When Yield Curves Invert Together: Currency Crash Risk*

Alfredo Effendy†

University of Washington

August 31, 2026

Click here for the latest version

Abstract

A yield-curve inversion is a well-documented leading indicator of domestic slowdown, but a single inversion does not reveal whether the expected slowdown is country-specific or global. To rule out a signal driven by a single country's conditions, I construct the IYC signal from clusters of fresh inversions across G10 countries. The signal predicts elevated currency crash risk. During signal episodes, currencies with high interest-rate differentials (IRD) and high exposure to the equity market ("currency beta") depreciate the most, explaining the carry trade crash risk. Inversions, however, may reflect either a growth scare or inflation fighting. I use a strictly downward-sloping average G10 curve to identify inflation-fighting episodes, during which the IRD compensates subsequent depreciation on average, while in other signal episodes, depreciation is largely uncompensated.

JEL: E43, F31, G01, G11, G15.
Keywords: Inverted Yield Curve; Disaster Risk; Exchange Rate; Currency Carry Trade.

*I am grateful to my committee members—Yu-chin Chen (chair), Lukas Kremens, Eric Zivot, and Stephan Siegel—for their invaluable advice and guidance. I also thank Fabio Ghironi, Chang-Jin Kim, Yang Song, Philip Brock, Robin Yifan Luo, Hossein Poorvasei, Erik T.J. Andersen, Wenping Zhang, Tyson Ramirez, and other seminar participants at the University of Washington Department of Economics for helpful comments. All errors are my own.

†Email: `aeffendy@uw.edu`.

1 Introduction

The currency carry trade borrows in low-interest-rate currencies and invests in high-interest-rate ones. It has been profitable on average because high-interest-rate currencies do not depreciate enough to offset the interest rate differential (IRD), in violation of uncovered interest parity (UIP) (Hansen and Hodrick, 1980; Fama, 1984; Engel, 2016). The modern literature interprets this premium as compensation for depreciation risk and organizes currency returns into priced factors (Lustig and Verdelhan, 2007; Lustig et al., 2011; Menkhoff et al., 2012). The premium accumulates gradually, but much of it can disappear during currency crashes arising during unwinds of crowded positions and funding spirals, which leads to carry trade negative skewness (Brunnermeier et al., 2008), a pattern that recurs across a century of monetary regimes (Accominotti et al., 2019). In disaster models, these losses are the realizations for which the interest rate differential compensates investors in advance (Farhi and Gabaix, 2016; Chernov et al., 2018). The question facing a carry investor each month is whether the IRD adequately compensates for depreciation risk in the spot leg.

Answering that question requires an indicator of when currency downside risk is elevated. Existing indicators each have a drawback. Option prices require liquid currency-option markets, which restricts both the cross section and the time series (Burnside et al., 2011; Jurek, 2014; Farhi et al., 2015). Volatility and funding indices tend to register a crash as it occurs, while data on positioning and crowdedness are rarely public. An ideal indicator would be observable in real time, forward-looking, and available across countries and over several decades.

Government yield curves offer such an indicator because they are observed in real time, encode expectations across maturities, and have long histories across countries. I use their cross-country configuration to construct the *IYC signal*, which turns on when a fresh cluster of G10 yield-curve inversions forms. Freshness decides which inversions still carry information, and clustering separates a common risk from a country-specific one.

To summarize each country's curve, I focus on the three-month, two-year, and ten-year yields, which are widely used benchmarks for the short, intermediate, and long ends of the government term structure. Under normal conditions, yields generally rise with maturity, producing an upward-sloping curve. A yield curve is inverted when a shorter-maturity yield stands above a longer-maturity yield. For the IYC signal, I measure inversion using the two-year and ten-year yields as the baseline and report the three-month and ten-year version in Appendix B.1. An inversion is the classic recession signal (Harvey, 1988; Estrella and Hardouvelis, 1991; Estrella and Mishkin, 1998; Bauer and Mertens, 2018; Campbell and Shiller, 1991), and yield curves forecast economic activity both country by country and jointly (Estrella and Mishkin, 1997; Harvey, 1991; Diebold et al., 2008).

An individual inversion does not reveal whether the anticipated slowdown is country-specific or shared across economies. It may reflect a domestic credit cycle, an idiosyncratic inflation episode, or country-specific monetary policy. A diversified currency portfolio, however, is exposed to movements that extend beyond any one country. When inversions remain informative across several countries at the same time, their overlap is more likely to contain a common component. Looking across countries therefore helps distinguish an isolated domestic warning from a broader change in expectations.

The existing forecasting evidence also leaves open when an inversion is useful for trading. Most yield-curve based results are reported at horizons ranging from several months to several quarters. Such forecasts are useful for predicting economic activity, but they do not tell an investor whether to hedge or stay invested.

A simple count of inverted curves, which I call the *naive count*, does not supply such a rule. Curves can remain inverted for years, allowing old information that might have been priced in the market to overlap. The naive count mixes recent and stale inversions, producing clusters with potentially diluted information. To preserve the timing content of the signal, I apply a freshness filter before aggregating across countries.

A country's inversion enters the active set when its slope crosses below zero and remains

active until the slope has steepened for two consecutive months, whether or not it has returned above zero. Steepening indicates that the slope has begun moving back toward its normal upward ordering. I require two consecutive months to filter temporary reversals and avoid the whipsaw produced when the slope hovers around zero. The requirement follows the confirmation logic of the two-quarter recession conventions used in the United Kingdom, Canada, and Australia.

The signal then aggregates these active inversions across countries and turns on when at least two fresh G10 inversions are active at the same time. I use two because it is the smallest count that rules out a strictly single-country event. Months with exactly one active inversion remain favorable carry trade environments, with total returns averaging 10.09% per year. Figure 1 illustrates how the freshness filter separates the IYC signal from the naive count, and Section 2 gives the full construction.

The cross-sectional currency returns show what happens when a fresh cluster forms. Currencies with greater exposure to the global cycle, measured by their beta to the U.S. equity market and IRD, depreciate more. Persistently high-rate economies are often commodity exporters and small open economies that are exposed to the global cycle (Ready et al., 2017; Della Corte et al., 2016).[1] The carry trade inherits this exposure because it is long the high-rate, high-beta currencies and short the low-rate, low-beta currencies. In spot terms, a portfolio that is long high-beta currencies and short low-beta currencies loses 9.24% per year when the signal is active, against a gain of 3.57% when it is inactive, a gap of 12.82%. The carry trade loses 8.99% against a gain of 2.64%, a gap of 11.63%. Taking into account the IRD, the beta portfolio loses 8.29% against a gain of 5.29%, a gap of 13.58%, and the carry trade loses 3.67% against a gain of 6.19%, a gap of 9.86%.

During signal months, portfolio skewness becomes more negative and returns deteriorate throughout the distribution, with the largest changes in the left tail. Leading indicators

[1] Complementary accounts based on country size (Hassan, 2013), position in the trade network (Richmond, 2019), and the static ranking of Hassan and Mano (2019) load many of the same currencies on world risk.

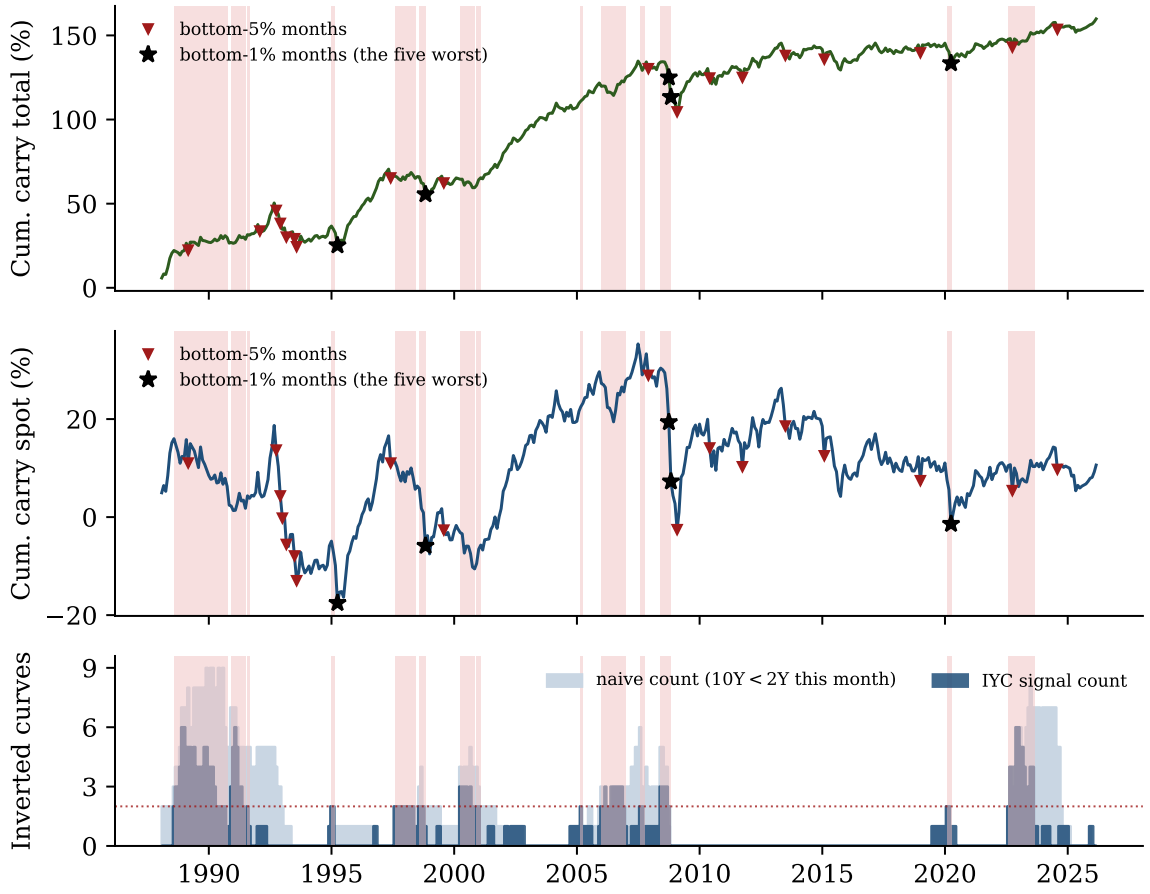


Figure 1: The carry trade and the IYC signal. Top to bottom, the carry portfolio's cumulative total return, its cumulative spot return, and two counts of inverted G10 curves. The *naive count* (light) counts curves with the 10-year yield below the 2-year this month, and the *IYC signal count* (dark) counts the inversions the signal counts, entered fresh and not yet confirmed re-steepened. The cluster threshold of two is dotted and IYC signal episodes are shaded throughout. Triangles mark each return series' own bottom-5% months, the crash months, of which the two sets share 22 of their 23. Black stars mark its bottom-1% months, the five worst of the sample, four of which fall inside the traded state.

also decline across G10 economies after cluster months, including in the United States. The currency and macroeconomic results both point to a shared weakening in expected activity rather than the isolated domestic slowdown that a single inversion may signal. In the language of the disaster-risk literature (Rietz, 1988; Barro, 2006; Wachter, 2013), I interpret the IYC signal as an observable indicator of elevated currency downside risk.

An inversion can arise in two distinct environments. Markets may expect rates to fall

because activity is weakening, or because a policy rate raised to fight inflation is expected to normalize. A growth-scare inversion tends to occur with the IRD near its ordinary level, leaving the subsequent depreciation poorly compensated. An inflation-fighting inversion occurs in a high-rate environment, in which the position collects unusually wide IRD. This leads to an important question. The signal identifies a period of elevated risk, but it does not say whether the IRD compensates enough of that risk to stay invested.

The shape of the full curve helps distinguish these cases. When the three-month yield stands above the two-year yield and the two-year yield stands above the ten-year yield, the short rate is the highest point on the curve. This shape is consistent with a policy rate held high while the market expects its eventual normalization. I call this strictly downward-sloping configuration the *downcurve* and measure it using the average G10 curve, consistent with the common risk construction.

During signal months in which the average G10 curve is in the downcurve configuration, the IRD is at its widest and, on average, covers most of the depreciation that follows. Outside the downcurve regime, the differential covers only a small fraction of the loss. The inversion cluster measures downside risk, while the IRD and the shape of the average curve indicate whether bearing that risk is adequately compensated.[2]

In terms of hedging, an investor who observes a cluster can buy options against the tail, unwind the carry position, or reverse into anti-carry (Caballero and Doyle, 2012). The carry premium can overwrite the spot signal in total returns. An investor who unwinds when the signal turns on forgoes the IRD. In the downcurve regime, that differential is large enough on average to cover most of the subsequent depreciation. I call the risk of surrendering this compensation *carry-premium risk* and the decision to remain invested the *carry overwrite*.

The cluster might sound like a repackaging of the global yield curve, the U.S. yield curve,

[2]This interpretation is consistent with the common component of term structures across currencies being largely a world-inflation object (Jotikasthira et al., 2015) and with curvature containing information about the policy stance (Mönch, 2012). Related work prices the persistent component of the differential (Kaebi and Martins, 2025), grounds it in habits (Verdelhan, 2010), and generalizes carry across asset classes (Koijen et al., 2018).

or other risk and volatility predictors such as the VIX, FX volatility, a commodity index, and the NFCI. I therefore run a horse race regression against all those measures, including each individual country's inversion, the United States included, and the results' power survives or even strengthens in some cases.

The closest yield-curve studies use relative Nelson–Siegel factors, curvature, and related shape measures to predict bilateral exchange rates and currency returns (Chen and Tsang, 2013; Dreher et al., 2020; Lloyd and Marin, 2024; Lustig et al., 2019).[3] These measures difference two countries' curves and remove the component they share, which is the component at issue here. A pairwise result against the dollar also does not map directly onto a diversified currency portfolio.

Conditioning carry performance on observable states is not new. Clarida et al. (2009) and Christiansen et al. (2011) allow the trade's returns and risk exposures to vary with volatility and liquidity regimes, while Barroso and Santa-Clara (2015) construct currency portfolios from observable characteristics. Related work conditions on average interest-rate levels (Lustig et al., 2014), macroeconomic inflation expectations (Andrews et al., 2024), and commodity, volatility, and liquidity predictors (Bakshi and Panayotov, 2013). The new element is the conditioning variable, a fresh cross-country cluster of government yield-curve inversions that identifies a common state before the associated currency losses are realized.

I am not aware of any prior currency work that uses overlapping fresh inversions across countries in this way. The construction is related to the global financial cycle, but it measures the common state from government yield curves rather than from realized asset prices and capital flows.[4]

The remainder of the paper proceeds as follows. Section 2 presents the portfolios, the signal construction, and the measure of risk compensation. Section 3 reports the predic-

[3]Curves and currencies have also been priced jointly in affine models (Backus et al., 2001; Ang and Chen, 2010).

[4]The global financial-cycle literature measures the common factor from prices and flows (Rey, 2013; Miranda-Agrippino and Rey, 2020), while intermediary-pricing models trace it to dealer balance sheets (He et al., 2017; Gabaix and Maggiori, 2015).

tive evidence, the cross-sectional exposure mechanism, the change in the conditional return distribution, and the hedging application. Section 4 compares the signal with rival curve measures and standard risk indicators. Section 5 extends the analysis to emerging markets, and Appendix B collects the robustness checks.

2 Framework and Signal Construction

2.1 Carry Crashes and Common Exposure

The portfolio construction follows the standard setting of Lustig et al. (2011), from the perspective of a U.S. investor. Let $s_{i,t}$ be the log spot price of currency i in dollars and $i_{i,t}$ its one-month money-market rate. Under UIP the expected appreciation would offset the IRD,

$$\mathbb{E}_t[\Delta s_{i,t+1}] \;=\; -\tfrac{1}{12}\left(i_{i,t} - i_{US,t}\right), \tag{1}$$

so a long position in currency i funded in dollars would earn nothing in expectation. The premium the carry literature documents is the unconditional violation of this parity, and Section 2.3 works with its conditional version. The monthly excess return on the position is

$$rx_{i,t+1} \;=\; \underbrace{\tfrac{1}{12}\left(i_{i,t} - i_{US,t}\right)}_{\text{income (carry) leg}} \;+\; \underbrace{\Delta s_{i,t+1}}_{\text{spot leg}}, \tag{2}$$

which under covered interest parity (CIP) equals the forward discount plus the realized appreciation, $s_{i,t+1} - f_{i,t}$. The income leg is known at t, and the spot leg is realized at $t+1$. Because Du et al. (2018) document post-2008 deviations between deposit and forward-implied rates, I construct the income leg from money-market differentials throughout.

Each month the currencies are ranked on $i_{i,t}$, and the *carry trade* goes long the k highest-rate currencies and short the k lowest-rate ones, equally weighted within each leg, so that

for a high-minus-low (HML) construction of size $k \times k$

$$r^c_{t+1} \;=\; \frac{1}{k} \sum_{i \in \text{ top-}k(i_{i,t})} rx_{i,t+1} \;-\; \frac{1}{k} \sum_{i \in \text{ bottom-}k(i_{i,t})} rx_{i,t+1} \;=\; C_t + X^c_{t+1}, \tag{3}$$

where applying the portfolio weights to the two components of (2) gives the income leg $C_t = \frac{1}{12k}\big(\sum_{\text{top-}k} i_{i,t} - \sum_{\text{bottom-}k} i_{i,t}\big)$ and the spot leg $X^c_{t+1} = \frac{1}{k}\big(\sum_{\text{top-}k} \Delta s_{i,t+1} - \sum_{\text{bottom-}k} \Delta s_{i,t+1}\big)$. I use $k = 3$ as the baseline, for the nine G10 currencies and for the sixteen-currency portfolio of Section 5.2 alike; the 1×1, 2×2, and 4×4 constructions are in Appendix B.1.

