 row, where boldface marks coefficients significant at the 5% level by that criterion; under the episode block bootstrap of Appendix B, the 95% intervals for the five-day estimates of 2024–2026 include zero for every currency. The periods contain 156, 128, 114, and 67 risk-off days (the five-day projection in Panel B uses the 66 days of 2024–2026 with a complete window). The EUR column reproduces the headline (+0.10 at $h = 0$, +0.75 at $h = 5$).

Turning from currencies to bonds, the cross-section tells a more muted story (Table 5). Among the four sovereigns, only the German Bund moves to a statistically significant haven gain in 2024–2026 by the HAC criterion ($\beta_{h=5} = -0.23$, $t = -5.4$), although its episode boot-

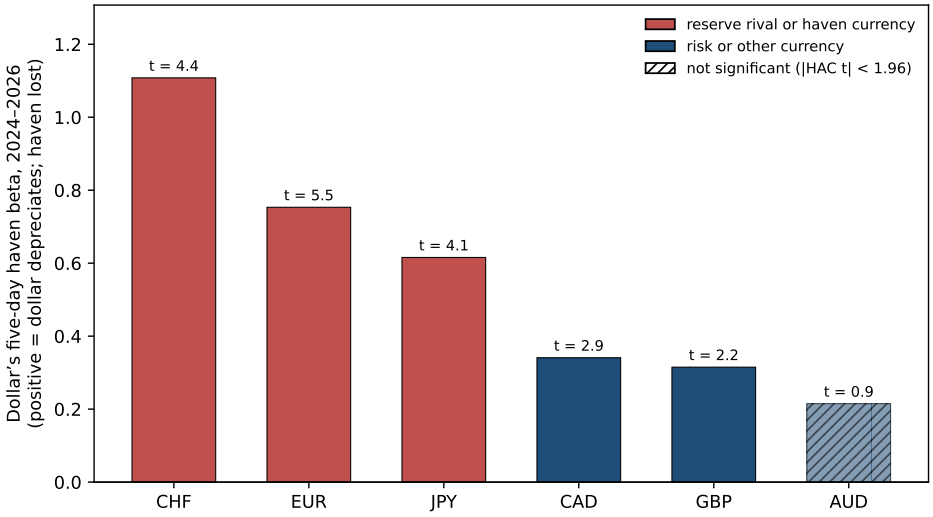


Figure 6: **The dollar's multi-day haven loss is ordered by reserve rival status.** Five-day haven beta ($\beta_{h=5}$) of the dollar against each currency in 2024–2026; positive values mean that the dollar depreciates (haven lost). The reserve rival and haven currencies (CHF, EUR, JPY) carry the largest losses; the pure risk currency (AUD, hatched) carries the smallest, which is insignificant. HAC t-statistics are shown above each bar. The episode block bootstrap intervals of Appendix B are wider and include zero for every currency, so the figure documents an ordering of point estimates.

strap interval, $[-0.49, +0.01]$, narrowly includes zero; the Eurex Bund futures corroborate the estimate (-0.22, $t = -2.4$, on the risk-off days since March 2025). The U.S. Treasury response is positive but insignificant, whether in duration-adjusted 10-year terms ($+0.26$, $t = 1.0$) or in 30-year futures ($+0.58$, $t = 1.4$), and the Japanese and Canadian responses are insignificant as well (-0.02 and $+0.16$). I read this conservatively: the bond evidence does *not* establish a loss of haven status for Treasuries that is specific to the United States. The multi-day comovement of stocks and bonds is also subject to the well-documented shift in the bond–equity correlation across inflation regimes (Campbell *et al.*, 2020), a confound that does not contaminate the currency cross-section. This is why the paper's central claim concerns the *dollar*, for which the evidence is considerably stronger, rather than Treasuries, for which it is weak.

The indeterminate Treasury response invites one more cut, across the maturity curve. Table 6 traces the five-day haven beta from 2-year to Ultra T-bond futures, and the pattern is monotone in duration. In 2024–2026 the 2-year contract is essentially unchanged from its mild pre-2024 haven response ($\beta_{h=5}$ from -0.02 to -0.01), and the response then rises steadily with maturity, to $+0.05$ at 5 years, $+0.18$ at 10 years, $+0.58$ for the long bond, and $+0.64$ for the Ultra T-bond, so the weakness is concentrated at the *long end*. Before 2024 the responses are small at every maturity but one. The exception in Table 6, a significant pre-2024 rally of the Ultra T-bond (-0.44, $t = -2.3$), reflects the construction of the continuous series rather than the market: the Ultra contract nearest to expiry is thinly traded late in each

Table 5: Five-day haven response of sovereign bonds to a U.S. risk-off shock

	U.S. (10-year)	Germany (10-year)	Japan (10-year)	Canada (10-year)	U.S. (30-year futures)
2010–2023	−0.06	+0.09	+0.02	−0.02	−0.07
HAC t	(−0.7)	(1.0)	(0.7)	(−0.3)	(−0.5)
2024–2026	+0.26	**−0.23**	−0.02	+0.16	+0.58
HAC t	(1.0)	(−5.4)	(−0.2)	(0.9)	(1.4)

Notes: Local projection of each sovereign's cumulative five-day bond return on the U.S. equity shock of day t, risk-off days, in comparable return units: the duration-adjusted 10-year price return (modified duration ≈ 8.5 applied to $-\Delta y$) for each country, with the United States also measured by 30-year Treasury futures. Positive values mean that the bond sells off when U.S. equities fall (haven lost); negative values mean that it rallies (haven gain); boldface marks coefficients significant at the 5% level by the HAC t-statistic. The earlier period pools 2010–2023 to mirror the formal break test. Only the German Bund shows a significant haven gain in 2024–2026 by the HAC criterion; its episode block bootstrap interval narrowly includes zero (Appendix B). The U.S., Japanese, and Canadian responses are statistically indeterminate. Yields are aligned to the futures calendar by carrying the value of the last business day over weekends and holidays. Ten-year yields are from FRED, the ECB AAA euro area yield curve, Japan's Ministry of Finance, and the Bank of Canada; Bund futures are from Eurex via Databento.

quarter and its rolls are not adjusted, and on a series that follows the most traded contract the estimate is −0.02 ($t = -0.1$; Appendix C). I therefore do not interpret it. The front end stays anchored, consistent with risk-off episodes still bringing expected policy easing, while the long end, where term premium and fiscal credibility are priced, is where the haven property weakens. No single maturity is individually significant in 2024–2026, so this locates the indeterminate Treasury weakness rather than establishing it; but its shape, monotone in duration and largest where fiscal credibility is priced, fits the credibility reading better than a uniform pattern or one led by the front end. The same maturity pattern emerges from an entirely different instrument in Acharya and Laarits (2026), who decompose Treasury yields around the tariff shock and find that the hedging property of the convenience yield eroded for long bonds while short-dated Treasuries retained it.

Table 6: The Treasury haven response across the maturity curve ($\beta_{h=5}$)

	2-year (ZT)	5-year (ZF)	10-year (ZN)	Long bond (ZB)	Ultra (UB)
2010–2023	−0.02 (−1.6)	−0.04 (−1.3)	−0.07 (−1.5)	−0.07 (−0.5)	**−0.44** (−2.3)
2024–2026	−0.01 (−0.1)	+0.05 (0.4)	+0.18 (0.9)	+0.58 (1.4)	+0.64 (1.2)

Notes: Local projection of the cumulative five-day log return of each Treasury futures contract on the U.S. equity shock of day t, risk-off days; $\beta > 0$ means that the contract sells off when equities fall (haven lost), and $\beta < 0$ that it rallies. Continuous CME futures (nearest expiry): 2-year (ZT), 5-year (ZF), 10-year (ZN), long bond (ZB), and Ultra T-bond (UB, the contract with the longest duration). HAC t-statistics (ten lags) in parentheses; boldface marks coefficients significant at the 5% level by that criterion. The two periods contain 398 and 67 risk-off days; the estimation samples are smaller where a contract lacks a quote at either end of the window (between 368 and 392 days before 2024 and between 59 and 66 in 2024–2026). Contracts with longer duration carry larger return magnitudes, so the informative comparison is the change in sign and the gradient across maturities, not the level. The pre-2024 Ultra T-bond estimate does not survive a series that follows the most traded contract and is adjusted at rolls (−0.02; Appendix C), whereas the 2024–2026 row does. The long bond column equals the 30-year futures figure in Table 5; the 10-year column here is the ZN futures return, whereas Table 5 reports the duration-adjusted measure based on cash yields, which is why the two 10-year numbers differ.