The *beta HML portfolio* replaces the sorting variable with the *currency beta*, the currency's exposure to the equity market in the sense of Lettau et al. (2014) and Kremens and Martin (2019). For each currency, $\hat{\beta}_{i,t}$ is the slope from the rolling regression of daily appreciation on the daily U.S. equity market return over the 90 trading days ending with month t, as in Kremens (2024),

$$\Delta s_{i,d} \;=\; \alpha_i + \beta_i\, r^M_d + \varepsilon_{i,d}, \qquad d \in \mathcal{D}_t, \tag{4}$$

where $\mathcal{D}_t$ is that window and r^M is the U.S. equity market return, and the portfolio applies the construction of (3) to the beta ranking,

$$r^\beta_{t+1} \;=\; \frac{1}{k} \sum_{i \in \text{ top-}k(\hat{\beta}_{i,t})} rx_{i,t+1} \;-\; \frac{1}{k} \sum_{i \in \text{ bottom-}k(\hat{\beta}_{i,t})} rx_{i,t+1}. \tag{5}$$

Its spot leg is X^β_{t+1}, defined as in (3). Appendix B.5 repeats the headline cells under a monthly 24-month rolling estimator and under an MSCI World benchmark.

In (3) the income leg C_t is known at t and strictly positive, so the trade loses in a month if and only if $X^c_{t+1} < -C_t$. Hence, every carry crash is an exchange-rate event in which the spot leg overwhelms the differential. In the sample the income leg averages +3.9%/yr, with a minimum of +0.7%/yr, while the spot leg carries all of the strategy's tail (Section 3.1).

The carry portfolio diversifies idiosyncratic risk by construction. Suppose the spot return of currency i is the sum of a common and an idiosyncratic component,

$$\Delta s_{i,t+1} \;=\; -\,\beta_i\, g_{t+1} \;+\; e_{i,t+1}, \tag{6}$$

where g_{t+1} is a realized common shock to the global cycle, such as a growth scare, and β_i is currency i's exposure to it. The currency beta of (4) measures this exposure empirically, using the equity market return as the proxy for g. Under the portfolio weights of (3), the spot leg is the exposure gap between the legs times the common shock, $-(\bar{\beta}^L - \bar{\beta}^S)\, g_{t+1}$, plus the difference of the legs' average idiosyncratic shocks. The second piece has a variance of order σ_e^2/k when each leg holds k currencies, so diversification and the monthly rebalancing keep it small. Consistent with this, widening the sort from 1×1 to the baseline 3×3 sharpens the carry portfolio's conditional loss from an annualized -10.0 (wild $p = 0.088$) to -11.6 ($p = 0.001$), and the single-currency legs carry the most idiosyncratic noise (Appendix B.1). Diversification bounds the size of the idiosyncratic piece but not its predictability. I therefore assume the piece is unforecastable, $\mathbb{E}[e_{i,t+1} \mid \mathcal{F}_t] = 0$, where $\mathcal{F}_t$ contains the time-t yield curves and the portfolio sorts, and Section 4.1 tests this restriction's empirical implication for the paper's own classifier. Only the common term carries a systematic loss, and the carry sort is exposed to it because the rate ranking tracks the exposure ranking (Section 3.2). A warning signal for the carry trade must accordingly be informative about the common shock and not about any single economy. Each curve prices its own economy's expected short-rate path, so a lone inversion can be country-specific while synchronization is evidence of an expected turn shared across economies. The data support the premise that the common component's leading content is real activity, measured by the OECD composite leading indicator (CLI), which is built to signal turning points six to nine months ahead.

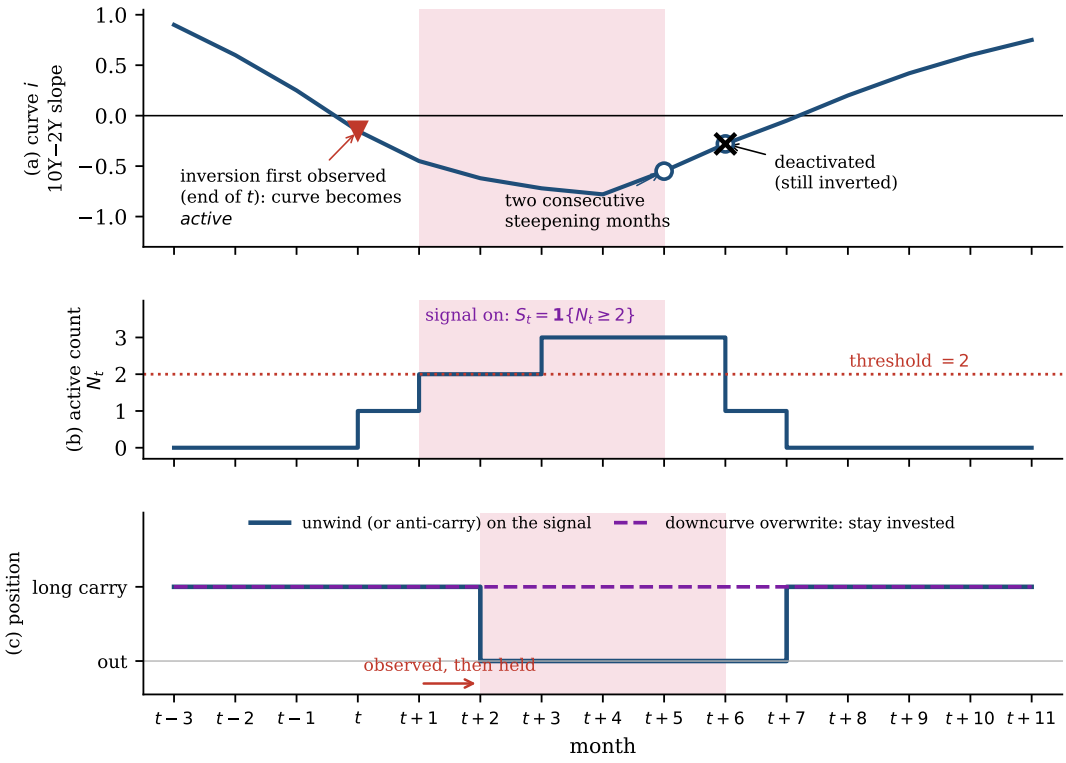


Figure 2: Signal timeline illustration. Panel (a) is one country's 10Y−2Y slope with the dated events of (8), where the inversion is first observed at the end of month t and the curve becomes active in that month (▼), its slope steepens in two consecutive months (circles), and it is deactivated (×) even though it is still inverted. Panel (b) is the active count N_t across the nine curves against the threshold of two, and the shaded band is the signal state IYC_t. Panel (c) is the position, unwinding (or going anti-carry) for the shaded months and returning when the signal turns off, while the downcurve overwrite stays invested throughout. Figure 3 shows a real episode.

2.2 Synchronized Yield Curves as a Signal of the Common State

Let $y_{i,t}^{(n)}$ be country i's n-period yield. Under the expectations hypothesis with a slow-moving term premium τ_i, the slope between a longer-dated tenor L and a shorter-dated tenor S is

$$y_{i,t}^{(L)} - y_{i,t}^{(S)} = \underbrace{\frac{1}{n_L}\sum_{h<n_L}\mathbb{E}_t[r_{i,t+h}] - \frac{1}{n_S}\sum_{h<n_S}\mathbb{E}_t[r_{i,t+h}]}_{\text{expected short-rate path}} + \tau_i + \varepsilon_{i,t}, \tag{7}$$

where $r_{i,t+h}$ is the one-period interest rate, n_L and n_S are the two horizons in months, τ_i the term premium differential between them, and $\varepsilon_{i,t}$ the transitory deviation from the hypothesis. The slope is negative when the market expects the short-rate path beyond the shorter tenor to average below its path up to that tenor, which happens when activity is expected to weaken (Harvey, 1988; Estrella and Mishkin, 1998) or when a policy rate already well above its normal level is expected to fall back toward it. The only asset-pricing input is that τ_i move slowly relative to the expected-path component; it need not be zero. The alternative explanation, an inversion produced by the premium rather than by the expected path, is addressed in Appendix B.9.

The signal is built curve by curve. Let $v_{i,t} = y_{i,t}^{(L)} - y_{i,t}^{(S)}$ be curve i's slope between the longer and shorter tenors of (7) at the end of month t, and $I_{i,t} = \mathbf{1}\{v_{i,t} < 0\}$ its inversion. The baseline takes L and S as the ten-year and two-year yields. Each curve carries an indicator $a_{i,t}$ of an *active* inversion and a counter $c_{i,t}$ of consecutive steepening months, $c_{i,t} = (c_{i,t-1} + 1)\,\mathbf{1}\{\Delta v_{i,t} > 0\}$, and

$$a_{i,t} = \begin{cases} 1 & \text{if } a_{i,t-1} = 0 \text{ and } I_{i,t} = 1,\ I_{i,t-1} = 0 \quad \text{(fresh inversion)}, \\ 0 & \text{if } a_{i,t-1} = 1 \text{ and } c_{i,t} \geq 2 \quad \text{(deactivated)}, \\ a_{i,t-1} & \text{otherwise}, \end{cases} \tag{8}$$

$$N_t = \sum_{i=1}^{9} a_{i,t}, \qquad \mathrm{IYC}_t = \mathbf{1}\{N_t \geq 2\}.$$

A curve becomes active in the month its inversion is first observed. An active inversion is deactivated only after two consecutive months in which that curve's slope steepens, whether or not the slope is back above zero, and a deactivated curve that stays inverted cannot re-enter until it has un-inverted and inverts afresh. The slopes are observed at the end of month t, the count and the state IYC_t are updated at that close, and the position is held over month $t+1$, so every return the paper regresses on the state is realized after the state is observed.

The nine curves belong to the nine portfolio currencies. The dollar is the numeraire, so its curve stays out of the count. Appendix B.1 repeats the headline cells with the U.S. curve in the count and the dollar in the portfolio, and none of the results change. The same recursion applied to one curve defines that country's *own signal*, which Section 4.1 uses. Appendix B.1 rebuilds the signal requiring one and three consecutive steepening months in place of two, and Appendix B.2 studies the naive count and an age-weighted count as an alternative measure of freshness, which Figure 3 illustrates on one episode. Figure 2 sets out the resulting sequence in event time.

If the signal carries information about a slowdown that is common, that information should be visible in the indicator that leads activity and not only in currency returns. Table 1 measures what the signal changes about an inversion, using the CLIs of all ten G10 economies, and Table B.10 in Appendix B.3 repeats it with each currency's spot return.[5] A curve that is active on its own predicts nothing (-0.01 pooled, $p \geq 0.90$, on 75 country-months)[6]. After a signal month the CLI falls in every G10 economy, whether or not its own curve is active. The pooled decline is -0.11 points, both on the own-country CLI and on the rest-of-G10 average (wild $p = 0.039$ and 0.036), and it reaches the United States, whose curve is not in the count (-0.13, wild $p = 0.022$). The local projections of Appendix B.10 extend the horizon, with the indicator about four-tenths of a point lower six months after an entry, which is consistent with the direct evidence that currency premia compensate business-cycle risk (Colacito et al., 2020).

[5]The indicator passes three checks. On point-in-time survey confidence, which is not a revised index, the cluster enters at -0.36 standard deviations of the monthly change (wild $p = 0.026$). Restricted to the six indicators carrying no 10Y$-$3M spread among their components, the pooled signal coefficient is -0.089 (wild $p = 0.097$). And the signal's return content is not spanned by the indicator, which leaves the signal at -12.43%/yr (wild $p = 0.001$) when added to the spot regression of Table 3.

[6]Interpreting the count as a summary of the common component requires that what remains in each curve be roughly independent across countries, and a common term premium is the obvious threat (Wright, 2011). When each curve is stripped of its fitted exposure to such a premium, most inversions remain standing and the beta HML loss holds, while the carry loss shrinks to a third of its size and loses significance (Appendix B.9).

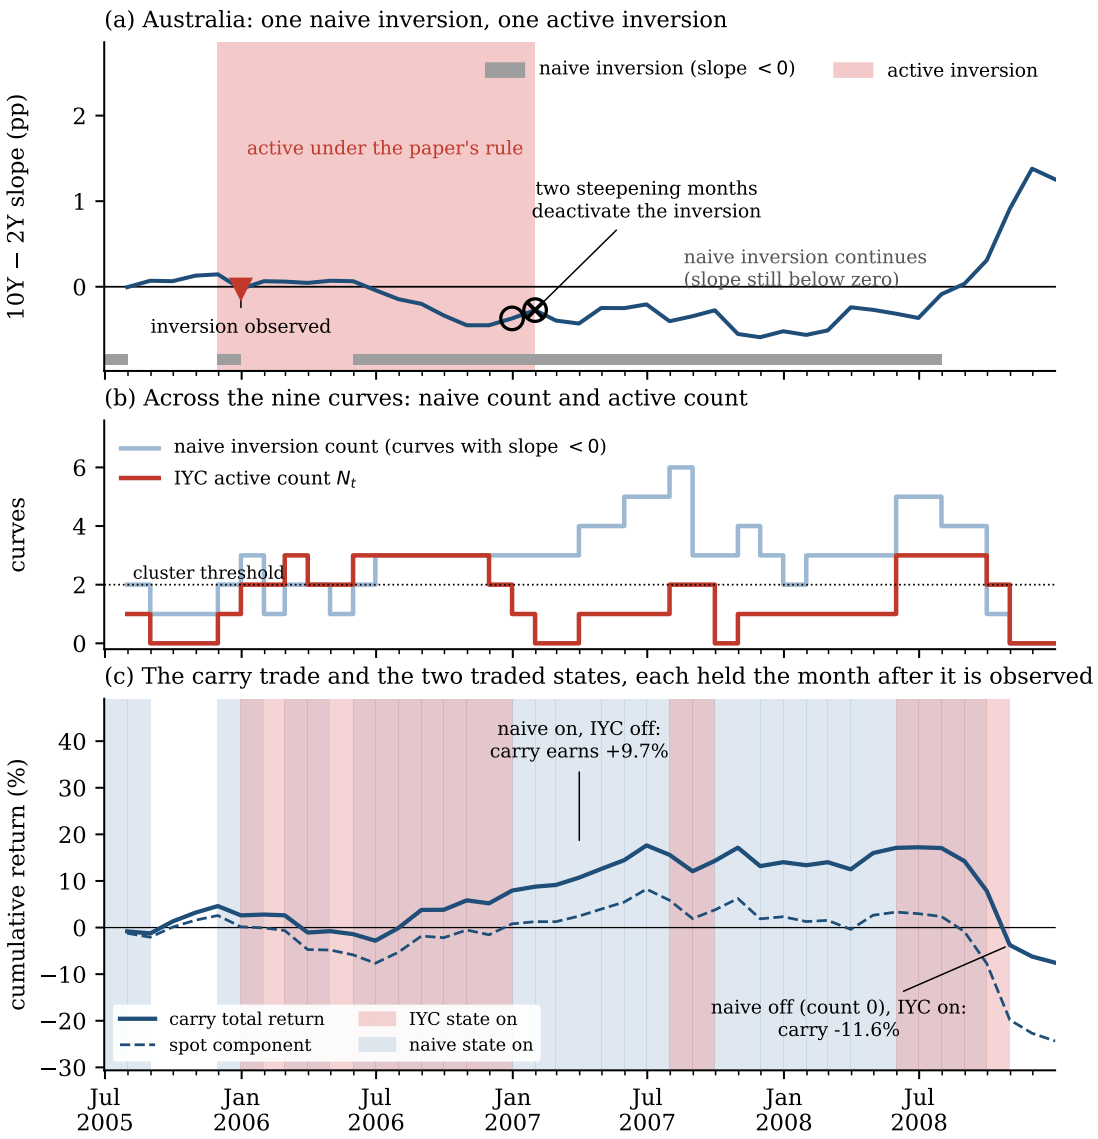


Figure 3: The naive count against the IYC signal, 2005:07–2008:12. (a) The Australian 10Y−2Y slope. The naive inversion is the grey bar (slope below zero); the active period under the paper's rule is shaded, from the month the inversion is first observed (▼) to the deactivation (×) after two consecutive steepening months (circles). The curve is deactivated in January 2007 while still inverted and stays naively inverted for another eighteen months. (b) Across the nine curves, the naive inversion count and the IYC active count N_t against the cluster threshold of two. (c) The carry trade's cumulative total return (solid) and its spot component (dashed), with the months in which each traded state is on shaded (each dated as it is traded). From January to June 2007 the naive state is on while the IYC state is off, and the carry trade earns +9.7%. In October 2008 the naive state is off, its count having fallen to one in September and to zero in October, while the IYC state is still on, and the carry trade loses 11.6%.

Table 1: Each economy's CLI after an own-signal month and after an IYC signal month. ON counts the months the country's own signal is on and No Cl. those of them with the IYC signal off. Each column is the state regression $A_{c,t+1} - A_{c,t} = \alpha + \gamma\,\mathrm{State}_t + e_{t+1}$, with the change in the OECD CLI in index points per month as the outcome. The three states are the country's own signal, that signal in the months the IYC signal is off, and the IYC signal. A dash marks a state with fewer than four months. The left column group takes the country's own indicator, the right the average indicator of the rest of the G10. Each coefficient reports its Newey–West (12 lag) p and the wild cluster bootstrap p ($B = 9{,}999$), and the pooled row adds country fixed effects with a Driscoll–Kraay (12 lag) p; 1988:01–2026:02.

			The Country's Own Indicator			Rest of the G10, Average Indicator		
	ON	No Cl.	Own	Own, No Signal	IYC signal	Own	Own, No Signal	IYC signal
AU	44	3	−0.045	—	−0.115	+0.026	—	−0.112
$NW_{12}[p]$			0.479	—	0.006	0.419	—	0.025
WCB[p]			0.642	—	0.096	0.552	—	0.032
CA	41	10	−0.185	−0.039	−0.189	−0.100	−0.033	−0.104
$NW_{12}[p]$			0.000	0.478	0.000	0.086	0.264	0.032
WCB[p]			0.011	0.561	0.019	0.031	0.162	0.047
CH	12	0	+0.063	—	−0.082	−0.010	—	−0.116
$NW_{12}[p]$			0.145	—	0.293	0.655	—	0.010
WCB[p]			0.751	—	0.264	0.739	—	0.033
EU	28	0	−0.025	—	−0.100	−0.077	—	−0.114
$NW_{12}[p]$			0.616	—	0.140	0.111	—	0.013
WCB[p]			0.618	—	0.133	0.686	—	0.033
GB	68	21	−0.026	−0.065	−0.055	−0.054	−0.012	−0.119
$NW_{12}[p]$			0.764	0.194	0.471	0.230	0.752	0.009
WCB[p]			0.778	0.211	0.484	0.348	0.766	0.018
JP	22	2	−0.092	—	−0.131	−0.088	—	−0.111
$NW_{12}[p]$			0.141	—	0.000	0.175	—	0.029
WCB[p]			0.244	—	0.004	0.698	—	0.060
NO	70	27	−0.127	+0.100	−0.142	−0.120	−0.072	−0.109
$NW_{12}[p]$			0.076	0.629	0.149	0.067	0.309	0.016
WCB[p]			0.368	0.819	0.191	0.027	0.403	0.049
NZ	54	12	−0.115	−0.073	−0.068	−0.027	+0.050	−0.118
$NW_{12}[p]$			0.238	0.262	0.346	0.401	0.355	0.017
WCB[p]			0.368	0.424	0.396	0.391	0.611	0.020
SE	24	0	−0.115	—	−0.105	−0.114	—	−0.113
$NW_{12}[p]$			0.234	—	0.052	0.184	—	0.020
WCB[p]			0.284	—	0.069	0.101	—	0.029
U.S.	42	9	−0.023	+0.016	−0.130	+0.007	−0.025	−0.111
$NW_{12}[p]$			0.547	0.731	0.004	0.834	0.420	0.023
WCB[p]			0.461	0.782	0.022	0.915	0.546	0.038
Pooled (nine)			−0.085	−0.009	−0.110	−0.064	−0.025	−0.113
$NW_{12}[p]$			0.004	0.903	0.000	0.001	0.426	0.000
WCB[p]			0.019	0.957	0.039	0.005	0.473	0.036
DK[p]			0.022	0.906	0.025	0.031	0.375	0.018

2.3 The Risk Compensation Ratio and Hedging

The natural application of the signal is to unwind the carry portfolio, or to reverse it into anti-carry, while the signal is on. An investor considering either would want to know the opportunity cost of hedging, which is the income the position would stop collecting, and how that cost compares with the crash risk the hedge avoids. The benchmark for that

comparison is the Fama (1984) regression of appreciation on the lagged differential, whose parity value is a slope of -1 and whose estimate sits near zero, the familiar puzzle. In most months the differential is collected with no systematic depreciation offsetting it, so a hedge is expensive on average, and the question is whether the signal marks the states in which the ordering reverses. Conditioning (3) on the signal makes the comparison explicit. The expected return is the income leg C_t plus the expected spot move, and if the spot risk in the state is a crash of size J arriving with probability λ_t, with no systematic move otherwise, then $\mathbb{E}_t[X^c_{t+1} \,|\, \text{signal}] = -\lambda_t J$ and

$$\mathbb{E}_t\left[r^c_{t+1} \,\middle|\, \text{signal}\right] \;=\; C_t - \lambda_t J \;\lessgtr\; 0 \quad\Longleftrightarrow\quad \kappa_t \;\equiv\; \frac{C_t}{\lambda_t J} \;\lessgtr\; 1, \tag{9}$$

so an unwind or an anti-carry position raises expected return if and only if $\kappa_t < 1$. In the same Fama regression, if the differential covers a fraction κ of the expected loss and crash exposure is aligned with the differential across currencies, the population slope inside the signal is $-1/\kappa$, so a slope below -1 is the regression counterpart of a conditional loss; Section 3.2 estimates the slope in and out of the signal. The curve informs both terms of the ratio. The threshold state of Section 2.2 is informative about λ_t, while the differential and, as Section 2.4 argues, the shape of the inverted curves are informative about κ_t.

Neither λ_t nor J is observable at t; the investor estimates them from the record of past signal months, and Sections 3.1 and 3.4 report the realized counterparts of both sides of the comparison. In Table 2, the IRD is higher in signal months, +5.32%/yr against +3.55 outside, so the compensation does rise with the risk, but the increase of +1.77 (wild $p = 0.17$) is a fraction of the spot loss of 11.63%/yr that arrives in the same state. An unwind forgoes $C_t + X^c_{t+1}$ for as long as the portfolio is flat, so its cost accumulates with the length of the unwind. An anti-carry position earns the crash if it comes but costs twice that if it does not, because the short pays the differential and loses the rally. Either way, the hedge forgoes the compensation for the risk it removes, and I call this cost the *carry-premium risk*, which is

largest when the differential is widest. An investor can therefore decide, from the observed C_t and an estimate of the crash risk it is being paid for, whether to unwind or to stay invested and *overwrite* the signal, staying in the trade because the compensation is high enough.

2.4 The Downcurve Regime

Inversion implies that the expected short-rate path falls; the shape of the whole curve distinguishes the environments in which it falls, and the environment determines κ_t. Tenor by tenor, the three-month yield reflects a high current policy setting, the two-year yield incorporates the expected near-term easing, and the ten-year yield places greater weight on the lower rate expected after normalization. Within the expectations hypothesis this is the profile of a rate raised to fight inflation and expected to normalize, the second case of Section 2.2. The paper calls this configuration *downcurve* and defines it on the average G10 curve, which is built tenor by tenor as the cross-sectional average of the nine countries' yields,

$$\bar{y}_t^{(n)} \;=\; \frac{1}{9}\sum_{i=1}^{9} y_{i,t}^{(n)}, \qquad n \in \{3M,\, 2Y,\, 10Y\}, \tag{10}$$

so that the regime is

$$D_t = \mathbf{1}\{\bar{y}_t^{3M} > \bar{y}_t^{2Y} > \bar{y}_t^{10Y}\}, \tag{11}$$

dated as the signal is. The regime is on in 52 of the 458 months, 37 of them inside the signal, and the split $\mathrm{IYC}_t \times D_t$ is the object Section 3.5 evaluates. When the front of the curve is not its highest point, the curve does not display the configuration associated with normalization from an elevated current policy rate. I treat these remaining signal months as the comparison regime without assigning them a unique structural cause. Figure 4 shows when the regime occurs. The downcurve months form two blocks, 1988–91 and 2023–24 (plus a single month in 2007), two well-documented inflation-fighting environments. The regime describes the environment in which an inversion arrives; the paper does not claim to identify what caused it.

Table 2: Stylized facts by state. Panel A reports months, the average G10 three-month, two-year, and ten-year yields, the average G10 money-market rate and the U.S. money-market rate and average CPI inflation (nine currencies, year on year, each row showing the reading from two months earlier, a conservative allowance for publication; all %), and the carry trade's IRD, spot return, and total return (annualized, %/yr), in the full sample, by downcurve regime, and by IYC signal state. Panel B reports the downcurve gap in the full sample, the gap between signal and no-signal months, and the downcurve gap within signal months, with the wild cluster bootstrap p on the WCB[p] row and its Newey–West p on the $\mathrm{NW}_{12}[p]$ row (regime runs, signal episodes, and signal runs as the respective clusters; $B = 9{,}999$). Sample 1988:01–2026:02.

		Rates and Inflation (%, Lagged)						Carry Trade (%/yr)		
State	Months	3M	2Y	10Y	G10 Rate	U.S. Rate	CPI YoY	IRD	Spot	Total
Panel A: Means by State										
Full sample	458	3.53	3.65	4.28	3.65	3.22	2.19	+3.91	+0.28	+4.19
non-downcurve months	406	2.94	3.13	3.90	3.06	2.71	1.86	+3.69	+0.41	+4.10
downcurve months	52	8.14	7.67	7.31	8.33	7.19	4.75	+5.62	−0.76	+4.86
No signal	365	2.84	3.00	3.80	2.96	2.50	1.84	+3.55	+2.64	+6.19
IYC signal	93	6.25	6.19	6.19	6.39	6.06	3.55	+5.32	−8.99	−3.67
non-downcurve months	56	4.18	4.39	4.67	4.25	4.92	2.57	+4.64	−14.95	−10.31
downcurve months	37	9.38	8.92	8.48	9.62	7.79	5.05	+6.35	+0.02	+6.37
Panel B: Gaps										
Downcurve − other, full sample		+5.21	+4.53	+3.41	+5.28	+4.48	+2.88	+1.93	−1.17	+0.76
$\mathrm{NW}_{12}[p]$		0.001	0.004	0.022	0.001	0.000	0.000	0.042	0.657	0.756
WCB[p]		0.327	0.425	0.672	0.339	0.188	0.033	0.524	0.846	0.716
Signal − no signal		+3.41	+3.19	+2.39	+3.43	+3.57	+1.71	+1.77	−11.63	−9.86
$\mathrm{NW}_{12}[p]$		0.002	0.002	0.010	0.003	0.000	0.002	0.004	0.003	0.013
WCB[p]		0.146	0.156	0.255	0.156	0.021	0.126	0.178	0.001	0.007
Downcurve − other, within signal		+5.20	+4.53	+3.81	+5.36	+2.86	+2.48	+1.71	+14.96	+16.68
$\mathrm{NW}_{12}[p]$		0.000	0.000	0.000	0.000	0.000	0.000	0.040	0.011	0.005
WCB[p]		0.007	0.040	0.150	0.007	0.011	0.002	0.527	0.063	0.053

Figure 5 shows how the average yield curve's shape changes across the states, and Table 2 documents the bond yields, the money-market rates, inflation, and the carry trade's compensation in each of them. The downcurve months are the high-rate, high-inflation months of the sample whether or not the signal is on. Three-month yields are near 9.4% in the downcurve signal months and near 5% in the downcurve no-signal months (mostly the 2023–24 block), against about 4% in the other signal months, and the carry trade's differential is at its largest there (+6.4 against +4.64%/yr within the signal). Panel B reports inference on the gaps. Within the signal, the downcurve gaps in rates and inflation are significant at the 5% level once signal runs are the unit of inference. Over the full sample only the

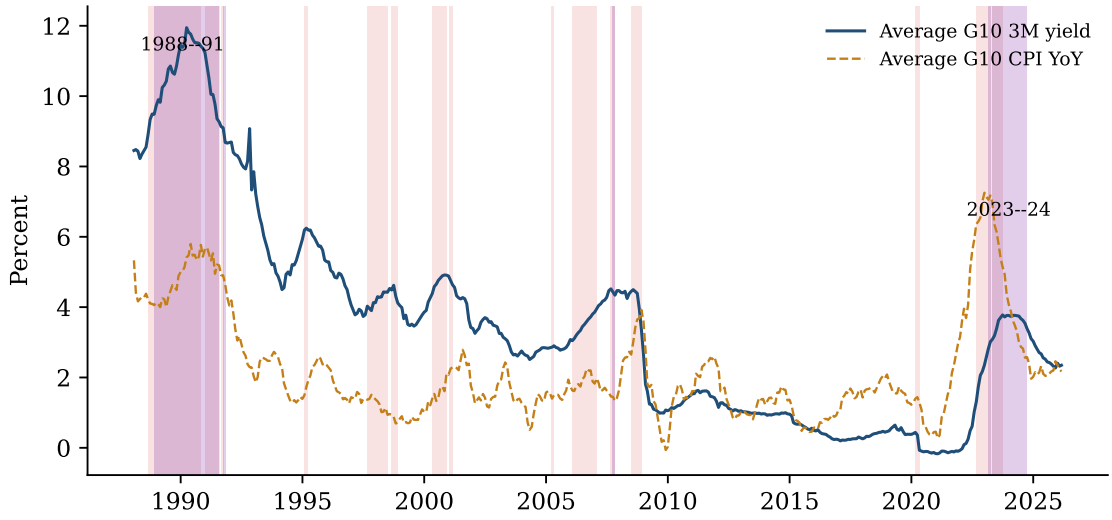


Figure 4: Downcurve, yields, and inflation. Average G10 three-month yield (solid) and CPI inflation (dashed); purple bands mark downcurve months (D_t of (11) on the average curve), pink bands the IYC signal episodes.

inflation gap is, because the 52 downcurve months form five runs, two of them blocks and three single months, and leave few clusters. The differential gap is positive in every split but not significant. Of the signal-versus-no-signal gaps in rates and inflation, only the U.S. rate's is significant (5%), with the inflation gap short of the 10% level ($p = 0.121$).

This compensation hypothesis links both terms of κ_t to the curve. The numerator is the differential, which Table 2 shows at its highest in the downcurve configuration. The denominator is lower there too, because the crash in (6) is a growth scare and an inflation-fighting inversion carries less of one, and because the currency is defended by the highest yields on offer, the inflation-fighting support channel of Clarida and Waldman (2008). Bearing the signal's risk should therefore be *paid* in the downcurve regime and *unpaid* in the rest, and Section 3.5 reports that this is the pattern in the data. The overwrite is exactly this, staying invested through those months instead of unwinding.

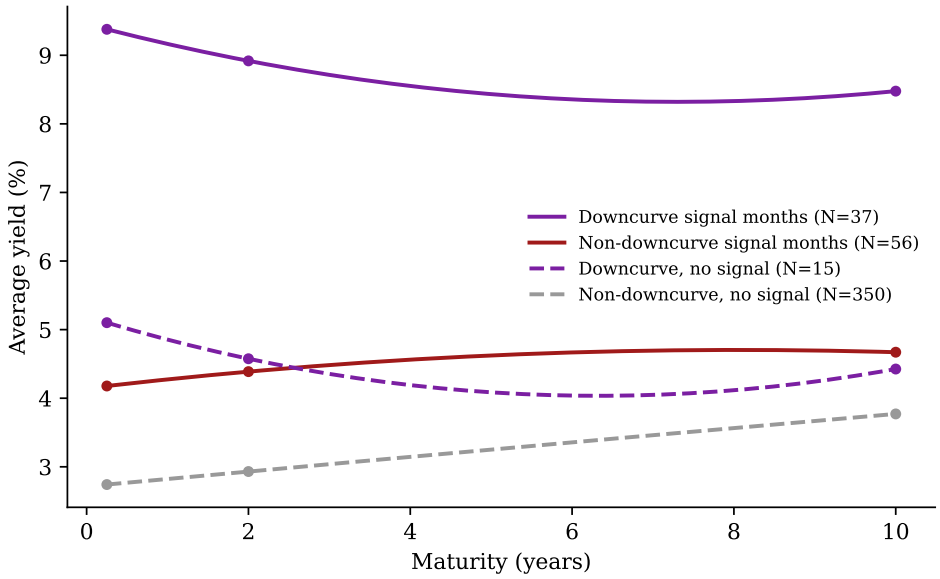


Figure 5: Average G10 yield curves by state and regime, the regime being D_t of (11), dated as the signal is. Solid lines are downcurve and non-downcurve signal months, dashed lines downcurve and non-downcurve no-signal months. Points are the average 3-month, 2-year, and 10-year yields.

2.5 Data

The main sample is monthly and runs from January 1988 to February 2026. The baseline universe is the nine G10 currencies (AUD, NZD, CAD, NOK, SEK, GBP, EUR, CHF, and JPY), and I use the DEM as a proxy for the EUR before 1999. For each currency I use its exchange rate against the dollar, its one-month money-market rate, and its 3M, 2Y, and 10Y government bond yields. The currency betas are estimated on daily (monthly) exchange rate and S&P 500 (MSCI World) returns, both of which are available 90 trading days (24 months) before the start of the sample.

For the EM extension I add the seven most liquid floating currencies, BRL, CLP, IDR, INR, MXN, KRW, and ZAR. Due to data availability, their sample is shorter. The last money-market rate to start is the Chilean, in January 1996, so the EM panel begins in February 1996, while the alternative signal built on the EM curves themselves begins in June 2003.

Among the other variables the NFCI, the CLI, and CPI inflation are available from

January 1988, while the VIX begins in October 1990. The CLI ends in November 2019 for New Zealand and in November 2022 for Switzerland, Norway, and Sweden, and I extend each with that country's business or consumer confidence index, which affects few of the signal months. All data are obtained from Bloomberg, Finaeon, central bank datasets, and the Chicago Fed. Appendix A sets out the economic regimes the sample spans.

3 IYC Signal Mechanism and Implementation

A synchronized inversion prices a slowdown the economies share, so the currencies most exposed to it should depreciate the most. The carry trade sorts on the IRD and the beta HML portfolio sorts on the currency beta, and both characteristics are higher for the currencies that carry that exposure, so in each portfolio the loss falls on the long leg. This section tests that mechanism through the two portfolios' returns (Section 3.1), the scaling of the loss with the IRD and with the currency beta across currencies (Section 3.2), the same conditioning in the Fama regression (Section 3.3), the shape of the conditional distribution (Section 3.4), and the state as a portfolio overlay (Section 3.5).

3.1 Portfolio Predictive Regressions

The paper's predictive regressions share one form. The outcome is a return realized in month $t+1$ and the regressor is the signal observed at the end of month t,

$$r_{t+1} \;=\; \alpha \;+\; \gamma\, \mathrm{IYC}_t \;+\; e_{t+1}, \tag{12}$$

where IYC_t is the signal of (8). I estimate (12) for the four portfolio returns the paper follows, the carry portfolio's spot leg X^{c}_{t+1} and total return r^{c}_{t+1} of (3) and the beta HML portfolio's spot leg X^{β}_{t+1} and total return r^{β}_{t+1} of (5), so that the exchange-rate leg and the income leg are separated. Section 4 adds controls to the same regression.