The two cuts point the same way. The dollar's estimated haven losses are largest against the rival reserve currencies and smaller against the risk currencies, the German Bund rallies, and gold shows no significant gain once it is measured in euros (its rally in dollar terms is the mechanical mirror image of dollar depreciation; Section 5.3). That pattern is what a

rotation *toward competing reserve assets* would produce (Farhi and Maggiori, 2018), rather than a generic collapse of diversification, and it is the reading I carry into the discussion of mechanisms below. In the language of the literature on the ranking of havens, the point estimates suggest that the dollar slipped down that ranking in 2025, the kind of shift that Fatum and Yamamoto (2016) document across the global financial crisis, when the yen displaced the franc as the safest haven.

5.5 What the options market priced

The evidence so far concerns realized returns: what the dollar *did* after risk-off shocks. A complementary question is forward-looking: did the options market *price* a rising risk of dollar depreciation? A realized reversal and an ex ante repricing of dollar tail risk are logically distinct, and finding both moves the reading from "the dollar happened to fall" toward "the market came to expect that it might."

I measure the latter with the 25-delta *risk reversal* on EUR/USD options: the difference in implied volatility between a 25-delta call and a 25-delta put on the euro futures contract, the standard option-implied gauge of directional crash risk in currency markets (Brunnermeier *et al.*, 2008). Because the contract is quoted in dollars per euro, a euro call is a bet on dollar *depreciation*, so a *positive* risk reversal means that the market pays more to insure against dollar weakness than against dollar strength. The classic "dollar smile" makes this skew *negative*, as investors pay up for dollar calls as a crash hedge, so a move toward zero or into positive territory is the options market analogue of the dollar shedding its haven premium. I build the skew from CME end-of-day option settlements, which the exchange marks for every listed strike. Because CME collects the option premium when the trade clears, I treat each settlement as a discounted price, invert it to an implied volatility with the model of Black (1976) for options on futures, take the forward from put–call parity so that the smile is internally consistent, and read the 25-delta wings off the daily smile at the listed expiry nearest to thirty days (Appendix A).

Table 7 and Figure 7 report the result, which tracks the realized reversal closely. In 2024 the EUR/USD risk reversal was negative on average (a yearly mean of −0.32 volatility points): the options market priced the dollar as the convenience yield literature would have it, as the asset whose value rises when risk does. In 2025 the skew *flipped sign* to +0.16, as the market paid more for protection against dollar downside than against dollar upside, and in the partial 2026 sample it reverted to −0.30. The quarterly path locates the flip: the skew turns positive in 2025Q2 (+0.62), the quarter of the tariff turmoil, stays positive for the rest of 2025, and is negative again by 2026Q2 (−0.52). The 2025 mean stands +0.48 volatility points above the pooled mean of 2024 and 2026 (HAC $t = 3.4$). The change involved the asymmetry of implied volatility as well as its level: at-the-money volatility was 6.4, 7.6, and 6.4 points in the three years, so 2025 saw both a rise in euro volatility and a repricing of the dollar's tail. In

2026 the haven pricing reappears where it should, on risk-off days, on which the EUR/USD skew averages −0.52, against −0.27 on all other days of that year.

Table 7 extends the skew to the other two reserve rivals. Against the Swiss franc, the risk reversal is positive throughout, as one would expect for a haven currency, but it *rose* markedly in 2025, from +0.38 to +0.92, before falling back to +0.46 in 2026: a step of +0.52 points relative to the surrounding years (HAC $t = 3.1$), almost exactly matching the step for the euro. Two reserve rivals, then, show the same repricing of dollar downside in 2025. The yen does not show the corresponding shift. Its skew is high in every year (+1.1 to +1.3, the highest of the three) and barely moves in 2025 (−0.05, insignificant). The yen's own haven premium may dominate the skew of its dollar cross, but the data do not establish why the yen differs. The evidence from the options market therefore rests on the euro and the franc, against which dollar tail risk was repriced in 2025.

Table 7: The 25-delta risk reversal against the three reserve rivals: the option-implied price of dollar tail risk

	2024	2025	2026	2025 vs. 2024 and 2026	ATM (2024/2025/2026)
EUR/USD	−0.32	+0.16	−0.30	**+0.48** (t=3.4)	6.4 / 7.6 / 6.4
CHF/USD	+0.38	+0.92	+0.46	**+0.52** (t=3.1)	7.2 / 8.0 / 7.5
JPY/USD	+1.34	+1.22	+1.14	−0.05 (t=−0.3)	10.3 / 10.2 / 8.6

Notes: Yearly mean of the daily 25-delta risk reversal, RR25 = IV(25Δ call) − IV(25Δ put), in implied volatility points, on options on the euro, franc, and yen futures. Each contract is quoted in dollars per unit of foreign currency, so RR25 > 0 means that the market prices dollar *downside* (upside in the foreign currency) more richly than dollar upside. Implied volatilities are inverted from CME end-of-day settlements with the model of Black (1976) for options on futures. CME options are margined equity-style, with the premium paid when the trade clears, so each settlement is treated as a discounted price with discount factor $D = e^{-rT}$ at the Secured Overnight Financing Rate r, and the forward is taken from put–call parity, $F = K + (C - P)/D$; the smile is read at the listed expiry nearest to thirty days. Setting $D = 1$ moves no yearly mean by more than 0.02 volatility points and no 2025 step by more than 0.01 (Appendix A). The fifth column is the difference between the 2025 mean and the pooled mean of 2024 and 2026 (HAC t, twenty lags; boldface marks significance at the 5% level); ATM is the at-the-money volatility by year. The franc and yen skews are positive throughout, consistent with those currencies being havens themselves; the signal specific to the dollar is the *change* in 2025, which is significant for the euro and the franc but not for the yen (see text).

The two lines of evidence coincide. Over the same window in which the dollar's realized multi-day haven beta turned positive, the option-implied price of insuring against dollar depreciation rose relative to the price of insuring against dollar strength, and when the realized effect faded, so did the skew. A repricing of this kind is consistent with the stress of 2025 having carried information about the dollar's safety as a reserve asset (a fall in s_1 in the model of Section 3), and its reversion in 2026 matches the realized evidence. The skew does not identify a cause any more than the returns do, but it provides complementary evidence of increased pricing of dollar downside risk.

5.6 Robustness

Table 8 collects a broad set of checks, which I group into breadth and insensitivity to the threshold; specificity to risk-off days; the rate channel and pooling; and the formal structural break with its placebos. Two remarks govern how the table should be read. First, because the 2024–2026 sample is small ($N = 66$ risk-off days in 35 episode blocks), I rely on inference that

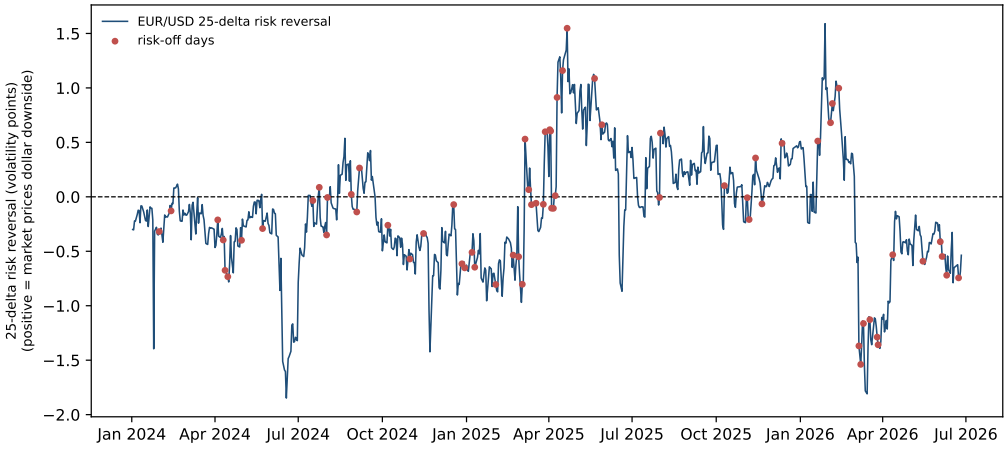


Figure 7: **The options market repriced the dollar's tail in 2025.** Daily EUR/USD 25-delta risk reversal (implied volatility points); negative values mean that the market pays up for dollar strength (the haven price), and positive values that it pays up for dollar weakness. Red markers are risk-off days. The skew is negative on average in 2024, turns positive in April 2025, stays positive in every month through February 2026, and returns to negative territory from March 2026, mirroring the realized multi-day reversal in an independent, forward-looking data source.

resamples whole episodes rather than on HAC standard errors alone. Second, the t-statistics in the table are HAC statistics, each measuring the precision of one specification. The stability of the point estimate is read from the estimates themselves, and no t-statistic in the table shows that a variant is significant at the episode level: as Appendix B documents, the episode bootstrap intervals for the 2024–2026 levels include zero, and the statistical case rests on the break test reported at the end of this section.