Table 3: The carry trade in and out of the IYC signal. For each series and state, the number of months, the annualized mean (%/yr), the annualized Sharpe ratio, the skewness, and the expected shortfall ($ES_{5\%}$, the mean of the worst 5% of monthly observations, %/month); states are the no-signal months, the IYC signal months, and the signal months outside and inside the downcurve regime. The γ (gap) column is the coefficient on the state dummy in (12), estimated over all months with the row's own state in place of IYC_t, with its Newey–West (12 lag) p and the wild cluster bootstrap p on the state's episodes beside it ($B = 9{,}999$). It is a difference from every other month and not a level, so it can carry the opposite sign to the mean in the same row. The *ex bottom decile* rows re-estimate the signal regression after dropping the series' full-sample bottom-decile months from both states, the ex-crash shift discussed in Section 3.4. Sample 1988:01–2026:02 ($N = 458$).

Series	State	N	Mean (%/yr)	Sharpe	Skew	$ES_{5\%}$	γ (Gap)	$NW_{12}[p]$	WCB[p]
Carry, spot	No signal	365	+2.64	0.35	−0.52	−5.06			
	IYC signal	93	−8.99	−1.01	−1.05	−7.59	−11.63	0.003	0.001
	non-downcurve	56	−14.95	−1.52	−1.10	−8.63	−17.35	0.002	0.002
	downcurve	37	+0.02	0.00	0.27	−3.71	−0.29	0.918	0.905
	ex bottom decile	80	−0.33	−0.06	0.76	−2.50	−8.38	0.001	0.010
Carry, total	No signal	365	+6.19	0.83	−0.45	−4.69			
	IYC signal	93	−3.67	−0.41	−1.04	−7.24	−9.86	0.013	0.007
	non-downcurve	56	−10.31	−1.05	−1.11	−8.27	−16.51	0.003	0.004
	downcurve	37	+6.37	0.96	0.41	−3.00	+2.37	0.351	0.396
	ex bottom decile	80	+5.09	0.85	0.81	−2.04	−6.45	0.009	0.022
Beta HML, spot	No signal	365	+3.57	0.43	0.04	−4.85			
	IYC signal	93	−9.24	−0.95	−0.98	−7.97	−12.82	0.001	0.006
	non-downcurve	56	−14.77	−1.62	−1.66	−8.60	−17.93	0.000	0.001
	downcurve	37	−0.88	−0.09	−0.58	−6.83	−2.01	0.740	0.858
	ex bottom decile	79	+1.04	0.16	0.83	−2.73	−7.69	0.011	0.035
Beta HML, total	No signal	365	+5.29	0.63	0.07	−4.67			
	IYC signal	93	−8.29	−0.85	−0.92	−7.86	−13.58	0.000	0.001
	non-downcurve	56	−13.29	−1.47	−1.60	−8.40	−18.03	0.000	0.001
	downcurve	37	−0.71	−0.07	−0.50	−7.05	−3.53	0.567	0.864
	ex bottom decile	79	+1.99	0.31	0.96	−2.53	−8.47	0.003	0.003

The *No signal* mean is the intercept of each regression, so mean = intercept + gap. A state can therefore show a positive mean and a negative gap. The *ex bottom decile* rows drop the bottom decile from both states, so their intercept is instead carry spot +8.05, carry total +11.54, beta HML spot +8.73, beta HML total +10.46.

Signal months are serially dependent, so each coefficient is reported with two p-values on labelled rows beneath it and no significance stars. The first is a Newey–West p at twelve lags (Newey and West, 1987), which treats the month as the unit; the second is a wild cluster bootstrap p with Rademacher weights, the null imposed, and entire signal runs resampled (Cameron et al., 2008; MacKinnon et al., 2023), which treats the episode as the unit (Appendix B.4 counts those episodes several ways). The episode standard is usually the more demanding, so a disagreement most often means that the estimate is significant

at the monthly frequency but not at the episode frequency, though a minority of cells run the other way. The γ (gap) column of Table 3 reports the four portfolio cells, the paper's headline, and Table 4 reports the per-currency version.

Outside the signal, the carry trade behaves as it does in the literature, earning +6.19%/yr in total returns at a Sharpe ratio of 0.83. Inside the signal, every moment of the return distribution deteriorates at once. The spot leg shows the repricing without the offset from the differential, and the beta HML portfolio shows the same shift. In the regression column all four portfolio gaps are negative, and each is significant on both the monthly and the episode standard.

The downcurve split shows where within the signal the loss is concentrated. In the signal months outside the downcurve regime, every cell sharpens. In contrast, the downcurve signal months show no loss in total returns and no significant loss in any cell. The differential is never negative in signal months, so unwinding always sacrifices income. Removing those months concentrates rather than dilutes the crash signal, because those months never carried the risk in the first place.

3.2 Interest Rate Differentials and Currency Betas

Crashes in the carry trade are concentrated in its long leg because a persistently high interest rate compensates a persistent exposure to global growth. Table 4 shows that the IRD and the currency beta are highly correlated across the sample ($\rho = 0.87$). The on-signal loss then scales with that persistent exposure, measured by a currency's sample-average beta or IRD.

In the cross-currency fit (Figure 6) the on-signal loss lines up with beta at $R^2 = 0.61$ with a small intercept (-1.6%/yr at zero beta), and it lines up with the average differential in the same way ($R^2 = 0.66$). Filtering out the downcurve months sharpens the picture, and the slope steepens to -80 per unit of beta at $R^2 = 0.81$. The commodity currencies lead the losers, with the New Zealand dollar at -18.73, the Australian dollar at -17.48 and the

Table 4: The mechanism currency by currency, sorted by beta. The first two columns are the sample-average money-market differential vs. the U.S. (%); sample-average real-time currency beta (the daily 90-day rolling estimator of (4)); and the currency's own on-state effect γ from $r_{i,t+1} = \alpha_i + \gamma_i \, \text{State}_t + e_{i,t+1}$ (annualized, %/yr). The left group takes the spot return $\Delta s_{i,t}$ and the right the total excess return $rx_{i,t}$ of (2), each for the IYC signal, for its non-downcurve months, and for its downcurve months (93, 56, and 37 months); no controls, 1988:01–2026:02. The Newey–West p and the wild cluster bootstrap p on the state's episodes are on the two labelled rows beneath each coefficient.

			Spot			Total		
Currency	Rate Diff.	Beta	IYC	Non-dc.	Downcurve	IYC	Non-dc.	Downcurve
AU	+1.87	+0.17	−11.53	−17.48	+0.14	−10.68	−18.72	+3.78
$NW_{12}[p]$			0.056	0.045	0.980	0.090	0.028	0.519
WCB[p]			0.046	0.094	0.973	0.142	0.073	0.517
NZ	+2.29	+0.15	−13.71	−18.73	−2.81	−12.93	−18.89	−0.87
$NW_{12}[p]$			0.006	0.017	0.535	0.012	0.013	0.858
WCB[p]			0.007	0.071	0.553	0.016	0.064	0.871
CA	+0.46	+0.14	−3.68	−8.55	+4.33	−3.20	−9.44	+6.68
$NW_{12}[p]$			0.260	0.057	0.108	0.358	0.031	0.017
WCB[p]			0.314	0.031	0.209	0.434	0.017	0.164
NO	+1.52	+0.10	−4.77	−6.15	−1.50	−5.36	−7.60	−0.70
$NW_{12}[p]$			0.395	0.409	0.837	0.331	0.285	0.924
WCB[p]			0.443	0.484	0.868	0.404	0.334	0.934
SE	+0.62	+0.09	−3.15	−4.06	−1.00	−3.36	−5.92	+1.25
$NW_{12}[p]$			0.545	0.567	0.876	0.513	0.384	0.849
WCB[p]			0.588	0.639	0.893	0.577	0.455	0.862
GB	+1.16	+0.04	−2.79	−2.37	−2.67	−1.48	−2.85	+0.90
$NW_{12}[p]$			0.554	0.689	0.714	0.753	0.621	0.905
WCB[p]			0.580	0.724	0.827	0.780	0.673	0.796
EU	−0.21	+0.02	+1.45	+2.76	−0.83	+0.36	+1.73	−1.72
$NW_{12}[p]$			0.775	0.656	0.918	0.942	0.771	0.828
WCB[p]			0.786	0.692	0.897	0.945	0.796	0.823
CH	−1.65	−0.06	+1.45	+3.18	−1.44	+0.43	+1.36	−1.03
$NW_{12}[p]$			0.748	0.486	0.867	0.924	0.758	0.906
WCB[p]			0.767	0.530	0.863	0.927	0.782	0.868
JP	−2.18	−0.11	−0.39	+5.35	−8.59	−2.12	+2.98	−8.94
$NW_{12}[p]$			0.920	0.359	0.132	0.604	0.623	0.171
WCB[p]			0.935	0.494	0.124	0.655	0.701	0.175

Canadian dollar at −8.55, while the safe havens gain, the yen at +5.35, the franc at +3.18 and the euro at +2.76 (Ranaldo and Söderlind, 2010). The carry trade is long the first group and short the second, which is where its loss comes from.

Consistent with its safe-haven role (Ranaldo and Söderlind, 2010), the yen does not lose in the growth scares in which the other currencies' losses go uncompensated. The downcurve

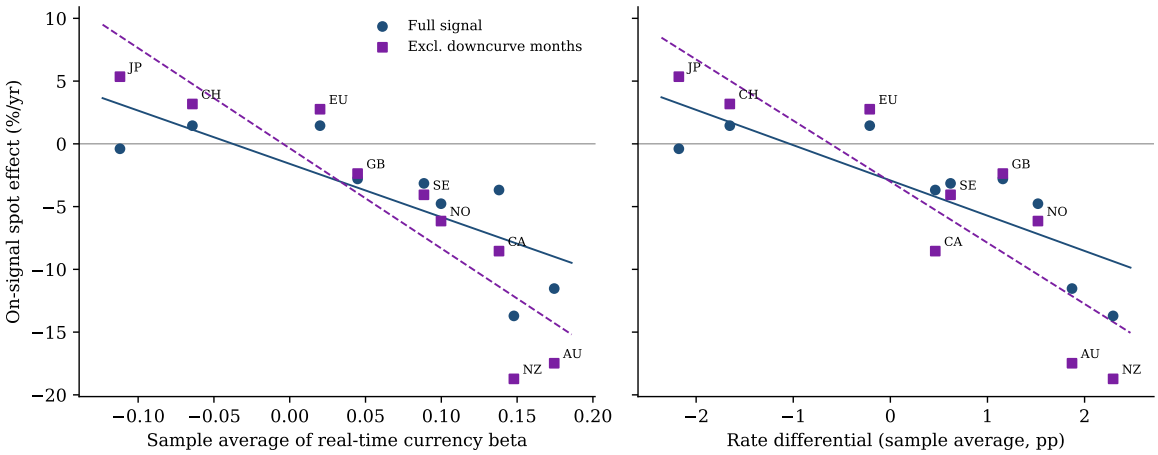


Figure 6: On-signal spot effect against the currency beta (left) and the average IRD (right), one point per currency, on a shared vertical scale. Circles and solid fits are for the full signal (slopes -42 and -2.8, $R^2 = 0.61$ and 0.66), squares and dashed fits for the non-downcurve months (slopes -80 and -4.9, $R^2 = 0.81$ and 0.75).

columns show the mirror image. No high-beta currency loses there, and the cross-currency fit on beta turns positive (slope $+23.4$, $R^2 = 0.45$). The yen, far from appreciating, carries the largest point-estimate *depreciation* of the nine, so the flight to safety that marks a growth scare does not merely disappear but runs the other way, although that estimate is not significant. Taking the differential into account does not change the qualitative result, and on the non-downcurve months the total-return losses are, if anything, slightly larger.

The cross-currency fits rest on nine points each, so I re-estimate both of them as panel regressions. Pooling the nine currencies, I estimate one regression for each measure of exposure,

$$\Delta s_{i,t+1} \;=\; \alpha_i \;+\; \gamma_1\,\mathrm{IYC}_t \;+\; \gamma_2\,\hat{\beta}_{i,t} \;+\; \gamma_3\left(\mathrm{IYC}_t \times \hat{\beta}_{i,t}\right) \;+\; e_{i,t+1}, \tag{13}$$

$$\Delta s_{i,t+1} \;=\; \alpha_i \;+\; \gamma_1\,\mathrm{IYC}_t \;+\; \gamma_2\,\mathrm{IRD}_{i,t} \;+\; \gamma_3\left(\mathrm{IYC}_t \times \mathrm{IRD}_{i,t}\right) \;+\; e_{i,t+1}, \tag{14}$$

with currency fixed effects α_i in both. I estimate them separately and not in one regression, because the two sortings rank the currencies almost identically and a regression containing both cannot say which of them the loss loads on. Each is estimated twice, once with the

exposure measure in real time at t and once with the currency's sample average in its place. Because the nine currencies share the dollar factor within a month, the panel replaces the Newey–West p with a Driscoll–Kraay (12 lag) p (Driscoll and Kraay, 1998), beside the same wild cluster bootstrap p on the signal's runs.

Table 5: Mechanism panel interactions. Each column estimates (13) or (14) on the nine G10 currencies, with the annualized spot appreciation (%/yr) as the outcome and currency fixed effects throughout. The exposure measure is named in the row label, $\hat{\beta}_{i,t}$ is the currency beta observable at t (the daily 90-trading-day estimator of (4)), $\bar{\beta}_i$ its sample average, $\mathrm{IRD}_{i,t}$ the money-market differential vs. the U.S. observable at t, and $\overline{\mathrm{IRD}}_i$ its sample average; the fixed effects absorb the level of either sample average. Beneath each coefficient are the Driscoll–Kraay p on the $\mathrm{DK}_{12}[p]$ row and the wild cluster bootstrap p on the WCB[p] row (B = 9,999); both are described in the text. Sample 1988:01–2026:02.

	(i)	(ii)	(iii)	(iv)	(v)
IYC_t	−4.13	−3.82	−1.59	−3.48	−2.91
$\mathrm{DK}_{12}[p]$	0.284	0.319	0.649	0.332	0.418
WCB[p]	0.352	0.413	0.685	0.440	0.478
$\hat{\beta}_{i,t}$		+9.04			
$\mathrm{DK}_{12}[p]$		0.144			
WCB[p]		0.168			
$\mathrm{IYC}_t \times \hat{\beta}_{i,t}$		+1.74			
$\mathrm{DK}_{12}[p]$		0.909			
WCB[p]		0.885			
$\mathrm{IYC}_t \times \bar{\beta}_i$			−42.42		
$\mathrm{DK}_{12}[p]$			0.035		
WCB[p]			0.024		
$\mathrm{IRD}_{i,t}$				+0.49	
$\mathrm{DK}_{12}[p]$				0.593	
WCB[p]				0.623	
$\mathrm{IYC}_t \times \mathrm{IRD}_{i,t}$				−1.76	
$\mathrm{DK}_{12}[p]$				0.190	
WCB[p]				0.330	
$\mathrm{IYC}_t \times \overline{\mathrm{IRD}}_i$					−2.80
$\mathrm{DK}_{12}[p]$					0.020
WCB[p]					0.012
Currency fixed effects	Yes	Yes	Yes	Yes	Yes
N	4122	4122	4122	4122	4122

In Table 5 the signal on its own is insignificant (column (i)), which is what a panel that averages the nine currencies should deliver, since it mixes the high-exposure currencies that lose with the funding currencies that do not. The two sample-average interactions

reproduce the cross-currency slope of Figure 6 and are significant on both standards. The monthly differential interaction keeps the sign of its sample-average counterpart but is not significant. The panel evidence is therefore consistent with the persistent component of currency exposure carrying the result, while its month-to-month variation does not.

3.3 Fama Regressions

The ranking in Table 4 suggests an implication for the forward premium puzzle, so I estimate the Fama (1984) regression of appreciation on the lagged differential in the same way. For each currency,

$$\Delta s_{i,t+1} \;=\; \alpha_i \;+\; \gamma_F \,\tfrac{1}{12}\,(i_{i,t} - i_{US,t}) \;+\; e_{i,t+1}, \tag{15}$$

where the parity benchmark is $\gamma_F = -1$. Table 6 estimates (15) separately on each of the four states. Outside the signal the pooled slope is positive, the forward premium puzzle in its familiar form. Inside it the implied slope is -1.27, past the benchmark and in the direction a conditional loss implies.

The interaction is confined to the non-downcurve months (-4.8, with per-currency slopes negative for eight of nine), and in the downcurve months it is essentially zero, which is where the differential covers the loss (Section 2.4). Lloyd and Marin (2024) document bilateral parity reversals on own-curve inversions at longer horizons, and the version here conditions on the global signal. The reversal inside the signal is a conditional counterpart to the disaster and habit accounts of the puzzle (Farhi and Gabaix, 2016; Chernov et al., 2018; Verdelhan, 2010), with the state observable in advance, although it is significant at the monthly frequency and not once whole episodes are the unit.

Table 6: Fama regressions by state. Specification: $\Delta s_{i,t+1} = \alpha_i + \gamma_F \frac{1}{12}(i_{i,t} - i_{US,t}) + e_{i,t+1}$, monthly appreciation vs. USD (%/month) on the lagged money-market differential (%/month), 1988:01–2026:02. Slopes are estimated separately on the no-signal months, the IYC signal months, and the signal months outside and inside the downcurve regime. Each slope reports its Newey–West (12 lag) p on the row beneath it and its wild cluster bootstrap p on the state's runs beneath that. The pooled row adds country fixed effects and interacts the differential with the signal (two-state) or with the non-downcurve and downcurve signal months (three-state); its entry is the implied slope in the state and the three p-values are those of the interaction (B = 9,999), the third a Driscoll–Kraay (12 lag) p robust to correlation across currencies within a month. UIP implies a slope of −1.