The reversal appears in the five-currency dollar basket ($\beta_{h=5} = +0.62$, HAC $t = 4.7$, on the days with all five quotes) and against each constituent (EUR +0.75, JPY +0.62, GBP +0.32, CHF +1.11, CAD +0.34; all HAC $t \geq 2.2$), each of which the dollar weakly hedged before 2024, so it is a dollar phenomenon and not an idiosyncrasy of the euro cross. Adding the Australian dollar, the risk currency deliberately excluded from the benchmark basket (Section 4), to form a six-currency basket attenuates the estimate only modestly (+0.55, $t = 3.8$), as the cross-section of Table 4 would predict, since the dollar's loss is smallest against the risk currencies. The estimate is similar across risk-off thresholds from −0.5% to −2% (from +0.59 to +0.81, all HAC $t \geq 3.8$). Under a definition of stress based on the VIX instead, it remains positive but is much smaller: +0.17 ($t = 1.9$) on days of VIX spikes (the top decile of daily changes in the VIX) and +0.23 ($t = 2.5$) on days with either a VIX spike or an S&P 500 decline of more than 1%. It also strengthens at the longer horizon ($h = 10$: +0.91, $t = 5.4$).

The sign flip is conditional on the severity of risk-off episodes, not an unconditional dollar trend. On calm days ($|r^{eq}| < 0.5\%$), the beta is insignificant (−0.19, $t = -1.5$); on risk-on days ($r^{eq} > +1\%$), it has the opposite sign (−0.28, $t = -5.0$), as the dollar keeps its ordinary inverse relation to equities. Estimated continuously over *all* days, the dollar's comovement with

equities moves from a hedging slope of −0.10 before 2024 to essentially zero in 2024–2026, a significant post-2024 interaction of +0.11 ($t = 2.2$), with the larger +0.75 concentrated in the risk-off tail.

The estimate survives controls for the cumulative change in the U.S. 10-year yield ($\beta = +0.80$, $t = 5.3$), for the expected policy path proxied by the cumulative change in the 2-year yield (+0.73, $t = 4.3$), and for the U.S.–German yield spread (+0.91, $t = 5.9$), so it does not merely reflect rate differentials or a repricing of the expected policy path, although, as stressed in Section 4, these controls are endogenous and condition on rather than rule out the rate channel. Because yield changes are an imperfect stand-in for monetary policy news, I also control for high-frequency FOMC surprises from the U.S. Monetary Policy Event-Study Database of Acosta *et al.* (2025), cumulating the surprises in the current policy rate and in the expected path from every FOMC communication window that falls inside each projection window: the estimate is essentially unchanged (+0.80, $t = 5.8$). Distinguishing pure policy surprises from central bank information surprises by the sign restriction of Jarociński and Karadi (2020) (a surprise that moves rates and equities in opposite directions is pure policy; one that moves them in the same direction is information) leaves the estimate at +0.75 ($t = 5.5$); no event classified as information falls inside the 2024–2026 risk-off windows, so the information control is identically zero there and the specification reduces to a control for the pure policy surprise alone. The bluntest version of the check is also the most transparent: only six of the 66 projection windows contain any FOMC communication, and dropping those six days gives +0.79 ($t = 5.6$). The estimate is thus little changed by controls for the measured FOMC surprises or by excluding the windows that contain FOMC communications. Nor is it an artifact of pooling across years: with year fixed effects the slope within years is still +0.60 ($t = 5.1$).[2]

Overlapping windows do not drive the inference either. Re-estimated on *non-overlapping* windows, in which consecutive risk-off days are more than five trading days apart so that no two windows share a daily return, the slope is +0.84 (OLS $t = 3.5$, $n = 44$). Dropping any one of the 35 episode blocks leaves the slope between +0.33 and +0.80, always positive. A formal break test rejects stability: the interaction of the equity shock with $\mathbf{1}[\text{year} \geq 2024]$, estimated on the full 2010–2026 sample, is +0.84 ($t = 5.5$), and its episode block bootstrap 95% interval, wide as a sample with few episodes warrants, excludes zero: $[+0.07, +1.08]$, with 98.5% of draws positive. This is the paper's formal statistical test. Placebo break cutoffs at 2018, 2020, and 2022 are all insignificant ($t \leq 1.8$), with significance only at 2024 ($t = 5.5$).

The break date can be probed more finely than annual placebos allow. Re-estimating the interaction on the full sample at every quarterly cutoff from 2021Q1 to 2026Q1 traces a distinct profile. The interaction is modest for every cutoff through 2022Q3 ($t \leq 1.7$ through

[2]The pooled +0.75 exceeds the slope in each single year (+0.50, +0.69, and −0.29) because the projection with a single intercept also loads on a positive covariation between years: years with deeper average risk-off shocks saw larger average dollar depreciation. The slope within years (+0.60), which nets out the year means, does not.

2022Q2 and a marginal $t = 2.1$ at 2022Q3), jumps at 2022Q4 (+0.79, $t = 5.0$), stays on a plateau of nearly constant magnitude for all cutoffs through 2025Q2 (peaking at +0.84, $t = 5.6$, at 2023Q2), and collapses to insignificance for cutoffs from 2025Q3 onward, that is, once the sample after the cutoff no longer contains the first half of 2025. Two conclusions follow. First, the data cannot distinguish a change that began in late 2022 or 2023 from one that began in 2024, because 2023 contributes only 22 complete risk-off windows and their stand-alone slope is an imprecise +0.13 ($t = 0.3$); the quiet stress calendar of 2023 simply carries little information. Second, and more important, every cutoff that produces a significant break does so by placing the 2025 episode after the cutoff, and no cutoff does so without it. The collapse of the profile after 2025Q2 is the clearest statement in the paper that 2025 carries the result. I therefore describe the change as estimated on the 2024–2026 window but *dated*, in substance, to the 2025 stress episode. Finally, the reversal is not the artifact of a single week. Dropping the entire April 2025 tariff turmoil roughly halves the estimate but does not eliminate it ($\beta_{h=5} = +0.42$, $t = 2.3$), and the estimate is positive and significant by the HAC criterion on risk-off days of U.S. origin: +0.50 ($t = 4.4$) on the 49 of the 66 days on which more than half of the whole day's equity decline, from the prior close to the close, came after the 09:30 open in New York. Two cruder rules, which compare the decline after the open with the decline in the 13:00–21:00 UTC window alone, give +0.51 ($t = 5.5$, 51 days, open fixed at 14:30 UTC) and +0.65 ($t = 5.2$, 58 days, 09:30 open in New York). The reversal is not, then, confined to shocks of foreign origin.

Table 8: Robustness of the dollar's multi-day reversal in 2024–2026 ($\beta_{h=5}$ unless noted)

Test	Estimate	HAC t
Baseline (dollar vs. EUR)	**+0.75**	5.5
episode block bootstrap 95% interval	[−0.01, +0.91] (Pr>0 = 97.4%)	
Five-currency dollar basket (EUR, JPY, GBP, CHF, CAD; $n = 62$)	**+0.62**	4.7
six-currency basket (adds AUD; $n = 62$)	**+0.55**	3.8
Threshold −0.5%	**+0.59**	4.0
Threshold −1.5%	**+0.70**	3.9
Threshold −2.0%	**+0.81**	3.8
Exclude the April 2025 tariff window (April 1–30)	**+0.42**	2.3
exclude the tariff fortnight (April 2–15)	**+0.41**	2.2
Risk-off days of U.S. origin only (whole day's decline, $n = 49$)	**+0.50**	4.4
window rule, open at 14:30 UTC ($n = 51$)	**+0.51**	5.5
window rule, open at 09:30 ET ($n = 58$)	**+0.65**	5.2
Horizon $h = 10$ (vs. $h = 5$ baseline)	**+0.91**	5.4
Non-overlapping five-day windows ($n = 44$, OLS t)	**+0.84**	3.5
Excluding Sunday evening sessions ($n = 62$)	**+0.69**	5.6
Range when one episode is left out (35 blocks)	[+0.33, +0.80], all > 0	
VIX spike days (top decile of ΔVIX; $n = 74$)	+0.17	1.9
VIX spike *or* S&P 500 below −1% ($n = 94$)	**+0.23**	2.5

The table continues with the checks of specificity and channel, the formal break, and the estimates by year; the notes follow its second part.

Table 8: Robustness of the dollar's multi-day reversal in 2024–2026 (continued)

Test	Estimate	HAC t
Specificity and channel:		
Placebo: calm days ($\lvert r^{eq}\rvert < 0.5\%$)	−0.19	−1.5
Asymmetry: risk-on days ($r^{eq} > +1\%$)	**−0.28**	−5.0
Control for cumulative U.S. 10-year yield change	**+0.80**	5.3
Control for cumulative U.S. 2-year yield change (policy path)	**+0.73**	4.3
Control for U.S.–German 10-year yield spread	**+0.91**	5.9
Control for cumulative FOMC surprises, USMPD (MP1 and MP2)	**+0.80**	5.8
pure policy vs. information split (JK sign restriction)	**+0.75**	5.5
exclude windows containing FOMC communication ($n = 60$)	**+0.79**	5.6
Slope within years (year fixed effects)	**+0.60**	5.1
Continuous, all days (no threshold): $r^{eq} \times \mathbf{1}[\text{year} \geq 2024]$	**+0.11**	2.2
Formal break: $r^{eq} \times \mathbf{1}[\text{year} \geq 2024]$ interaction	**+0.84**	5.5
episode block bootstrap 95% interval	[+0.07, +1.08]	(Pr>0 = 98.5%)
placebo cutoff at 2018 (full sample)	+0.25	1.7
placebo cutoff at 2020 (full sample)	+0.27	1.8
placebo cutoff at 2022 (full sample)	+0.44	1.7
Year by year (dollar $\beta_{h=5}$):		
2024 ($n = 21$)	+0.50	1.6
2025 ($n = 30$)	**+0.69**	5.8
2026, partial ($n = 15$)	−0.29	−1.1

Notes: Local projection of the cumulative five-day dollar return (against the euro unless noted) on the equity shock of day t, risk-off days, 2024–2026 ($N = 66$ in 35 episode blocks separated by more than five trading days) unless noted. Newey–West (HAC) t-statistics with ten lags (the baseline t ranges from 4.7 to 6.5 for bandwidths of five to fifteen lags); boldface marks estimates significant at the 5% level by that criterion. Each HAC t-statistic measures the precision of an estimate within its own specification; the stability of the estimate across specifications is read by comparing the estimates themselves. The episode block bootstrap (50,000 resamples, percentile intervals; Appendix B) is more conservative: its interval includes zero for the baseline level but excludes zero for the formal break. The baskets are equally weighted averages of the dollar's appreciation against each constituent on days with all constituent quotes. The VIX rows define risk-off days by a daily increase in the VIX in the top decile (FRED series VIXCLS; decile computed over 2010–2026) instead of the equity threshold; the VIX and the yields are aligned to the futures calendar by carrying the value of the last business day over weekends and holidays, so that a Monday's change is measured from the preceding Friday. The row that uses all days regresses the five-day return on the equity return, $\mathbf{1}[\text{year} \geq 2024]$, an indicator for negative equity returns, and the interaction of the equity return with $\mathbf{1}[\text{year} \geq 2024]$. The FOMC surprise rows cumulate the USMPD surprises in the current policy rate (MP1) and in the expected path (MP2) of Acosta *et al.* (2025) over each projection window, with the split into pure policy and information surprises based on the sign restriction of Jarociński and Karadi (2020). The non-overlapping windows select risk-off days in date order so that consecutive selected days are more than five trading days apart; the Sunday evening sessions are described in Appendix A. Days of U.S. origin are defined in Appendix A: the rule based on the whole day requires more than half of the day's equity decline to accrue after the 09:30 ET open; the two window rules compare the decline after the open with the decline between 13:00 and 21:00 UTC only (crude timing proxies, not identification claims). The formal break interaction is estimated on the full 2010–2026 risk-off sample ($N = 455$) and switches on for year ≥ 2024; the placebo rows re-estimate it with the cutoff moved to 2018, 2020, or 2022, and the text reports the full profile of quarterly cutoffs.

5.7 Mechanisms and the 2025 caveat

Two interpretations fit the evidence, and the data do not distinguish conclusively between them. Under the first, the intraday bid is a *liquidity* reflex. Dollars and Treasuries remain the deepest instruments to hold in the first minutes of stress, which is the convenience yield channel of Krishnamurthy and Vissing-Jorgensen (2012) and Jiang *et al.* (2021) and the dollar funding premium of Du *et al.* (2018), while the more deliberate multi-day reallocation increasingly favors substitutes whose issuers are not the source of the shock. That reallocation moves the exchange rate through the imperfectly elastic international financial markets of

Gabaix and Maggiori (2015). Under the second interpretation, the trade policy turmoil of 2025 made the United States itself the *origin* of risk, so the customary flight *to* U.S. assets was offset by a flight *from* U.S. policy. The two readings are not mutually exclusive: a liquidity reflex at the open and a reallocation driven by credibility over the following days can coexist, and in 2025 they appear to have pulled in opposite directions.

Controlling for the U.S. rate move, and for the U.S.–German yield spread, leaves the dollar selloff intact, so it does not merely reflect a relative repricing of yields. But because yields are themselves endogenous to the shock, this conditions on rather than rules out a monetary channel. The cross-asset rotation is consistent with the credibility reading but does not establish it. The Bund rally over the same horizon, significant by the HAC criterion although its episode bootstrap interval narrowly includes zero, suggests that safety was still sought beyond the first day, just not in the dollar, which is harder to reconcile with a story in which all havens fail together; a single foreign sovereign return, however, does not identify aggregate demand for safe assets. The gradient across reserve rivals in Section 5.4 points the same way: the point estimates show the dollar losing the most ground to competing reserve havens, less to risk currencies, and no significant ground to gold measured in euros, the pattern that a flight *from* U.S. policy *into* substitute reserve assets would produce.

The most direct cut I can offer, and the closest the data come to a test of Prediction P2, splits the 2024–2026 risk-off days by the *type* of shock that triggered them. If the multi-day reversal is a mechanical response to any equity selloff, it should appear whatever the trigger; if it reflects something specific to U.S. policy or credibility, it should be absent when the trigger is a routine data surprise. I classify a risk-off day as a *scheduled release* day if it falls on an FOMC announcement, an Employment Situation (payrolls) release, or a CPI release, in every case on the *actual* release date taken from the published FOMC calendar, the BLS release archive, and the FRED release date record. This matters in the sample because the government shutdowns of 2025–2026 displaced several releases from their customary slots at the start and in the middle of the month. All other days are *unscheduled*; a day that is both a scheduled release and a documented credibility event, for instance April 4, 2025, the day of China's tariff retaliation, which fell on a payrolls Friday, is assigned to credibility. One caution about definitions is essential: *unscheduled* is a residual category, not a synonym for "driven by U.S. policy." It contains the tariff and downgrade days but also foreign, geopolitical, and financial market news; the split therefore shows *when* the reversal appears, and the attribution to credibility is an interpretation that the rest of this subsection argues, not a fact that the classification delivers.