	No Signal	IYC Signal	Non-Downcurve	Downcurve
Currency	γ_F	γ_F	γ_F	γ_F
AU	+1.44	+0.26	−6.48	+0.33
$NW_{12}[p]$	0.112	0.834	0.223	0.677
WCB[p]	0.268	0.853	0.727	0.868
CA	−0.87	+1.14	−4.56	−0.47
$NW_{12}[p]$	0.343	0.152	0.494	0.623
WCB[p]	0.536	0.286	0.748	0.937
CH	+2.19	−2.06	−6.66	+0.40
$NW_{12}[p]$	0.035	0.399	0.020	0.916
WCB[p]	0.063	0.510	0.092	0.868
EU	+1.66	−5.89	−9.61	−2.02
$NW_{12}[p]$	0.206	0.047	0.001	0.630
WCB[p]	0.148	0.447	0.176	1.000
GB	+0.25	−0.79	−4.65	+0.56
$NW_{12}[p]$	0.881	0.641	0.252	0.818
WCB[p]	0.807	0.741	0.516	1.000
JP	+1.10	+1.59	+2.57	+3.55
$NW_{12}[p]$	0.225	0.120	0.050	0.000
WCB[p]	0.280	0.403	0.366	0.203
NO	−0.73	−3.43	−5.51	−2.35
$NW_{12}[p]$	0.544	0.185	0.182	0.412
WCB[p]	0.818	0.421	0.526	0.867
NZ	+0.88	−1.45	−6.50	+3.05
$NW_{12}[p]$	0.242	0.489	0.011	0.001
WCB[p]	0.441	0.701	0.222	0.438
SE	−0.73	−1.76	−6.22	−1.00
$NW_{12}[p]$	0.616	0.283	0.200	0.587
WCB[p]	0.903	0.472	0.520	0.743
Pooled, Currency Fixed Effects: Implied Slope by State				
Pooled	+0.49	−1.27	−4.37	+0.34
$NW_{12}[p]$		0.012	0.000	0.953
WCB[p]		0.319	0.019	0.976
$DK_{12}[p]$		0.190	0.032	0.968

3.4 Distributional Shift and Tail Risk

A further question is whether the signal fattens the tail or shifts the whole return distribution, and Table 7 answers it. I find that the signal produces both, and the two portfolios agree. For the *carry portfolio*, the signal shifts the entire distribution. Because the confirmed deactivation keeps the state on through four of the five worst months of the sample, the shift might rest on those months alone. It does not. Every reported quantile from the fifth through the seventy-fifth percentile moves left. The *ex bottom decile* rows of Table 3, which drop all 46 bottom-decile months—including the four the confirmed deactivation keeps—from both states, leave the carry portfolio's deficit negative and significant in spot and in total. For the *beta HML* portfolio the picture is the same. The probability of a crash month roughly doubles (4.1% to 9.7% at a -4%/month cutoff, wild $p = 0.022$; 8.8% to 15.1% at the bottom-decile definition, $p = 0.105$), the ex-crash shift is negative and significant in both legs, and the variance ratio (1.37) is not significantly above one (permutation $p = 0.21$). Table B.6 in Appendix B.2 runs the same accounting at the sample's own tail cuts, where a signal month is more than four times as likely as an unconditional one to fall in the worst one percent of carry-total months. Filtering on the downcurve divides the carry portfolio's distribution cleanly, with the non-downcurve months carrying the left tail and the downcurve months' lower tail thinner than the no-signal months'. The beta HML portfolio does not share the split, since its downcurve months hold the more negative fifth and tenth percentiles and the higher volatility.

The signal raises the odds of a crash without catching every one. Taking a crash to be a monthly loss beyond 4%, the carry spot leg crashes in 8.60% of signal months against 5.68% of months overall, and the beta HML leg roughly doubles its rate. The state is on in only a fifth of the sample, so it cannot contain the whole tail, and the inference the paper leans on is the shift in the distribution and not the conditional tail cells. Table B.6 also sets the coverage against the naive count, which is on for a third of the sample and still holds fewer

Table 7: Return distribution conditional on the signal. Quantiles of monthly returns (%/month; spot and total for both portfolios), annualized volatility, and skewness, out of the signal, in its 93 months, and in the signal months outside and inside the downcurve regime; 1988:01–2026:02.

Series	State	N	q5	q10	q25	q50	q75	Vol (Ann.)	Skew
Carry, spot	Off	365	−3.97	−2.41	−0.97	0.44	1.55	7.44	−0.52
	On	93	−4.96	−3.79	−1.75	−0.75	0.74	8.94	−1.05
	On, non-downcurve	56	−6.80	−4.11	−2.39	−1.03	0.02	9.85	−1.10
	On, downcurve	37	−2.41	−1.71	−1.26	−0.17	1.37	6.68	0.27
Carry, total	Off	365	−3.69	−2.18	−0.65	0.69	1.89	7.48	−0.45
	On	93	−4.44	−3.38	−1.41	−0.16	1.19	8.98	−1.04
	On, non-downcurve	56	−6.36	−3.71	−2.00	−0.58	0.33	9.82	−1.11
	On, downcurve	37	−1.85	−1.35	−0.93	0.27	1.77	6.67	0.41
Beta HML, spot	Off	365	−3.65	−2.69	−1.03	0.30	1.78	8.33	0.04
	On	93	−5.87	−3.94	−1.84	−0.49	0.79	9.75	−0.98
	On, non-downcurve	56	−5.20	−3.93	−2.32	−0.89	0.15	9.10	−1.66
	On, downcurve	37	−6.11	−4.13	−1.20	−0.06	1.62	10.32	−0.58
Beta HML, total	Off	365	−3.47	−2.58	−0.97	0.46	1.92	8.38	0.07
	On	93	−5.49	−3.83	−1.68	−0.29	0.72	9.72	−0.92
	On, non-downcurve	56	−5.05	−3.80	−2.28	−0.91	0.30	9.06	−1.60
	On, downcurve	37	−5.67	−4.27	−1.21	0.02	1.67	10.38	−0.50

of its worst months.[7]

3.5 Portfolio Trading and Risk Management Application

The question a carry trader faces is what to do with the signal. Section 2.3 gives three answers, unwinding the portfolio, reversing into anti-carry, and the downcurve overwrite, which stays invested in the months the differential pays for the risk. Table 8 reports what each does to the portfolio, as an overlay on the carry trade's total returns.[8] Figure 7 plots the cumulative returns.

Unwinding on the signal raises the carry trade's Sharpe ratio from 0.53 to 0.73 and cuts

[7] Across all 458 circular placements of the signal's run structure, which hold its length and episode pattern fixed, a state covering four or more of the five worst months arises six times, an exact p of 0.013. The bottom-decile coverage does not clear the same test, at 13 of the 46 months inside and an exact p of 0.18, so the decile comparison rests on the regressions rather than on the count.

[8] All returns in this paper are reported before transaction costs, and the strategies of Table 8 do not account for them.

Table 8: Trading the IYC signal. The same five strategies applied to the carry trade and to the beta HML portfolio (total returns, 1988:01–2026:02, $N = 458$ months, no transaction costs). Mean is annualized (%/yr); $ES_{5\%}$ and the worst month are in %/month; the maximum drawdown is the largest peak-to-trough cumulative loss (%). The overwrite strategies stay invested when the average curve is in the downcurve configuration. The last three columns give each strategy's Sharpe-ratio, mean-return and skewness difference from always being in the trade, with the p-values beneath. Asy is the asymptotic p-value from the Ledoit and Wolf (2008) standard error and Boot the studentized circular block-bootstrap p-value at a block length of six months ($B = 4{,}999$), block lengths of three and twelve giving the same conclusions; skewness, which has no convenient asymptotic standard error under serial dependence, is tested by the bootstrap alone.

	The Strategy						Difference from Always In		
Strategy	Mean	Sharpe	Skew	$ES_{5\%}$	Max DD	Worst Mo.	ΔSR	ΔMean	ΔSkew
Panel A: The Carry Trade (Total Returns)									
Always in	+4.19	0.53	-0.68	-5.30	-30.1	-11.61			
Unwind on signal	+4.93	0.73	-0.34	-4.48	-25.9	-7.27	+0.20	+0.75	+0.35
Asy[p]							0.071	0.293	
Boot[p]							0.122	0.292	0.164
Anti-carry on signal	+5.68	0.73	-0.02	-4.71	-25.9	-7.27	+0.20	+1.49	+0.66
Asy[p]							0.287	0.293	
Boot[p]							0.318	0.292	0.139
Overwrite: unwind unless downcurve	+5.45	0.78	-0.33	-4.51	-25.9	-7.27	+0.25	+1.26	+0.36
Asy[p]							0.022	0.057	
Boot[p]							0.079	0.066	0.115
Overwrite: anti-carry unless downcurve	+6.71	0.87	-0.01	-4.59	-25.9	-7.27	+0.34	+2.52	+0.68
Asy[p]							0.051	0.057	
Boot[p]							0.094	0.066	0.119
Panel B: The Beta HML Portfolio (Total Returns)									
Always in	+2.54	0.29	-0.27	-5.62	-30.1	-12.51			
Unwind on signal	+4.22	0.56	+0.20	-4.45	-14.0	-6.55	+0.13	+0.73	+0.47
Asy[p]							0.044	0.114	
Boot[p]							0.048	0.130	0.036
Anti-carry on signal	+5.90	0.68	+0.32	-4.73	-14.8	-6.55	+0.19	+1.46	+0.59
Asy[p]							0.147	0.114	
Boot[p]							0.125	0.130	0.124
Overwrite: unwind unless downcurve	+4.16	0.52	+0.02	-4.94	-22.1	-7.64	+0.18	+1.18	+0.29
Asy[p]							0.001	0.003	
Boot[p]							0.011	0.008	0.178
Overwrite: anti-carry unless downcurve	+5.79	0.67	+0.19	-4.97	-22.1	-7.64	+0.31	+2.36	+0.46
Asy[p]							0.004	0.003	
Boot[p]							0.012	0.008	0.309

its worst month from −11.61 to −7.27%, and the two overwrites raise it further, to 0.78 and 0.87, by staying invested through the months the differential pays for. The skewness moves with it, from −0.68 always invested to −0.34 under the unwind and to −0.01 under the anti-carry overwrite, so the negative skewness that defines the carry trade in the literature (Brunnermeier et al., 2008) is largely removed. The gain is economically meaningful on both counts, though it is marginal in statistical terms for the carry portfolio and clearer for the beta HML portfolio.

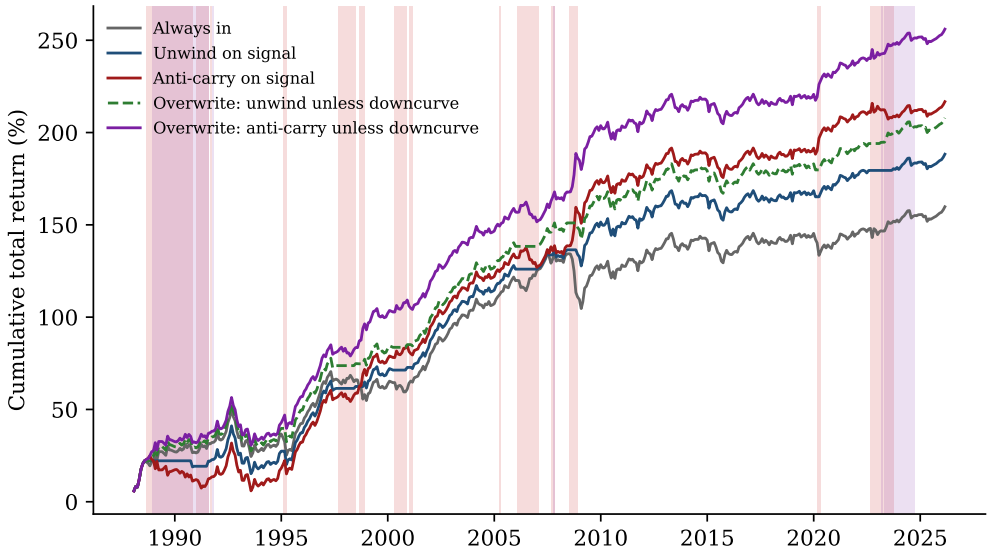


Figure 7: Cumulative total returns of the strategies in Table 8. Pink bands mark the IYC signal episodes, in which the unwind and anti-carry strategies depart from the always-invested line, and purple bands the downcurve months, in which the overwrite strategies differ from the plain unwind and anti-carry, most visibly in 1988–91 and 2023–24.

The downcurve split explains the asymmetry between the two hedges, and it is priced on the same core portfolio in every state (Appendix B.6). Over the full sample the regime does not predict returns, but inside signal months the gap between the two sets is +16.68%/yr (wild bootstrap $p = 0.05$). The spot leg still depreciates in downcurve months, so the regime does not remove the risk, though the crash is a fraction of its size elsewhere and the month's income covers 40% of the loss against 17% in the other signal months. The overwrite exploits that difference, and a rule that unwinds whenever risk is elevated, without asking whether it

is compensated, forgoes about two-fifths of the mean improvement over the always-invested position.

4 The Signal against Its Rivals

Section 3 estimates (12) without controls. This section puts the competing states and the standard risk gauges into the same regression,

$$r_{t+1} \;=\; \alpha \;+\; \gamma\, \mathrm{IYC}_t \;+\; \delta' W_t \;+\; e_{t+1}, \tag{16}$$

where W_t collects the controls named in each table, dated t apart from the two contemporaneous currency factors of Section 4.2.

4.1 Common versus Idiosyncratic, and versus the U.S. Curve

A natural concern about the signal is that it may be no more than a proxy for the global slope or for the U.S. curve, which is often deemed a global proxy. The concern is fair, because the signal aggregates those same slopes in the first place. It differs from both by construction, in the way Section 2.2 sets out, and Table 9 confirms that the difference matters.

Racing the signal against the level of the global average slope and its one-month change, the U.S. inversion, and the level and one-month change of the U.S. slope leaves none of the five rivals statistically significant. Adding them in fact strengthens the signal's own coefficient in every leg.

The third candidate is a currency's own curve, which the domestic interpretation of Section 2.2 would make the natural conditioning variable, and which matters here because bilateral exchange rates load heavily on common components (Verdelhan, 2018). The same conclusion holds currency by currency against the dollar, in Table B.10 of the appendix. A

Table 9: The portfolio horse race. Specification: $r_{t+1} = \alpha + \gamma\,\mathrm{IYC}_t + \delta' W_t + e_{t+1}$, with r the portfolio's annualized monthly return and W the rivals of the column, namely (1) none; (2) the continuous global average 10Y−2Y slope and its one-month change, the U.S. inversion (2s10s < 0, 57 months), and the U.S. 2s10s slope and its one-month change. Every regressor is dated t and the return is the following month's; 1988:01–2026:02. Each variable reports its coefficient, with the $\mathrm{NW}_{12}[p]$ row giving its Newey–West (12 lag) p and the WCB[p] row the wild cluster bootstrap p on signal runs ($B = 9{,}999$, with the null imposed on that coefficient in turn). Specification (1) is the signal on its own, repeated from Table 3 as a baseline for the adjusted R^2 of the second column.

	Carry Spot		Carry Total		Beta HML Spot		Beta HML Total	
	(1)	(2)	(1)	(2)	(1)	(2)	(1)	(2)
IYC signal	−11.63	−14.62	−9.86	−13.10	−12.82	−17.31	−13.58	−18.20
$\mathrm{NW}_{12}[p]$	0.003	0.002	0.013	0.006	0.001	0.000	0.000	0.000
WCB[p]	0.001	0.003	0.007	0.020	0.006	0.005	0.001	0.002
Global slope		+3.39		+2.48		−1.73		−1.25
$\mathrm{NW}_{12}[p]$		0.329		0.492		0.651		0.736
WCB[p]		0.311		0.528		0.775		0.823
Δ Global slope (1m)		−36.54		−35.31		−15.18		−14.97
$\mathrm{NW}_{12}[p]$		0.062		0.071		0.477		0.479
WCB[p]		0.287		0.290		0.554		0.560
U.S. inversion		+2.75		+2.19		+1.06		+1.95
$\mathrm{NW}_{12}[p]$		0.513		0.611		0.798		0.629
WCB[p]		0.616		0.692		0.847		0.719
U.S. slope		−3.07		−2.84		−2.01		−2.23
$\mathrm{NW}_{12}[p]$		0.144		0.184		0.254		0.191
WCB[p]		0.185		0.227		0.424		0.365
Δ U.S. slope (1m)		+21.98		+20.99		+10.24		+10.08
$\mathrm{NW}_{12}[p]$		0.073		0.086		0.443		0.450
WCB[p]		0.221		0.237		0.490		0.505
$\bar{R}^2$ (%)	2.7	4.2	1.9	3.1	2.7	2.5	3.0	2.9
N	458	458	458	458	458	458	458	458

currency whose own signal is active while the IYC signal is off does not depreciate, at a pooled +2.90%/yr and insignificant, and when currencies do depreciate on the IYC signal the losses fall on the high-beta names, not on the countries whose curves are in the count. [9]

Table 10: The signal against the rivals, compact form. Each cell is the coefficient γ on the signal in $r_{t+1} = \alpha + \gamma\,\mathrm{IYC}_t + \delta' W_t + e_{t+1}$ for the leg named by the row, with the control blocks entering cumulatively across the columns: (i) none; (ii) + financial conditions; (iii) + the dollar and commodity returns; (iv) + FX volatility; (v) + the leading indicator; (vi) + equity volatility. Re-estimated on the shorter column-(vi) sample, the signal alone is larger than its full-sample value in both carry legs, so the weakening in that column comes from the controls and not from the sample. Appendix B.3, Tables B.8 and B.9, gives every control coefficient and re-runs the full control set on the downcurve-filtered state. Each coefficient reports its Newey–West (12 lag) p and the wild cluster bootstrap p on signal runs ($B = 9{,}999$); 1988:01–2026:02.

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(i)′
Carry, spot	−11.63	−10.60	−9.82	−10.08	−10.11	−7.39	−13.05
$\mathrm{NW}_{12}[p]$	0.003	0.000	0.000	0.001	0.001	0.038	0.010
WCB$[p]$	0.001	0.001	0.001	0.002	0.002	0.062	0.004
Carry, total	−9.86	−9.04	−8.27	−8.66	−8.70	−6.54	−12.17
$\mathrm{NW}_{12}[p]$	0.013	0.002	0.004	0.003	0.003	0.061	0.015
WCB$[p]$	0.007	0.003	0.012	0.013	0.012	0.098	0.009
Beta HML, spot	−12.82	−12.63	−11.90	−12.15	−12.20	−12.55	−17.07
$\mathrm{NW}_{12}[p]$	0.001	0.000	0.000	0.000	0.000	0.001	0.000
WCB$[p]$	0.006	0.004	0.004	0.002	0.002	0.003	0.001
Beta HML, total	−13.58	−13.67	−12.96	−13.29	−13.34	−13.00	−17.33
$\mathrm{NW}_{12}[p]$	0.000	0.000	0.000	0.000	0.000	0.001	0.000
WCB$[p]$	0.001	0.000	0.001	0.001	0.001	0.002	0.001
N	458	458	458	458	458	423	423

4.2 Global Risk Gauges and Currency Factors

Another way to test the signal's robustness is to put it against the predictors the literature already uses. One might argue that it is only a proxy for the VIX, for funding stress, or for the credit cycle. Bakshi and Panayotov (2013) also time the carry trade with a commodity return, an average currency volatility, and a funding-liquidity measure. To test the signal against all of these, I run a ladder of nested specifications, where each column adds a new block of controls and keeps the ones before it. I order the blocks by how far back their series run, so that only the last column loses months. Table 10 gives the compact version and

[9] Individual-currency inference is weak by construction, since each regression uses one currency and fourteen episodes. The evidence for the mechanism comes instead from the cross-sectional alignment of these coefficients with beta (Section 3.2).

Appendix B.3, Tables B.8 and B.9, the full ones.[10]

Through column (v), where the sample still holds at all 458 months, no cell of the two portfolios falls by more than a sixth, at a worst-case wild p of 0.013. The only column that moves anything is (vi), where the VIX enters and the sample shortens by thirty-five months. There the carry cells clear 10% but not 5%, so they are marginal and not established, while the beta HML cells are unmoved. Applying the downcurve filter would restore the power for carry. Column (vi)′ of Tables B.8 and B.9 repeats the full control set on the non-downcurve signal, and the carry cells return to -10.61 and -9.70 at bootstrap p-values of 0.017 and 0.034.

Several of these variables do predict the carry trade on their own. The change in financial conditions and the two FX volatility terms are significant on both standards, and the adjusted R^2 rises from 2.7% with the state alone to 27.13% with the full set. The dollar factor of Lustig et al. (2011) and the FX volatility innovation of Menkhoff et al. (2012) enter contemporaneously with the return, so the columns holding them fixed are spanning tests, and the volatility innovation is the strongest term in the table, although it is built from the absolute daily moves of the same nine currencies whose returns form the portfolio and is mechanically high in the months the carry trade loses. The signal survives all of them, so it carries information these variables do not.

Table 11: EM spot returns on the IYC signal, per currency. Specification: $\Delta s_{i,t+1} = \alpha_i + \gamma_i \, \mathrm{IYC}_t + e_{i,t+1}$ (1996:02–2026:02; annualized). Each currency reports its sample-average money-market differential vs. the U.S. (percentage points), its sample-average currency beta, its coefficient, the Newey–West (12 lag) p and the wild cluster bootstrap p on signal runs ($B = 9{,}999$); the pooled row adds currency fixed effects.

Currency	Rate Diff.	Beta	γ	$\mathrm{NW}_{12}[p]$	$\mathrm{WCB}[p]$
BRL	+11.75	+0.29	−6.62	0.449	0.479
CLP	+3.15	+0.17	−11.04	0.180	0.166
IDR	+7.40	+0.06	−28.22	0.246	0.547
INR	+4.87	+0.06	−7.25	0.050	0.145
MXN	+7.47	+0.30	−11.44	0.203	0.252
KRW	+1.54	+0.12	−17.10	0.111	0.170
ZAR	+5.89	+0.31	−15.39	0.059	0.024
Pooled (FE)			−13.87	0.116	0.175

5 Emerging Market Extensions

5.1 Emerging Market Currencies: Beta and IRD

The natural extension is to the emerging markets, where the IRDs the carry trade collects are several times the G10's. The average IRD against the dollar reaches +11.75 percentage points for the BRL and exceeds +7 for the IDR and the MXN (Table 11), against a G10 maximum of +2.29 for the New Zealand dollar (Table 4). The compensation comes with the tail that the peso-problem literature attaches to it (Burnside et al., 2011), and UIP behaves differently in these currencies to begin with (Bansal and Dahlquist, 2000).

The design question is whether to take the signal built from the G10 curves to these currencies or to rebuild it from their own. EM curves carry sovereign and political risk in their expected short-rate paths that G10 curves do not, so a count over them measures

[10]Their volatility input is option-implied and no such series exists here before 2004, so I substitute a realized measure that correlates 0.92 in innovations with it over the overlap. Their liquidity measure is the TED spread, which the LIBOR transition ends in 2022:02, so it would cost the 2022–23 episode, and the financial-conditions pair spans it in any case. Neither their commodity return nor their lagged volatility innovation carries anything ($p \geq 0.22$).

something other than the common expected slowdown of Section 2.2. The financial cycle that governs these currencies also originates in the developed markets (Rey, 2013; Miranda-Agrippino and Rey, 2020). I therefore take the G10 signal as the baseline and report the EM-native alternative in Appendix B.8. Holding the signal fixed also makes the exercise an out-of-sample test of the global interpretation, because nothing in the EM data went into building it.

The universe is the seven floating EM currencies of Section 2.5, floating because a peg suppresses the response being predicted, following the classification of Ilzetzki et al. (2019). Table 11 repeats the per-currency spot regression of Section 3.2 on them. All seven load negatively on the signal and the pooled coefficient is slightly larger than the G10 beta HML portfolio's, so the result holds qualitatively. Statistical power, however, is weaker. These currencies are far more volatile than the G10's, and the standard errors widen with them.

What does not carry over from the G10 is the ordering in average beta and average IRD. Neither the seven EM currencies nor all sixteen reproduce the G10 gradient of Table 5 (Appendix B.7, Table B.15), and all seven depreciate together. A likely reason is that these currencies do not hold a stable position in the cross-section. Political instability and sovereign risk move their exposures over time in a way the G10's do not, and their IRDs compensate domestic inflation as well as global exposure.

5.2 The Portfolio with Emerging Markets

In the portfolio setting, I combine the seven EM currencies with the nine G10 and build the same three-long, three-short carry and beta portfolios of (3) over the combined sixteen. This shortens the sample to 1996:02–2026:02, for the reason given in Section 2.5. The resulting portfolio's long leg is dominated by the EM currencies, as their IRDs and betas imply, and they take 99% of the carry portfolio's long-leg slots and 72% of the beta HML portfolio's.

Table 12: Conditional moments with and without EM, 1996:02–2026:02. For each portfolio and leg, monthly returns off and on the (lagged) IYC signal, summarized by months, annualized mean (%/yr) and Sharpe ratio, skewness, and the monthly 5% value-at-risk, 5% expected shortfall, and worst month (%/month). Panel A: the paper's G10 portfolios on this subsample; Panel B: the same sorts on the sixteen-currency universe.

Series	State	N	Mean (%/yr)	Sharpe	Skew	$\text{VaR}_{5\%}$	$\text{ES}_{5\%}$	Worst
Panel A: The G10 Portfolio								
Carry, spot	No signal	304	+2.63	0.39	−0.41	−3.35	−4.49	−6.04
	Signal on	57	−11.38	−1.13	−1.23	−6.80	−8.63	−12.09
Carry, total	No signal	304	+5.84	0.86	−0.39	−3.15	−4.21	−5.78
	Signal on	57	−7.06	−0.70	−1.22	−6.36	−8.27	−11.61
Beta HML, spot	No signal	304	+2.11	0.27	−0.26	−3.87	−4.98	−6.84
	Signal on	57	−12.01	−1.24	−1.30	−5.16	−8.60	−12.77
Beta HML, total	No signal	304	+3.78	0.48	−0.23	−3.68	−4.80	−6.55
	Signal on	57	−10.36	−1.07	−1.24	−5.02	−8.40	−12.51
Panel B: The G10+EM Portfolio (16 Currencies)								
Carry, spot	No signal	304	+0.28	0.03	−0.80	−5.26	−7.01	−19.25
	Signal on	57	−12.15	−0.61	−1.64	−10.77	−17.16	−26.18
Carry, total	No signal	304	+11.26	1.04	−0.33	−4.08	−5.73	−16.55
	Signal on	57	+3.75	0.19	−1.17	−9.35	−15.02	−23.43
Beta HML, spot	No signal	304	+0.85	0.09	−0.63	−5.03	−6.80	−12.53
	Signal on	57	−23.44	−1.26	−1.99	−11.79	−17.05	−25.92
Beta HML, total	No signal	304	+6.34	0.63	−0.49	−4.64	−6.23	−12.00
	Signal on	57	−14.85	−0.82	−1.68	−10.26	−15.95	−23.88

Meanwhile, the funding legs are relatively unchanged, because no EM currency here sits low enough on either the IRD or beta to enter them.

The spot-leg result carries over qualitatively, as we can see in Table 12, where both portfolios crash on the signal and the beta-sorted one crashes harder than its G10 counterpart. The IRD, however, is large enough here that the carry trade's total return in signal months is still positive. This is the same idea as the downcurve overwrite of Section 3.5, except that the compensation now arrives in every state instead of only inside the regime. That compensation shows up in the mean but not in the tail. The worst month on the signal is about twice as deep as the G10's.

The beta-sorted portfolio is the stronger of the two here. Its total return in signal months is still deeply negative, where the carry portfolio's is not. Table 13 tells the same

Table 13: State regressions with and without EM. Specification: $r_{t+1} = \alpha + \gamma\,\text{State}_t + e_{t+1}$, r the portfolio's annualized monthly return, for the signal, its non-downcurve months, and its downcurve months; 1996:02–2026:02. Each portfolio reports its coefficient in %/yr, the Newey–West (12 lag) p on the row beneath it, and the wild cluster bootstrap p on the state's episodes on the row beneath that ($B = 9{,}999$).

	Signal	Non-downcurve	Downcurve
Panel A: The G10 Portfolio			
Carry, spot	−14.01	−17.00	+8.68
$\text{NW}_{12}[p]$	0.015	0.006	0.010
WCB[p]	0.008	0.003	0.136
Carry, total	−12.91	−15.80	+8.90
$\text{NW}_{12}[p]$	0.023	0.010	0.010
WCB[p]	0.017	0.010	0.142
Beta HML, spot	−14.12	−17.17	+8.98
$\text{NW}_{12}[p]$	0.001	0.000	0.399
WCB[p]	0.003	0.000	0.430
Beta HML, total	−14.13	−17.32	+9.89
$\text{NW}_{12}[p]$	0.001	0.000	0.360
WCB[p]	0.003	0.000	0.313
Panel B: The G10+EM Portfolio			
Carry, spot	−12.43	−16.31	+15.45
$\text{NW}_{12}[p]$	0.132	0.047	0.002
WCB[p]	0.223	0.121	0.084
Carry, total	−7.51	−10.58	+13.86
$\text{NW}_{12}[p]$	0.281	0.137	0.003
WCB[p]	0.351	0.191	0.093
Beta HML, spot	−24.29	−28.35	+8.04
$\text{NW}_{12}[p]$	0.023	0.012	0.313
WCB[p]	0.019	0.014	0.333
Beta HML, total	−21.19	−24.46	+5.30
$\text{NW}_{12}[p]$	0.014	0.007	0.539
WCB[p]	0.014	0.018	0.548

story. The signal regression holds up across the board for the beta HML portfolio, while for carry it weakens considerably and the cells no longer clear conventional significance. The downcurve split is still visible in both, with the loss concentrated in the non-downcurve months and the downcurve months positive in every leg. I interpret this as the signal acting mostly on beta risk exposure, which the IRD sort's monthly rebalancing may partly take into account. One potential explanation for this is that the market prices a wider IRD for these currencies because of the peso problem discussed in Section 5.1, so the EM carry trade ends up overcompensating for the risk the signal marks.

6 Conclusion

This paper contributes to the yield curve and international asset pricing literature by showing that clusters of fresh G10 yield curve inversions, identify periods of elevated currency downside risk. Taking into account the freshness of the signal, it separates inversions that still carry timing information from those that have gone stale. The depreciation falls on high-beta currencies, which also tend to have high interest rates, so the carry trade inherits this exposure through its long leg. The signal is active in about 20% of the sample and is able to capture the carry trade's five worst months. The predictive power remains after controlling for the average G10 slope, the U.S. yield curve, and other standard measures of global risk.

The paper also shows why elevated downside risk does not need to imply that the carry position should always be unwind. Inversions may reflect expected weakness in economic activity or high policy rates set to fight inflation that is expected to normalize. The shape of the average G10 curve helps to distinguish these environments. When the curve is strictly downward sloping ("downcurve"), the IRDs are unusually wide and, on average, compensate most of the subsequent depreciation. Therefore, the opportunity cost of unwinding in the downcurve regime is unusually expensive, and an investor might want to consider overwriting the signal. The evidence from EM currencies extension reinforces this trade-off, although their portfolio depreciation is slightly larger than the G10 portfolio, their wide IRDs offer enough compensation to the risk and remains profitable.

Lastly, the paper translates information from international yield curve into a monthly conditioning variable relevant to portfolio risk management. Much of the existing literature evaluates exchange-rate predictability over horizons of several months or quarters, making it unclear how those forecasts map into an investor's current position. The IYC signal identifies when common currency downside risk is elevated, while the shape of the average curve indicates whether the IRD provides adequate compensation. In conclusion, freshness

determines whether an inversion remains informative, clustering identifies whether its content is common across countries, and shape of the curve distinguishes compensated from uncompensated currency risk.

References

Accominotti, O., Cen, J., Chambers, D., and Marsh, I. W. (2019). Currency regimes and the carry trade. *Journal of Financial and Quantitative Analysis*, 54(5), 2233–2260.

Adrian, T., Crump, R. K., and Moench, E. (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics*, 110(1), 110–138.

Andrews, S., Colacito, R., Croce, M. M., and Gavazzoni, F. (2024). Concealed carry. *Journal of Financial Economics*, 159, 103874.

Ang, A. and Chen, J. S. (2010). Yield curve predictors of foreign exchange returns. Working paper, Columbia Business School.

Backus, D. K., Foresi, S., and Telmer, C. I. (2001). Affine term structure models and the forward premium anomaly. *Journal of Finance*, 56(1), 279–304.

Bakshi, G. and Panayotov, G. (2013). Predictability of currency carry trades and asset pricing implications. *Journal of Financial Economics*, 110(1), 139–163.

Bansal, R. and Dahlquist, M. (2000). The forward premium puzzle: Different tales from developed and emerging economies. *Journal of International Economics*, 51(1), 115–144.

Barro, R. J. (2006). Rare disasters and asset markets in the twentieth century. *Quarterly Journal of Economics*, 121(3), 823–866.

Barroso, P. and Santa-Clara, P. (2015). Beyond the carry trade: Optimal currency portfolios. *Journal of Financial and Quantitative Analysis*, 50(5), 1037–1056.

Bauer, M. D. and Mertens, T. M. (2018). Economic forecasts with the yield curve. *FRBSF Economic Letter*, 2018-07.

Brunnermeier, M. K., Nagel, S., and Pedersen, L. H. (2008). Carry trades and currency crashes. *NBER Macroeconomics Annual*, 23, 313–348.

Burnside, C., Eichenbaum, M., Kleshchelski, I., and Rebelo, S. (2011). Do peso problems explain the returns to the carry trade? *Review of Financial Studies*, 24(3), 853–891.

Caballero, R. J. and Doyle, J. B. (2012). Carry trade and systemic risk: Why are FX options so cheap? NBER Working Paper 18644.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements

for inference with clustered errors. *Review of Economics and Statistics*, 90(3), 414–427.

Campbell, J. Y. and Shiller, R. J. (1991). Yield spreads and interest rate movements: A bird's eye view. *Review of Economic Studies*, 58(3), 495–514.

Chen, Y.-C. and Tsang, K. P. (2013). What does the yield curve tell us about exchange rate predictability? *Review of Economics and Statistics*, 95(1), 185–205.

Chernov, M., Graveline, J., and Zviadadze, I. (2018). Crash risk in currency returns. *Journal of Financial and Quantitative Analysis*, 53(1), 137–170.

Christiansen, C., Ranaldo, A., and Söderlind, P. (2011). The time-varying systematic risk of carry trade strategies. *Journal of Financial and Quantitative Analysis*, 46(4), 1107–1125.

Clarida, R. and Waldman, D. (2008). Is bad news about inflation good news for the exchange rate? In *Asset Prices and Monetary Policy*, University of Chicago Press.

Clarida, R., Davis, J., and Pedersen, N. (2009). Currency carry trade regimes: Beyond the Fama regression. *Journal of International Money and Finance*, 28(8), 1375–1389.