Table 9 reports the split. The reversal is concentrated on unscheduled days ($\beta_{h=5} = +0.82$, $t = 5.9$), where it is significant under both criteria: the episode block bootstrap interval, $[+0.09, +0.95]$, excludes zero. It remains after the curated cluster of tariff and downgrade days is removed ($+0.90$, $t = 2.6$). On the 14 scheduled release days, the point estimate is small ($\beta_{h=5} = +0.08$) and imprecise (bootstrap interval $[-0.70, +1.38]$), and the dollar's same-day

response stays on the haven side ($\beta_{h=0} = -0.13$). The difference between the two slopes, +0.73, is not statistically significant (HAC $t = 1.6$). Before 2024, the dollar did not sell off after *either* type of risk-off day, and scheduled release days were, if anything, the *stronger* haven ($\beta_{h=5} = -0.33$, $t = -2.1$, against −0.06 on unscheduled days). The larger estimate on unscheduled days is consistent with the credibility interpretation, in which the dollar still hedges ordinary macroeconomic news but sells off when the United States itself is the source of the risk. However, the scheduled group is small, the classification rests on a curated event list, and unscheduled days include events unrelated to U.S. policy.

Table 9: The reversal by shock type

	$\beta_{h=0}$	$\beta_{h=5}$	n
2024–2026			
All risk-off days	+0.10 (1.1)	+**0.75** (5.5)	66
Scheduled release (FOMC, payrolls, CPI)	−0.13 (−1.0)	+0.08 (0.4)	14
Unscheduled	+0.12 (1.3)	+**0.82** (5.9)	52
excluding the tariff and downgrade cluster	+0.03 (0.4)	+**0.90** (2.6)	47
2010–2023 (placebo: no selloff in either group)			
Scheduled release		−**0.33** (−2.1)	50
Unscheduled		−0.06 (−0.9)	339

Notes: Haven beta of the dollar against the euro on risk-off days, by shock type; $\beta > 0$ means that the dollar depreciates (haven lost), and $\beta < 0$ that it appreciates (haven). HAC t-statistics (ten lags) in parentheses; boldface marks coefficients significant at the 5% level by that criterion; n is the number of five-day windows. Scheduled release = a risk-off day with an FOMC announcement, an Employment Situation (payrolls) release, or a CPI release, all on *actual* release dates: FOMC from the Federal Reserve's published calendar, payrolls from the BLS news release archive (which records, among other deviations, the releases delayed by the shutdowns of 2013 and 2025–2026), and CPI from the FRED release date record (release 10). Unscheduled = all other days; a day that is both a scheduled release and a documented credibility event is assigned to credibility and so counts as unscheduled. The full date lists are in the replication package. Under the episode block bootstrap of Appendix B, the 2024–2026 interval at $h = 5$ is [+0.09, +0.95] for unscheduled days and [−0.70, +1.38] for the 14 scheduled release days; the difference between the two slopes is +0.73 (HAC $t = 1.6$), so the scheduled estimate is best read as showing no evidence of a reversal rather than as a precise level.

The reversal is concentrated in 2025. Year by year, the dollar's $\beta_{h=5}$ is +0.50 ($t = 1.6$, $n = 21$) in 2024, +0.69 ($t = 5.8$, $n = 30$) in 2025, and −0.29 ($t = -1.1$, $n = 15$) in early 2026; the quarterly profile of break dates in Section 5.6 makes the same point from the other direction, collapsing to insignificance for any cutoff that leaves the first half of 2025 out of the period after the cutoff. The estimation window is 2024–2026, but the effect is carried by 2025; the point estimate for early 2026 turns negative but is insignificant, so I cannot say whether it marks a genuine return to hedging or merely a pause. I therefore interpret the result as a pronounced *stress episode*, coinciding with the trade policy turmoil of 2025, in which the dollar temporarily ceded its multi-day haven function to foreign substitutes, and not as a permanent erosion of exorbitant privilege. This is a more conservative reading than that of Jiang *et al.* (2025), whose Treasury basis evidence from April suggested that investors perceived the repricing as potentially long-lasting; the reversion of both the realized beta and the option-implied skew in 2026 is the margin on which the two readings part, and more data will settle it. The paper's claim is the narrower one: the dollar's multi-day haven persistence *can* reverse, and did for a turbulent year, not that it has reversed for good.

6 Conclusion

The dollar did not stop hedging in the first minutes of stress; what changed in 2025 was its multi-day response. The intraday bid that the literature usually invokes was intact in recent stress, while the multi-day persistence reversed. The estimated contribution of the shock day's session window, in which the reflex operates, is small and insignificant; most of the multi-day depreciation accrued in the rest of the shock day and over the following days. Haven returns shifted toward German Bunds and rival reserve currencies, with the point estimates largest for competing reserve havens and smaller for risk currencies and gold, while the multi-day response of Treasuries was statistically indeterminate. The estimated pattern weakened in early 2026.

The statistical evidence has a clear structure. The change in the dollar's multi-day response between 2010–2023 and 2024–2026 is significant even when whole stress episodes are resampled, as is the reversal on unscheduled days; the Bund rally narrowly misses that standard. The level of the 2024–2026 response is large but rests on few independent episodes: its episode bootstrap interval includes zero, narrowly against the euro and more widely against the other currencies and the baskets, and the tariff episode of spring 2025 carries more than half of the point estimate. The finding refines, rather than confirms, the narrative of de-dollarization: the dollar's safe haven function was impaired at the multi-day horizon, not at the intraday frequency at which it is usually invoked. This conclusion is conditional on the paper's definition of stress based on equity returns, days on which S&P 500 futures fell by more than 1%. The reversal does not depend on the threshold, and under a definition of stress based on the VIX it remains positive but is much smaller. It also holds on futures series that follow the most traded contract and on Federal Reserve spot exchange rates (Appendix C). For policy, the relevant margin is the multi-day one. Should the reversal recur, the automatic dollar bid that the United States has long enjoyed in stress would become more conditional on the credibility of its own policy. The larger estimate on unscheduled stress days than on scheduled release days is consistent with such a credibility margin, but the difference is imprecisely estimated, and unscheduled days include more than policy news.

Three extensions would sharpen what a single episode can only suggest. The conditional comovement is not a causally identified shock to haven demand, so dating and classifying stress shocks by origin, and extending the cross-section of foreign sovereigns beyond the Bund, would tighten the link from U.S. policy news to the dollar's multi-day response. The cross-asset evidence rests on return responses rather than measured capital flows; the flow evidence to date finds diversification rather than exodus (Boissay and Huang, 2025), and a fuller accounting of who did the rotating (reserve managers, hedgers, or leveraged investors) would pin down the mechanism. And because the result is carried by a single year, evidence from subsequent stress episodes will help establish whether the 2025 pattern recurs.

References

Acharya, V. V. and Laarits, T. (2026) Tariff war shock and the convenience yield of US Treasuries: A hedging perspective, NBER Working Paper 34640.

Acosta, M., Ajello, A., Bauer, M. D., Loria, F. and Miranda-Agrippino, S. (2025) Financial market effects of FOMC communication: Evidence from a new event-study database, Federal Reserve Bank of San Francisco Working Paper 2025-30.

Black, F. (1976) The pricing of commodity contracts, *Journal of Financial Economics*, **3**, 167–179.

Boissay, F. and Huang, W. (2025) The safe haven properties of US Treasuries and portfolio flows of global investors, BIS Quarterly Review, September 2025.

Brunnermeier, M. K., Nagel, S. and Pedersen, L. H. (2008) Carry trades and currency crashes, in *NBER Macroeconomics Annual 2008, Volume 23* (Eds.) D. Acemoglu, K. Rogoff and M. Woodford, University of Chicago Press, pp. 313–347.

Caballero, R. J., Farhi, E. and Gourinchas, P.-O. (2017) The safe assets shortage conundrum, *Journal of Economic Perspectives*, **31**, 29–46.

Campbell, J. Y., Pflueger, C. and Viceira, L. M. (2020) Macroeconomic drivers of bond and equity risks, *Journal of Political Economy*, **128**, 3148–3185.