Colacito, R., Riddiough, S. J., and Sarno, L. (2020). Business cycles and currency returns. *Journal of Financial Economics*, 137(3), 659–678.

Della Corte, P., Riddiough, S. J., and Sarno, L. (2016). Currency premia and global imbalances. *Review of Financial Studies*, 29(8), 2161–2193.

Diebold, F. X., Li, C., and Yue, V. Z. (2008). Global yield curve dynamics and interactions: A dynamic Nelson–Siegel approach. *Journal of Econometrics*, 146(2), 351–363.

Dreher, F., Gräb, J., and Kostka, T. (2020). From carry trades to curvy trades. *The World Economy*, 43(3), 758–780.

Driscoll, J. C. and Kraay, A. C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4), 549–560.

Du, W., Tepper, A., and Verdelhan, A. (2018). Deviations from covered interest rate parity. *Journal of Finance*, 73(3), 915–957.

Engel, C. (2016). Exchange rates, interest rates, and the risk premium. *American Economic Review*, 106(2), 436–474.

Engstrom, E. C. and Sharpe, S. A. (2019). The near-term forward yield spread as a leading indicator: A less distorted mirror. *Financial Analysts Journal*, 75(4), 37–49.

Estrella, A. and Hardouvelis, G. A. (1991). The term structure as a predictor of real economic activity. *Journal of Finance*, 46(2), 555–576.

Estrella, A. and Mishkin, F. S. (1997). The predictive power of the term structure of interest rates in Europe and the United States: Implications for the European Central Bank. *European Economic Review*, 41(7), 1375–1401.

Estrella, A. and Mishkin, F. S. (1998). Predicting U.S. recessions: Financial variables as leading indicators. *Review of Economics and Statistics*, 80(1), 45–61.

Fama, E. F. (1984). Forward and spot exchange rates. *Journal of Monetary Economics*, 14(3), 319–338.

Farhi, E. and Gabaix, X. (2016). Rare disasters and exchange rates. *Quarterly Journal of Economics*, 131(1), 1–52.

Farhi, E., Fraiberger, S. P., Gabaix, X., Rancière, R., and Verdelhan, A. (2015). Crash risk in currency markets. NBER Working Paper 15062.

Gabaix, X. and Maggiori, M. (2015). International liquidity and exchange rate dynamics. *Quarterly Journal of Economics*, 130(3), 1369–1420.

Hansen, L. P. and Hodrick, R. J. (1980). Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *Journal of Political Economy*, 88(5), 829–853.

Harvey, C. R. (1988). The real term structure and consumption growth. *Journal of Financial Economics*, 22(2), 305–333.

Harvey, C. R. (1991). The term structure and world economic growth. *Journal of Fixed Income*, 1(1), 7–19.

Hassan, T. A. (2013). Country size, currency unions, and international asset returns. *Journal of Finance*, 68(6), 2269–2308.

Hassan, T. A. and Mano, R. C. (2019). Forward and spot exchange rates in a multi-currency world. *Quarterly Journal of Economics*, 134(1), 397–450.

He, Z., Kelly, B., and Manela, A. (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1), 1–35.

Ilzetzki, E., Reinhart, C. M., and Rogoff, K. S. (2019). Exchange arrangements entering the twenty-first century: Which anchor will hold? *Quarterly Journal of Economics*, 134(2), 599–646.

Jotikasthira, C., Le, A., and Lundblad, C. (2015). Why do term structures in different currencies co-move? *Journal of Financial Economics*, 115(1), 58–83.

Jurek, J. W. (2014). Crash-neutral currency carry trades. *Journal of Financial Economics*, 113(3), 325–347.

Kaebi, M. M. and Ferreira Batista Martins, I. (2025). Long-run interest rate differentials and the profitability of currency carry. Örebro University Working Paper 10/2025.

Koijen, R. S. J., Moskowitz, T. J., Pedersen, L. H., and Vrugt, E. B. (2018). Carry. *Journal of Financial Economics*, 127(2), 197–225.

Kremens, L. (2024). Positioning risk. Working paper, University of Washington.

Kremens, L. and Martin, I. (2019). The quanto theory of exchange rates. *American Economic Review*, 109(3), 810–843.

Ledoit, O. and Wolf, M. (2008). Robust performance hypothesis testing with the Sharpe ratio. *Journal of Empirical Finance*, 15(5), 850–859.

Lettau, M., Maggiori, M., and Weber, M. (2014). Conditional risk premia in currency markets and other asset classes. *Journal of Financial Economics*, 114(2), 197–225.

Lloyd, S. P. and Marin, E. A. (2024). Exchange rates, yield curves and transitory risk. Working paper, November 2024. Previously circulated as "Exchange-Rate Risk and Business Cycles."

Lustig, H. and Verdelhan, A. (2007). The cross section of foreign currency risk premia and consumption growth risk. *American Economic Review*, 97(1), 89–117.

Lustig, H., Roussanov, N., and Verdelhan, A. (2011). Common risk factors in currency markets. *Review of Financial Studies*, 24(11), 3731–3777.

Lustig, H., Roussanov, N., and Verdelhan, A. (2014). Countercyclical currency risk premia. *Journal of Financial Economics*, 111(3), 527–553.

Lustig, H., Stathopoulos, A., and Verdelhan, A. (2019). The term structure of currency carry trade risk premia. *American Economic Review*, 109(12), 4142–4177.

MacKinnon, J. G., Nielsen, M. Ø., and Webb, M. D. (2023). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, 232(2), 272–299.

Menkhoff, L., Sarno, L., Schmeling, M., and Schrimpf, A. (2012). Carry trades and global foreign exchange volatility. *Journal of Finance*, 67(2), 681–718.

Miranda-Agrippino, S. and Rey, H. (2020). U.S. monetary policy and the global financial cycle. *Review of Economic Studies*, 87(6), 2754–2776.

Mönch, E. (2012). Term structure surprises: The predictive content of curvature, level, and slope. *Journal of Applied Econometrics*, 27(4), 574–602.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.

Ranaldo, A. and Söderlind, P. (2010). Safe haven currencies. *Review of Finance*, 14(3), 385–407.

Ready, R., Roussanov, N., and Ward, C. (2017). Commodity trade and the carry trade: A tale of two countries. *Journal of Finance*, 72(6), 2629–2684.

Rey, H. (2013). Dilemma not trilemma: The global financial cycle and monetary policy independence. In *Global Dimensions of Unconventional Monetary Policy*, Jackson Hole Economic Policy Symposium, Federal Reserve Bank of Kansas City.

Richmond, R. J. (2019). Trade network centrality and currency risk premia. *Journal of*

Finance, 74(3), 1315–1361.

Rietz, T. A. (1988). The equity risk premium: A solution. *Journal of Monetary Economics*, 22(1), 117–131.

Verdelhan, A. (2010). A habit-based explanation of the exchange rate risk premium. *Journal of Finance*, 65(1), 123–146.

Verdelhan, A. (2018). The share of systematic variation in bilateral exchange rates. *Journal of Finance*, 73(1), 375–418.

Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance*, 68(3), 987–1035.

Wright, J. H. (2011). Term premia and inflation uncertainty: Empirical evidence from an international panel dataset. *American Economic Review*, 101(4), 1514–1534.

A Economic Regimes

Table A reports statistics across the different eras. The sample opens in the post-inflation high-rate world, with average G10 three-month yields of 8.5%, and passes through the Great Moderation and the pre-crisis expansion. It then enters the Great Recession and a decade at the zero lower bound, during which the carry trade's IRD thins to +2.49%/yr and the signal falls silent for ten years. Average short rates are negative at the COVID trough. It ends in the post-COVID inflation era, where rates, inflation, the signal and the downcurve regime all return.

B Robustness

This appendix covers robustness checks on the signal's construction, competing variables, and subsample stability.

Table A: The sample's economic regimes. For each era the table reports the number of months, the average G10 three-month yield (%), G10 CPI inflation (YoY %, nine currencies), the carry trade's income and total returns (annualized, %/yr), and the number of signal and downcurve months. Era boundaries are calendar conventions.

Era	Months	3M Yield	CPI YoY	Income	Total	Signal	Downcurve
Post-inflation high rates (1988–94)	84	8.52	3.73	+5.94	+5.24	34	33
Great moderation / dot-com (1995–2000)	72	4.52	1.44	+4.86	+4.59	23	0
Pre-crisis expansion (2001–07)	84	3.41	1.66	+4.42	+9.56	16	1
Great Recession (2008–09)	24	2.66	2.03	+3.85	−1.36	5	0
Zero lower bound (2010–19)	120	0.80	1.39	+2.49	+1.70	0	0
COVID (2020–21)	24	−0.05	1.25	+0.94	−0.17	2	0
Post-COVID inflation (2022–26)	50	2.60	4.02	+3.11	+3.54	13	18

B.1 Signal Construction

This appendix reports the signal under alternative constructions, namely different thresholds for clustering and for steepening, a different tenor combination, a different number of currency pairs, and the U.S. curve and the dollar added to the count. Appendix B.2 then takes up the naive-count benchmark and an alternative freshness filter.

Replacing the 2-year tenor with the 3-month bill (Table B.1) preserves every sign at roughly two-thirds the size. The carry-spot and both beta HML coefficients remain significant, while carry total is marginal on episode inference ($p = 0.062$). The 10Y−3M signal is also a substantially different classifier (104 months, only 52 shared with the baseline). The 3-month end is dominated by the realized policy rate, so it inverts whenever policy sits above its long-run average (Engstrom and Sharpe, 2019).

On the cluster size (Table B.2), accepting *any single inversion* dilutes the state to 169 months, or 37% of the sample, and washes out the carry effect (−4.02%/yr, $p = 0.16$). To check for a country-specific effect on the portfolio, I isolate the months with *exactly one* active inversion, and those turn out to be positive carry months (+6.84 spot, +7.08 total, $p \leq 0.021$). The "any inversion" state therefore averages a benign idiosyncratic state with the crash state. The lone inverter is a high-rate country, and it sits in the carry portfolio's

Table B.1: Slope-definition robustness: the signal rebuilt on 10Y−3M slopes instead of 10Y−2Y, identical activation and deactivation rules. Specification as in Table 3.

Slope Definition	ON Mo.	Carry Spot	Carry Total	Beta HML Spot	Beta HML Total
10Y−2Y (baseline)	93	−11.63	−9.86	−12.82	−13.58
$NW_{12}[p]$		0.003	0.013	0.001	0.000
WCB[p]		0.001	0.007	0.006	0.001
10Y−3M	104	−7.46	−6.73	−8.41	−8.89
$NW_{12}[p]$		0.054	0.081	0.038	0.024
WCB[p]		0.035	0.062	0.024	0.012
10Y−2Y, excl. downcurve	56	−17.35	−16.51	−17.93	−18.03
$NW_{12}[p]$		0.002	0.003	0.000	0.000
WCB[p]		0.002	0.004	0.001	0.001
10Y−3M, excl. downcurve	69	−9.66	−9.94	−10.10	−10.23
$NW_{12}[p]$		0.062	0.048	0.031	0.025
WCB[p]		0.038	0.036	0.023	0.018

The two signals overlap in 52 months. The 10Y−3M downcurve months (35 of them) are flat, so the downcurve filter behaves the same way on either construction.

Table B.2: Cluster-size grid: the state at each activation threshold and at each exact count of active inversions. ON Mo. counts the position months; each cell is the state regression of Table 3, $r_{t+1} = \alpha + \gamma\,\mathrm{State}_t + e_{t+1}$ (annualized, %/yr) over 1988:01–2026:02, with the Newey–West (12 lag) p and the wild cluster bootstrap p on the state's own runs on the labelled rows beneath each coefficient ($B = 9{,}999$). The paper's baseline is the nested ≥ 2 rule.

Cluster Size	ON Mo.	Carry Spot	Carry Total	Beta HML Spot	Beta HML Total
≥ 1 (any inversion)	169	−4.02	−2.65	−9.52	−9.96
$NW_{12}[p]$		0.125	0.308	0.000	0.000
WCB[p]		0.155	0.374	0.003	0.001
$= 1$ (exactly one)	76	+6.84	+7.08	−1.03	−0.88
$NW_{12}[p]$		0.003	0.002	0.712	0.751
WCB[p]		0.014	0.021	0.762	0.797
≥ 2 (baseline)	93	−11.63	−9.86	−12.82	−13.58
$NW_{12}[p]$		0.003	0.013	0.001	0.000
WCB[p]		0.001	0.007	0.006	0.001
$= 2$ (exactly two)	41	−15.16	−14.23	−16.82	−17.16
$NW_{12}[p]$		0.006	0.011	0.002	0.001
WCB[p]		0.010	0.018	0.003	0.002
≥ 3	52	−6.43	−4.33	−6.99	−7.93
$NW_{12}[p]$		0.041	0.184	0.072	0.031
WCB[p]		0.116	0.264	0.182	0.081
$= 3$ (exactly three)	22	−11.63	−10.45	−8.98	−9.49
$NW_{12}[p]$		0.064	0.089	0.010	0.004
WCB[p]		0.197	0.245	0.065	0.032
≥ 4	30	−1.89	+0.68	−4.78	−5.96
$NW_{12}[p]$		0.548	0.845	0.426	0.300
WCB[p]		0.320	0.862	0.465	0.315

long leg in 86% of those months against 36% in the beta HML portfolio's, so an idiosyncratic inversion moves the carry portfolio systematically and the beta HML portfolio hardly at all. Exactly two active inversions is the sharpest cell, and exactly three still gives a crash cell of essentially the baseline's size, though on only 22 months, which may be what leaves the carry cells short of significance there.

Raising the threshold to three weakens the results, and the row below it shows why, since at four or more the effect disappears in every leg. This may sound counterintuitive, because more inversions ought to mean more risk, but breadth also arrives later in an episode, by which time freshness has gone. Four of the thirty four-or-more months fall in an episode's first three, against about half of the two- and three-curve months, and twenty-four of the thirty are downcurve months.

Two things follow from the grid. The beta HML cells hold up in more of it than the carry cells, surviving at any single inversion and at exactly three where the carry cells do not. And the paper's threshold is not the product of data mining. A search over this grid would have chosen exactly two, which is the largest cell in all four columns; the nested ≥ 2 rule was chosen on economic grounds, and it is the one an investor can implement, because an exact-count state switches off in the month a third curve inverts.

The deactivation grid (Table B.3) behaves as one would expect. Deactivating on a single steepening month admits too much noise and weakens the results substantially, leaving a state of 40 scattered months, while requiring three consecutive months holds the signal on far too long, swelling it to 178 and diluting every cell. Once again the beta HML portfolio is the more robust of the two. It stays significant under the three-month rule, where the carry cells do not, and it keeps more of its size under the one-month rule.

Table B.3: Deactivation-confirmation grid: the signal rebuilt with the deactivation requiring $\bar{c}$ consecutive re-steepening months, which replaces the threshold $c_{i,t} \geq 2$ of (8) with $c_{i,t} \geq \bar{c}$ for $\bar{c} = 1, 2, 3$; identical activation rule. Specification as in Table 3.

Deactivation Rule	ON Mo.	Carry Spot	Carry Total	Beta HML Spot	Beta HML Total
1 re-steepening month	40	−2.94	−0.81	−5.78	−6.64
$NW_{12}[p]$		0.436	0.837	0.076	0.039
WCB$[p]$		0.416	0.843	0.114	0.074
2 consecutive (baseline)	93	−11.63	−9.86	−12.82	−13.58
$NW_{12}[p]$		0.003	0.013	0.001	0.000
WCB$[p]$		0.001	0.007	0.006	0.001
3 consecutive	178	−5.31	−3.90	−5.67	−5.80
$NW_{12}[p]$		0.075	0.190	0.019	0.015
WCB$[p]$		0.061	0.221	0.026	0.013

The number of currencies per leg. This section tests the result across portfolio constructions (Table B.4). Both portfolios use $k = 3$ in the paper (Section 2.1), and running them at one, two and four pairs per leg leaves the picture unchanged, with signs and magnitudes stable across widths. What weakens is inference in the narrow carry sorts, where carry spot is marginal at one pair and carry total insignificant at one and at two, while the beta HML cells hold at every width, the narrowest spot cell only just. That is what one would expect, since a one-pair sort is a bilateral trade and carries country-specific risk and not the common exposure the signal is about.

Adding the U.S. curve and the dollar. Table B.5 repeats the headline cells with the U.S. included on both margins, in the curve construction and by adding the dollar to the portfolio.

Adding the U.S. curve changes none of the results. The baseline signal's 93 months are a strict subset of the ten-curve signal's 99, and every combination of portfolio and signal remains significant. Those six extra months leave every coefficient where it was, so most of the common signal comes from outside the U.S. curve.

Table B.4: Reconstructing the portfolio with 1×1, 2×2, 3×3, and 4×4. Cells are the state regressions of Table 3, $r_{t+1} = \alpha + \gamma \, \text{State}_t + e_{t+1}$ (annualized, %/yr) over 1988:01–2026:02, with the wild cluster bootstrap p on the state's runs on the WCB[p] row and its Newey–West p on the $\text{NW}_{12}[p]$ row ($B = 9{,}999$).