De Bock, R. and de Carvalho Filho, I. (2015) The behavior of currencies during risk-off episodes, *Journal of International Money and Finance*, **53**, 218–234.

Du, W., Keerati, R. and Schreger, J. (2026) Decoupling dollar and Treasury privilege, NBER Working Paper 35000.

Du, W., Tepper, A. and Verdelhan, A. (2018) Deviations from covered interest rate parity, *Journal of Finance*, **73**, 915–957.

European Central Bank (2025) What safe haven after the April US tariff announcement? Implications for euro area financial stability, Financial Stability Review, November 2025, Special Feature.

Farhi, E. and Maggiori, M. (2018) A model of the international monetary system, *Quarterly Journal of Economics*, **133**, 295–355.

Fatum, R. and Yamamoto, Y. (2016) Intra-safe haven currency behavior during the global financial crisis, *Journal of International Money and Finance*, **66**, 49–64.

Gabaix, X. and Maggiori, M. (2015) International liquidity and exchange rate dynamics, *Quarterly Journal of Economics*, **130**, 1369–1420.

Gopinath, G. and Stein, J. C. (2021) Banking, trade, and the making of a dominant currency, *Quarterly Journal of Economics*, **136**, 783–830.

Gourinchas, P.-O. and Rey, H. (2007) From world banker to world venture capitalist: US external adjustment and the exorbitant privilege, in *G7 Current Account Imbalances: Sustainability and Adjustment* (Ed.) R. H. Clarida, University of Chicago Press, pp. 11–66.

Grisse, C. and Nitschka, T. (2015) On financial risk and the safe haven characteristics of Swiss franc exchange rates, *Journal of Empirical Finance*, **32**, 153–164.

Habib, M. M. and Stracca, L. (2012) Getting beyond carry trade: What makes a safe haven currency?, *Journal of International Economics*, **87**, 50–64.

He, Z., Krishnamurthy, A. and Milbradt, K. (2019) A model of safe asset determination, *American Economic Review*, **109**, 1230–1262.

He, Z., Nagel, S. and Song, Z. (2022) Treasury inconvenience yields during the COVID-19 crisis, *Journal of Financial Economics*, **143**, 57–79.

Jarociński, M. and Karadi, P. (2020) Deconstructing monetary policy surprises: The role of information shocks, *American Economic Journal: Macroeconomics*, **12**, 1–43.

Jiang, Z., Krishnamurthy, A. and Lustig, H. (2021) Foreign safe asset demand and the dollar exchange rate, *Journal of Finance*, **76**, 1049–1089.

Jiang, Z., Krishnamurthy, A., Lustig, H., Richmond, R. J. and Xu, C. (2025) Dollar upheaval: This time is different, Working paper, SSRN 5220444.

Jordà, Ò. (2005) Estimation and inference of impulse responses by local projections, *American Economic Review*, **95**, 161–182.

Kaul, A. and Sapp, S. (2006) Y2K fears and safe haven trading of the U.S. dollar, *Journal of International Money and Finance*, **25**, 760–779.

Krishnamurthy, A. and Vissing-Jorgensen, A. (2012) The aggregate demand for Treasury debt, *Journal of Political Economy*, **120**, 233–267.

Lilley, A., Maggiori, M., Neiman, B. and Schreger, J. (2022) Exchange rate reconnect, *Review of Economics and Statistics*, **104**, 845–855.

Maggiori, M., Neiman, B. and Schreger, J. (2020) International currencies and capital allocation, *Journal of Political Economy*, **128**, 2019–2066.

Montiel Olea, J. L. and Plagborg-Møller, M. (2021) Local projection inference is simpler and more robust than you think, *Econometrica*, **89**, 1789–1823.

Ranaldo, A. and Söderlind, P. (2010) Safe haven currencies, *Review of Finance*, **14**, 385–407.

Sakurai, Y. (2021) How has the relationship between safe haven assets and the US stock market changed after the global financial crisis?, *Journal of International Financial Markets, Institutions and Money*, **75**, 101351.

A Data construction and variable definitions

This appendix collects the exact definitions used in the main text. The replication package, which contains the derived data behind every table and figure, the code that regenerates them from those data, the code that rebuilds the derived data from the licensed raw files, and a verification script that recomputes the table entries and the estimates quoted in the text and checks each against the printed value, is deposited at the Harvard Dataverse (https://doi.org/10.7910/DVN/6GRG61).

Trading calendar and risk-off days. Daily returns are computed from Databento's daily bars of the continuous futures, which aggregate trading by UTC calendar day. Because CME Globex reopens on Sunday evening, the Sunday session forms a separate daily observation, so the futures calendar has six trading days per week; horizons, window lengths, and the gaps that define episode blocks are all counted in trading days of this calendar. A risk-off day is a trading day on which front-month E-mini S&P 500 futures fall by more than 1% relative to the prior close ($N = 465$ from June 7, 2010, to June 26, 2026). Thirty-six of these are Sunday sessions. Twenty-three are risk-off days under both constructions compared in Appendix C and reflect gaps on weekend news (four of them in 2024–2026, among them April 6, 2025); the other 13, all between 2010 and 2021, fall two days after a quarterly expiry of the E-mini contract and reflect the construction of the continuous series (Appendix C). Excluding all Sunday sessions leaves the results intact: the 2024–2026 estimate is +0.69 (HAC $t = 5.6$, $n = 62$; Table 8), and the break interaction on the full sample is +0.88 ($t = 6.3$), with an episode block bootstrap interval of $[+0.20, +1.08]$. The intraday analysis uses the 429 weekday risk-off days (426 with both dollar and 30-year quotes), since a Sunday session contains no U.S. cash session. The multi-day projection uses the 455 risk-off days with a complete window from the prior close to $t+5$; of the ten without one, nine lack a euro quote for a Sunday session at one end of the window, and the window of the tenth, June 23, 2026, runs past the end of the sample.

Returns. The dollar return is appreciation against the euro, $-\Delta \ln(\text{EUR/USD})$; the basket return is the equally weighted average of $-\Delta \ln(\text{foreign/USD})$ over the euro, Swiss franc, yen, sterling, and Canadian dollar, computed on days with all five quotes; the six-currency robustness basket adds the Australian dollar. Unless a table or figure states otherwise, dollar results refer to the euro pair. In the rotation table, assets are expressed in comparable percentage return units: dollar appreciation; gold as a log return (in USD and in EUR); and bonds as duration-adjusted price returns, applying a modified duration of ≈ 8.5 to $-\Delta y$ for the U.S. and German 10-year bonds, with the United States also measured directly by 30-year Treasury futures returns.

Intraday window. Intraday betas use five-minute returns within 13:30–21:00 UTC, a fixed window that contains the U.S. cash session under both daylight saving and standard time (and, depending on daylight saving time, an extra hour before the open or after the close), winsorized at the 0.1st and 99.9th percentiles; standard errors are clustered by day. The New York robustness check in Table 1 instead restricts the window to 09:30–16:00 New York time, the exact cash session.

Clock decomposition. The session decomposition (Section 5.2, Table 2) splits the dollar's cumulative five-day appreciation against the Swiss franc into three non-overlapping segments: the session window of the shock day (the return within the fixed 13:30–21:00 UTC window, from the one-minute futures), the rest of that day outside the window (the daily return at h=0 less the session return), and all hours of the following five trading days (the return at h=5 less the return at h=0). Each segment is projected on the same equity shock of day t, so the three sum to the total at h=5. The franc is used because its continuous one-minute series is well covered (on the risk-off days of 2024–2026, a median of 464 of the 480 one-minute bars in the 13:00–21:00 UTC window from which each risk-off day is downloaded, against 92 for the euro, whose continuous series then often follows a thinly traded serial contract).

Monetary policy surprises. The FOMC surprise controls in Section 5.6 use the U.S. Monetary Policy Event-Study Database (Acosta *et al.*, 2025), taking the surprise in the current policy rate (MP1) and in the expected policy path (MP2) for each monetary event window and summing them over the FOMC communications that fall inside each risk-off projection window. The split into pure policy and information surprises classifies each event by the sign restriction of Jarociński and Karadi (2020), using the S&P 500 return in the same intraday event window: a surprise that moves rates and equities in opposite directions is pure policy, and one that moves them in the same direction is information.

Option risk reversal. The 25-delta risk reversal (Section 5.5) is built from CME end-of-day settlement prices of the European-style options on the euro, Swiss franc, and yen futures (product codes EUU, CHU, and JPU), which the exchange marks for every listed strike. CME options are margined equity-style: the buyer pays the premium in full when the trade clears, so a settlement price is a discounted value, $C = D \cdot \text{Black}(F, K, T, \sigma)$ with $D = e^{-rT}$, where T is the time to expiry and r is the Secured Overnight Financing Rate (FRED series SOFR) on that day. For each day and the listed expiry nearest to thirty days, I read the forward off put–call parity, $C - P = D(F - K)$, as the median of $K + (C - P)/D$ over the strikes within 3% of the futures price, invert each out-of-the-money settlement to an implied volatility with the model of Black (1976), compute the forward delta of each strike ($N(d_1)$ for calls and $N(d_1) - 1$ for puts), and interpolate the implied volatility at a call delta of 0.25 and a put delta of −0.25;

RR25 is their difference, in volatility points. Setting $D = 1$ changes the yearly means of Table 7 by less than 0.02 volatility points, the at-the-money levels by less than 0.05, and the 2025 steps by less than 0.01, and leaves their t-statistics unchanged to the first decimal. Each contract is quoted in dollars per unit of foreign currency, so a call is a bet on dollar depreciation and a positive RR25 is the option-implied price of dollar downside. The sample is daily, from January 2024 to June 2026.

Episode blocks. For the block bootstrap and the range of estimates when one episode is left out, risk-off days are grouped into episode blocks: runs of risk-off days in which each day falls within five trading days of the previous one, so that a new block starts after a gap of more than five trading days (35 blocks in 2024–2026). Because the dependent variable at $h = 5$ spans the returns of day t and of the five trading days that follow, days in different blocks share no daily return, while days whose windows overlap always fall in the same block. The block bootstrap resamples whole blocks (50,000 resamples, percentile intervals). Appendix B reports the distribution of block sizes, the diagnostics for influential episodes, and a comparison with blocks defined by calendar gaps of more than five days.

Classification by origin. Three rules classify a weekday risk-off day as of U.S. origin. The rule based on the whole day, reported first in Table 8, requires more than half of the day's equity decline, from the prior daily close to the daily close, to accrue after the cash open at 09:30 New York time; it combines the one-minute price at the open with the daily closes. The two window rules use only the one-minute bars in the 13:00–21:00 UTC window from which each risk-off day is downloaded and require more than half of the decline within that window to accrue after the open, fixed at 14:30 UTC or at 09:30 New York time; because the window begins shortly before the open, they classify most days as of U.S. origin. The rule based on the whole day and the window rule based on the New York open agree on 83% of the 429 weekday risk-off days, and the two window rules on 87%. Sunday sessions, which contain no U.S. cash session, are never classified as of U.S. origin. All three are crude timing proxies used for a robustness split, not identification claims.

Classification by shock type. For the split by shock type (Section 5.7), a risk-off day is a *scheduled release* day if it coincides with an FOMC announcement, an Employment Situation (payrolls) release, or a CPI release, and *unscheduled* otherwise. All three lists use actual release dates. FOMC announcement days come from the Federal Reserve's published meeting calendar. Payrolls days come from the BLS Employment Situation news release archive, whose dated filenames record every actual release from 2010 through 2026, including the departures from the usual first Friday of the month: releases on the second Friday, Thursday releases ahead of the Independence Day holiday, the release of October 22, 2013, delayed by that year's shutdown, and the reschedulings during the lapses in appropriations of 2025 and 2026,

under which no release occurred on October 3, November 7, or December 5, 2025, or on February 6, 2026, and delayed releases occurred on November 20 and December 16, 2025, and on February 11, 2026. CPI days come from the FRED release date record (release 10), which likewise carries the revised dates. A day that is both a scheduled release and a documented U.S. policy credibility event (sovereign downgrades, debt ceiling brinkmanship, and the 2025 tariff sequence) is assigned to credibility. The full date lists are in the replication package.

B Block bootstrap inference for the principal estimates

This appendix extends the episode block bootstrap used for the formal break test to the level estimates themselves. Table B.1 reports, for each principal estimate for 2024–2026, the percentile 95% confidence interval from resampling whole episode blocks (50,000 draws; blocks as defined in Appendix A) and the share of draws with the same sign as the point estimate. The last column repeats the exercise with blocks separated by calendar gaps of more than five days.

Table B.1: Episode block bootstrap 95% intervals for the principal estimates of 2024–2026

			Blocks by trading days (baseline)		Blocks by calendar days
Estimate	Point	HAC t	95% interval	Same sign	95% interval
Dollar vs. EUR (headline)	**+0.75**	5.5	[−0.01, +0.91]	97.4%	[−0.04, +0.91]
Dollar vs. EUR, same day ($\beta_{h=0}$)	+0.10	1.1	[−0.33, +0.17]	68.5%	[−0.33, +0.17]
Dollar vs. CHF	**+1.11**	4.4	[−0.38, +1.38]	85.4%	[−0.41, +1.36]
Dollar vs. JPY	**+0.62**	4.1	[−0.25, +0.94]	93.7%	[−0.26, +0.91]
Dollar vs. GBP	**+0.32**	2.2	[−0.22, +0.53]	91.4%	[−0.23, +0.54]
Dollar vs. CAD	**+0.34**	2.9	[−0.26, +0.46]	70.4%	[−0.27, +0.46]
Dollar vs. AUD	+0.21	0.9	[−0.39, +0.39]	75.9%	[−0.42, +0.40]
Five-currency basket	**+0.62**	4.7	[−0.19, +0.76]	91.5%	[−0.19, +0.76]
Six-currency basket (with AUD)	**+0.55**	3.8	[−0.21, +0.69]	89.9%	[−0.21, +0.69]
German Bund (duration-adjusted)	**−0.23**	−5.4	[−0.49, +0.01]	97.2%	[−0.52, +0.04]
Unscheduled days only	**+0.82**	5.9	[+0.09, +0.95]	98.4%	[−0.02, +0.96]
Scheduled release days only	+0.08	0.4	[−0.70, +1.38]	59.4%	[−0.70, +1.38]
Unscheduled minus scheduled	+0.73	1.6	[−0.91, +1.49]	84.6%	[−0.98, +1.49]
Memo: break interaction, full sample	**+0.84**	5.5	[+0.07, +1.08]	98.5%	[+0.04, +1.09]

Notes: All rows are five-day estimates ($\beta_{h=5}$) except the second, which is the same-day response. Percentile intervals from 50,000 resamples of whole episode blocks of 2024–2026 risk-off days. In the baseline, blocks are runs of risk-off days in which each day falls within five trading days of the previous one, so that blocks share no daily return; there are 35 blocks for most rows (between 34 and 36, depending on which days have complete data), 31 for the 52 unscheduled days, and 12 for the 14 scheduled release days. The last column instead starts a new block after a calendar gap of more than five days (39 blocks over the 66 days of the headline row). The difference row is the interaction of the equity shock with an indicator for unscheduled days in a regression on all 66 days. The memo row is the interaction $r^{\text{eq}} \times \mathbf{1}[\text{year} \geq 2024]$ of Table 8, whose bootstrap resamples the 201 blocks of the full 2010–2026 sample (234 with calendar gaps). HAC t has ten lags; boldface marks point estimates significant at the 5% level by that criterion. With 50,000 draws, the bounds of the headline and unscheduled intervals move by less than 0.02 across random seeds (replication package).

Four items accompany the intervals. First, *the grouping.* Because the dependent variable at $h = 5$ spans six daily returns, only a gap of more than five trading days guarantees that two risk-off days in different blocks share no daily return. The calendar grouping separates

four pairs of days that lie exactly five trading days apart and therefore share one daily return (April 4 and 10, 2024; August 28 and September 3, 2024; February 21 and 27, 2025; and April 10 and 16, 2025). Grouping by trading days places each of these pairs in a single block and reduces the number of blocks from 39 to 35. The two groupings give similar intervals; the one conclusion that depends on the choice is the reversal on unscheduled days, whose lower bound is +0.09 with blocks defined by trading days and −0.02 with blocks defined by calendar days. Second, *the distribution of block sizes*: of the 35 blocks, 22 are single days, seven have two days, three have three, one has four, one has seven (February 21 to March 18, 2025), and one, the tariff episode from March 26 to April 21, 2025, has ten; the small number of multi-day blocks is why the resampled intervals are wide. Third, *the influential episodes*: dropping the tariff block moves the headline beta from +0.75 to +0.33, by far the largest effect of any single block (no other block moves it by more than 0.05), which quantifies how much of the point estimate that episode carries; the full range of estimates when one block is left out, [+0.33, +0.80], remains strictly positive. Fourth, *the reading*. The level intervals for the dollar against each currency and against both baskets include zero; the headline interval does so narrowly, and those for the franc, the yen, and the baskets by a wider margin. The estimates whose intervals exclude zero are the reversal on unscheduled days and, most importantly, the break interaction estimated on the full sample, which is the paper's formal test; the interval for the Bund rally, [−0.49, +0.01], narrowly includes zero. The difference between the slopes on unscheduled and scheduled release days is not significant under either criterion. This is the same combination of strength and fragility that the estimates for single years convey, stated in resampling form: the change in the dollar's multi-day response between 2010–2023 and 2024–2026 is statistically robust, while the level of the response in 2024–2026 leans, for its precision, on the 2025 stress episode.

C Robustness to the construction of the futures series

The futures series used throughout are Databento's continuous contracts ranked by expiry, which follow the contract nearest to expiry and are not adjusted at rolls (Section 4). Two features of this construction could matter. First, the nearest expiry is not always the most traded contract: the euro, yen, sterling, Canadian dollar, and Australian dollar futures and gold list monthly serial contracts that trade thinly, and for every product the series holds the expiring contract until its last trading day, after most trading has moved to the next one. Second, each roll enters one daily return as the price gap between the old and the new contract. This appendix re-estimates the principal results on two alternatives, both produced by the replication package. The first holds, for each product, the contract with the largest trading volume, switching to a later expiry when that contract's volume on the previous day exceeds the held contract's, and computes every daily return within a single contract, so that no roll gap enters a return. The second uses the Federal Reserve's H.10 noon buying rates in

New York for the six currencies and measures the dollar's return from the business day before a weekday risk-off day to five business days after it, with the equity shock taken from the first alternative.

The first alternative changes the sample of risk-off days only slightly: it drops 15 of the 465 days and adds 12. Thirteen of the dropped days are Sunday sessions, all between 2010 and 2021, that fall two days after a quarterly expiry of the E-mini contract. The expiring contract stops trading at 09:30 New York time on the expiry Friday, so the baseline series books both the rest of that Friday's move and the price gap between contracts on the following Sunday. For the same reason, ten of the twelve added days are expiry Fridays on which the baseline series missed an afternoon selloff; one of them, March 20, 2026, falls in 2024–2026.

Table C.1 shows that the principal results do not depend on the construction. On the series rolled by volume, the headline slope is unchanged (+0.75, $t = 5.6$), the break interaction is larger and its bootstrap interval further from zero (+0.92, [+0.16, +1.11]), the ordering of the currency cross-section is the same, the reversal on unscheduled days is unchanged, and the pattern across years is the same. The spot rates, measured at noon in New York rather than at the close of the futures day, give somewhat smaller slopes but the same pattern: +0.63 ($t = 5.4$) against the euro, a larger loss against the franc (+0.94) than against the yen (+0.55) or the three risk currencies (+0.32 to +0.35), and a break of +0.82 ($t = 5.2$). The one estimate that does not survive is the pre-2024 rally of the Ultra T-bond, which falls from −0.44 to −0.02 and which Section 5.4 therefore does not interpret. Two statements shift at the margin. On the series rolled by volume, the dollar's pre-2024 slope against the euro is a significant hedge (−0.16, $t = -2.9$, against −0.09), which sharpens the contrast with 2024–2026, and the placebo break at 2022 reaches $t = 2.1$, consistent with the quarterly profile of Section 5.6, in which the break first becomes visible in late 2022.

The intraday estimates use one-minute bars of the same continuous contracts. For the euro, more than half of the five-minute returns are zero on 158 of the 429 risk-off days with intraday data, because the continuous series then follows a thinly traded contract. Restricting the reflex regressions to days with at least 400 of the 480 one-minute euro bars leaves the pattern of Table 1 intact and strengthens the recent estimates: −0.116 ($t = -4.6$) in 2010–2015, +0.015 ($t = 0.9$) in 2016–2020, −0.156 ($t = -8.4$) in 2021–2023, and −0.070 ($t = -3.4$, 18 days) in 2024–2026. Stale quotes therefore attenuate the reflex rather than create it.

Table C.1: Principal estimates under three constructions of the market data

	Baseline (nearest expiry)	Futures rolled by volume, adjusted at each roll	Federal Reserve H.10 spot rates
Risk-off days	465	462	433 (weekdays)
Dollar vs. EUR, 2024–2026	+0.75 (5.5)	+0.75 (5.6)	+0.63 (5.4)
bootstrap 95% interval	[−0.01, +0.91]	[+0.01, +0.92]	
Dollar vs. EUR, 2010–2023	−0.09 (−1.3)	−0.16 (−2.9)	−0.19 (−1.8)
Dollar vs. CHF, 2024–2026	+1.11 (4.4)	+1.16 (4.8)	+0.94 (4.4)
Dollar vs. JPY, 2024–2026	+0.62 (4.1)	+0.67 (4.6)	+0.55 (4.4)
Dollar vs. GBP, 2024–2026	+0.32 (2.2)	+0.35 (2.5)	+0.34 (2.6)
Dollar vs. CAD, 2024–2026	+0.34 (2.9)	+0.38 (3.4)	+0.35 (3.1)
Dollar vs. AUD, 2024–2026	+0.21 (0.9)	+0.28 (1.1)	+0.32 (1.2)
Five-currency basket, 2024–2026	+0.62 (4.7)	+0.66 (5.0)	
Break interaction, full sample	+0.84 (5.5)	+0.92 (6.2)	+0.82 (5.2)
bootstrap 95% interval	[+0.07, +1.08]	[+0.16, +1.11]	
Placebo break at 2022	+0.44 (1.7)	+0.52 (2.1)	
Unscheduled days, 2024–2026	+0.82 (5.9)	+0.82 (6.2)	
bootstrap 95% interval	[+0.09, +0.95]	[+0.14, +0.95]	
Scheduled release days, 2024–2026	+0.08 (0.4)	+0.11 (0.5)	
Dollar vs. EUR, 2024 only	+0.50 (1.6)	+0.50 (1.5)	+0.38 (1.4)
Dollar vs. EUR, 2025 only	+0.69 (5.8)	+0.69 (6.1)	+0.57 (5.9)
Dollar vs. EUR, 2026 only (partial)	−0.29 (−1.1)	−0.20 (−0.8)	−0.23 (−0.7)
German Bund, 2024–2026	−0.23 (−5.4)	−0.24 (−5.5)	
30-year Treasury futures, 2024–2026	+0.58 (1.4)	+0.57 (1.4)	
Ultra T-bond futures, 2010–2023	−0.44 (−2.3)	−0.02 (−0.1)	
Ultra T-bond futures, 2024–2026	+0.64 (1.2)	+0.69 (1.4)	
Gold (USD), 2024–2026	−0.92 (−2.0)	−0.99 (−2.4)	

Notes: Local projections of the cumulative five-day return (the dollar's appreciation unless the row names another asset) on the equity shock of day t, risk-off days; HAC t-statistics (ten lags) in parentheses; episode block bootstrap intervals as in Appendix B; the Bund is the duration-adjusted 10-year return of Table 5. Baseline: the series used in the main text. Futures rolled by volume and adjusted at each roll: each product holds the contract with the largest trading volume, and every return is computed within one contract; the risk-off days are redefined on this series. H.10 spot rates: Federal Reserve noon buying rates, on the weekday risk-off days of the equity series rolled by volume, from the business day before to five business days after (64 of the 433 days fall in 2024–2026). The Bund rows use the same yields in both futures columns and differ only through the set of risk-off days. Empty cells: not computed for that construction.