		1×1	2×2	3×3	4×4
Carry, spot	IYC signal	−10.02	−11.00	−11.63	−10.14
	$\text{NW}_{12}[p]$	0.085	0.016	0.003	0.001
	WCB[p]	0.088	0.005	0.001	0.001
	non-downcurve	−19.51	−18.91	−17.35	−14.70
	$\text{NW}_{12}[p]$	0.019	0.006	0.002	0.001
	WCB[p]	0.024	0.005	0.002	0.001
	downcurve	+6.37	+3.37	−0.29	−0.86
	$\text{NW}_{12}[p]$	0.291	0.301	0.918	0.694
	WCB[p]	0.488	0.241	0.905	0.858
Carry, total	IYC signal	−7.74	−8.97	−9.86	−8.77
	$\text{NW}_{12}[p]$	0.200	0.059	0.013	0.007
	WCB[p]	0.276	0.112	0.007	0.005
	non-downcurve	−18.17	−17.76	−16.51	−14.07
	$\text{NW}_{12}[p]$	0.029	0.010	0.003	0.002
	WCB[p]	0.047	0.014	0.004	0.001
	downcurve	+9.39	+6.10	+2.37	+1.23
	$\text{NW}_{12}[p]$	0.153	0.081	0.351	0.526
	WCB[p]	0.497	0.273	0.396	0.594
Beta HML, spot	IYC signal	−13.22	−15.32	−12.82	−10.73
	$\text{NW}_{12}[p]$	0.013	0.001	0.001	0.002
	WCB[p]	0.051	0.011	0.006	0.026
	non-downcurve	−21.49	−22.40	−17.93	−15.27
	$\text{NW}_{12}[p]$	0.001	0.000	0.000	0.000
	WCB[p]	0.007	0.001	0.001	0.004
	downcurve	+2.26	−1.00	−2.01	−1.32
	$\text{NW}_{12}[p]$	0.754	0.857	0.740	0.810
	WCB[p]	0.741	0.892	0.858	0.815
Beta HML, total	IYC signal	−14.87	−16.29	−13.58	−11.44
	$\text{NW}_{12}[p]$	0.004	0.000	0.000	0.001
	WCB[p]	0.009	0.001	0.001	0.009
	non-downcurve	−21.70	−22.22	−18.03	−15.42
	$\text{NW}_{12}[p]$	0.001	0.000	0.000	0.000
	WCB[p]	0.007	0.001	0.001	0.003
	downcurve	−1.05	−3.38	−3.53	−2.65
	$\text{NW}_{12}[p]$	0.892	0.548	0.567	0.633
	WCB[p]	0.931	0.821	0.864	0.788

Table B.5: Headline cells with the U.S. included in the signal count and/or the portfolio. State regressions $r_{t+1} = \alpha + \gamma\,\mathrm{IYC}_t + e_{t+1}$, with the Newey–West (12 lag) p and the wild cluster bootstrap p on contiguous signal runs ($B = 9{,}999$); no controls.

Portfolio	Signal	Leg	γ	$\mathrm{NW}_{12}[p]$	WCB$[p]$
G10 carry	Baseline (9 curves)	spot	−11.63	0.003	0.001
G10 carry	Baseline (9 curves)	total	−9.86	0.013	0.007
G10 carry	+US curve (10)	spot	−11.11	0.003	0.001
G10 carry	+US curve (10)	total	−9.43	0.013	0.008
Dollar-incl. carry	Baseline (9 curves)	spot	−11.83	0.007	0.000
Dollar-incl. carry	Baseline (9 curves)	total	−10.04	0.026	0.016
Dollar-incl. carry	+US curve (10)	spot	−11.56	0.006	0.000
Dollar-incl. carry	+US curve (10)	total	−9.87	0.023	0.012
G10 beta HML	Baseline (9 curves)	spot	−12.82	0.001	0.006
G10 beta HML	Baseline (9 curves)	total	−13.58	0.000	0.001
G10 beta HML	+US curve (10)	spot	−12.49	0.001	0.005
G10 beta HML	+US curve (10)	total	−13.18	0.000	0.001
Dollar-incl. beta HML	Baseline (9 curves)	spot	−11.24	0.004	0.036
Dollar-incl. beta HML	Baseline (9 curves)	total	−11.93	0.002	0.010
Dollar-incl. beta HML	+US curve (10)	spot	−10.84	0.005	0.030
Dollar-incl. beta HML	+US curve (10)	total	−11.47	0.002	0.008

Signal months: baseline (9 curves) 93, +US curve (10) 99, overlap 93
USD sits in the carry funding leg 31% of months and in the beta short leg 49%
Three worst beta HML months (2008:10, 1991:03, 2020:03): 3 of 3 in the baseline signal, 3 of 3 in the ten-curve signal

B.2 The Naive-Count Benchmark

This section tests the signal against the simplest benchmark, clustering by the naive count.

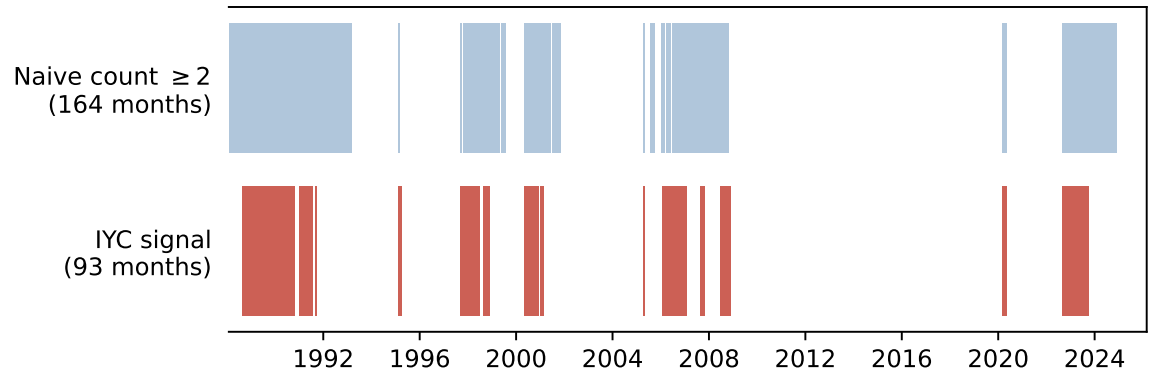


Figure B.1: IYC against the naive signal. The naive count (≥ 2 of 9 inverted, 164 months, top) is shown against this paper's signal (93 months, bottom), 1988–2026.

The naive count is on for 36% of the sample, nearly twice the signal's 20%, because once a cluster forms it tends to stand for years (Figure B.1). Staying on longer dilutes the state

and costs it accuracy on the tail. Table B.6 puts a number on that, where a *concentration* of 1.0 means a state captures tail months no better than chance. At the broader cuts the naive count does catch more tail months than the signal, but only by being on for nearly twice as much of the sample. Overall, a cluster of inversions catches tail months better than a random month whatever the construction, in both panels and at every cut, the one exception being the downcurve-filtered naive count at the extreme beta HML tail.

Another way to filter for freshness is to penalise inversions by their age. The simplest version weights the count by 1/age, which Table B.7 reports. The weighting puts the emphasis on the months just after a cluster forms, and it undoes the dilution the stale inversions cause, restoring the effect on every leg. That it works at all is evidence that most of the crashes arrive soon after the cluster forms.

B.3 Additional Robustness and Timing Checks

Controls. Tables B.8 and B.9 give the full regression tables behind Section 4.2, with every control coefficient and the adjusted R^2, while Table B.10 gives the currency-level version of the horse race discussed in Section 4.1.

Subsample stability. Table B.11 runs a leave-one-out study over the episodes, dropping each one's months in turn and re-estimating on what is left. Every cell keeps its sign and its significance in all fourteen deletions, the financial crisis and COVID among them, so the result is not the work of one or two episodes.

Table B.12 then splits the sample in half at the end of 2004. Because the rule was formulated with knowledge of the full sample, this is a stability check and not a held-out test, and both portfolios lose in both halves at nearly identical size. The downcurve regime

Table B.6: Tail capture per classifier. Concentration = capture share ÷ time share. Panel A takes the tail months to be the full-sample worst 1%, 5%, and 10% of carry total returns; Panel B the same cuts on the beta HML portfolio's spot returns, whose total-return tails select the same months because its income leg is too small to reorder them. 1988:01–2026:02.

Panel A: carry portfolio

Classifier	Months On	Time Share	Tail Captured	Concentration	In-State Rate
Worst 1% of Carry-Total Months (5 Months ≤ −6.03%/month)					
Naive count ≥ 2	164	36%	3/5 (60%)	1.68×	1.8%
Naive count excl. downcurve	112	24%	3/5 (60%)	2.45×	2.7%
IYC signal	93	20%	4/5 (80%)	3.94×	4.3%
IYC signal excl. downcurve	56	12%	4/5 (80%)	6.54×	7.1%
Worst 5% of Carry-Total Months (23 Months ≤ −3.75%/month)					
Naive count ≥ 2	164	36%	10/23 (43%)	1.21×	6.1%
Naive count excl. downcurve	112	24%	8/23 (35%)	1.42×	7.1%
IYC signal	93	20%	6/23 (26%)	1.28×	6.5%
IYC signal excl. downcurve	56	12%	5/23 (22%)	1.78×	8.9%
Worst 10% of Carry-Total Months (46 Months ≤ −2.32%/month)					
Naive count ≥ 2	164	36%	20/46 (43%)	1.21×	12.2%
Naive count excl. downcurve	112	24%	17/46 (37%)	1.51×	15.2%
IYC signal	93	20%	13/46 (28%)	1.39×	14.0%
IYC signal excl. downcurve	56	12%	12/46 (26%)	2.13×	21.4%

Panel B: beta HML portfolio

Classifier	Months On	Time Share	Tail Captured	Concentration	In-State Rate
Worst 1% of Beta-HML-Spot Months (5 Months ≤ −6.53%/month)					
Naive count ≥ 2	164	36%	2/5 (40%)	1.12×	1.2%
Naive count excl. downcurve	112	24%	1/5 (20%)	0.82×	0.9%
IYC signal	93	20%	3/5 (60%)	2.95×	3.2%
IYC signal excl. downcurve	56	12%	2/5 (40%)	3.27×	3.6%
Worst 5% of Beta-HML-Spot Months (23 Months ≤ −4.16%/month)					
Naive count ≥ 2	164	36%	12/23 (52%)	1.46×	7.3%
Naive count excl. downcurve	112	24%	7/23 (30%)	1.24×	6.2%
IYC signal	93	20%	9/23 (39%)	1.93×	9.7%
IYC signal excl. downcurve	56	12%	5/23 (22%)	1.78×	8.9%
Worst 10% of Beta-HML-Spot Months (46 Months ≤ −2.84%/month)					
Naive count ≥ 2	164	36%	19/46 (41%)	1.15×	11.6%
Naive count excl. downcurve	112	24%	13/46 (28%)	1.16×	11.6%
IYC signal	93	20%	14/46 (30%)	1.50×	15.1%
IYC signal excl. downcurve	56	12%	9/46 (20%)	1.60×	16.1%

Table B.7: Freshness filter on the naive count. State regressions $r_{t+1} = \alpha + \gamma\,\mathrm{State}_t + e_{t+1}$ as in Table 3, with the conditioning variable the naive count, the count restricted to non-downcurve months, the count weighted by 1/age of the standing cluster, and the paper's signal (with and without the same filter); bare, 1988:01–2026:02. Each coefficient reports its Newey–West (12 lag) p and the wild bootstrap p on the respective state's episodes.

Conditioning	Carry Spot	Carry Total	Beta HML Spot	Beta HML Total
Naive count ≥ 2	−4.01	−1.89	−2.92	−2.99
$\mathrm{NW}_{12}[p]$	0.096	0.457	0.295	0.308
WCB[p]	0.085	0.556	0.400	0.328
Naive count, excl. downcurve	−4.35	−2.76	−3.02	−2.60
$\mathrm{NW}_{12}[p]$	0.158	0.390	0.402	0.486
WCB[p]	0.174	0.439	0.407	0.500
Naive count, 1/age-weighted	−14.30	−13.61	−24.10	−24.43
$\mathrm{NW}_{12}[p]$	0.008	0.012	0.000	0.000
WCB[p]	0.028	0.047	0.000	0.000
IYC signal	−11.63	−9.86	−12.82	−13.58
$\mathrm{NW}_{12}[p]$	0.003	0.013	0.001	0.000
WCB[p]	0.001	0.007	0.006	0.001
IYC signal, excl. downcurve	−17.35	−16.51	−17.93	−18.03
$\mathrm{NW}_{12}[p]$	0.002	0.003	0.000	0.000
WCB[p]	0.002	0.004	0.001	0.001

is a different matter. It was identified within the sample, so its two-era evidence carries less weight than the signal's.

B.4 The Persistence of the Signal

The state's 93 months arrive in fourteen runs of very unequal length, so the effective number of clusters is smaller than fourteen. Merging runs that begin within three or six months of one another leaves ten events and merging within twelve months leaves seven, with the carry cells significant at every level (wild $p \leq 0.005$) and the beta HML cells at $p \leq 0.001$. The count implied by the dispersion of run lengths, $(\sum_g n_g)^2 / \sum_g n_g^2$, is 7.0, because the longest run holds 28% of all signal months and the longest three hold 55%.

Collapsing each episode to its mean return removes the within-episode variation entirely and leaves fourteen observations to compare against the off months. A permutation test that reassigns the on-label to blocks of the same lengths gives a spot difference of −16.82%/yr

Table B.8: The carry portfolio against the rivals. Every cell estimates $r_{t+1} = \alpha + \gamma\,\text{State}_t + \delta' W_t + e_{t+1}$ on the leg named by the column. The specifications nest, so column (i) is the state alone and each later column adds the block named at its head while keeping those above it. Column (vi)′ repeats the full control set of column (vi) with the downcurve-filtered state, the non-downcurve signal of Table 3 (56 months), in place of the full signal, and its bootstrap resamples that state's own runs. Each variable reports its coefficient, the Newey–West (12 lag) p beneath it, and the wild cluster bootstrap p beneath that ($B = 9{,}999$, with the null imposed on that coefficient in turn). N falls in the last two columns of each leg, where the VIX enters and the estimation sample begins in 1990:12. Every control is predetermined except DOL and ΔFX vol. (t), which are contemporaneous with the return and make those columns spanning tests.

	Spot							Total						
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vi)′	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vi)′
State	−11.63	−10.60	−9.82	−10.08	−10.11	−7.39	−10.61	−9.86	−9.04	−8.27	−8.66	−8.70	−6.54	−9.70
NW$_{12}$[p]	0.003	0.000	0.000	0.001	0.001	0.038	0.007	0.013	0.002	0.004	0.003	0.003	0.061	0.014
WCB[p]	0.001	0.001	0.001	0.002	0.002	0.062	0.017	0.007	0.003	0.012	0.013	0.012	0.098	0.034
Financial Conditions														
NFCI		+0.56	+0.35	+7.90	+7.80	+2.47	+1.61		+1.20	+1.01	+9.03	+8.89	+2.08	+1.33
NW$_{12}$[p]		0.870	0.913	0.031	0.033	0.497	0.631		0.746	0.777	0.019	0.022	0.549	0.678
WCB[p]		0.875	0.902	0.166	0.166	0.572	0.689		0.719	0.745	0.094	0.097	0.621	0.725
ΔNFCI		−46.40	−43.28	−31.24	−31.87	−49.05	−48.30		−46.50	−43.37	−31.44	−32.35	−48.33	−47.62
NW$_{12}$[p]		0.000	0.000	0.001	0.002	0.000	0.000		0.000	0.000	0.001	0.001	0.000	0.000
WCB[p]		0.040	0.035	0.004	0.003	0.003	0.002		0.036	0.030	0.003	0.002	0.003	0.004
Dollar and Commodity Returns														
DOL			+2.84	+2.38	+2.39	+2.76	+2.77			+2.80	+2.34	+2.35	+2.74	+2.75
NW$_{12}$[p]			0.004	0.006	0.007	0.000	0.000			0.005	0.007	0.007	0.000	0.000
WCB[p]			0.219	0.318	0.318	0.145	0.140			0.218	0.322	0.322	0.141	0.140
Commodity return			−0.15	−0.12	−0.11	−0.04	−0.02			−0.15	−0.11	−0.10	−0.03	−0.02
NW$_{12}$[p]			0.219	0.286	0.317	0.747	0.830			0.232	0.298	0.341	0.786	0.860
WCB[p]			0.370	0.380	0.409	0.813	0.867			0.398	0.401	0.442	0.838	0.890
FX Volatility														
ΔFX vol. (t)				−78.29	−78.21	−89.63	−87.67				−79.17	−79.06	−88.54	−86.66
NW$_{12}$[p]				0.000	0.000	0.000	0.000				0.000	0.000	0.000	0.000
WCB[p]				0.001	0.001	0.000	0.000				0.000	0.000	0.000	0.000
FX volatility				−49.70	−49.75	−58.42	−56.70				−52.42	−52.50	−55.64	−54.11
NW$_{12}$[p]				0.000	0.000	0.000	0.000				0.000	0.000	0.000	0.000
WCB[p]				0.007	0.007	0.005	0.008				0.005	0.005	0.007	0.012
ΔFX vol. ($t-1$)				+1.39	+1.50	−2.48	−3.37				+2.91	+3.08	−2.90	−3.65
NW$_{12}$[p]				0.916	0.909	0.862	0.812				0.828	0.817	0.841	0.798
WCB[p]				0.918	0.913	0.862	0.810				0.832	0.826	0.842	0.804
Leading Indicator														
ΔCLI (1m)					−1.11	−4.22	−4.80					−1.59	−4.83	−5.37
NW$_{12}$[p]					0.831	0.392	0.300					0.755	0.314	0.236
WCB[p]					0.977	0.830	0.777					0.962	0.772	0.712
Equity Volatility														
VIX						+0.69	+0.72						+0.66	+0.68
NW$_{12}$[p]						0.007	0.004						0.012	0.008
WCB[p]						0.068	0.058						0.121	0.098
ΔVIX						+0.54	+0.56						+0.55	+0.57
NW$_{12}$[p]						0.088	0.077						0.084	0.075
WCB[p]						0.011	0.016						0.017	0.020
$\bar{R}^2$ (%)	2.7	6.8	12.1	20.7	20.5	27.1	27.8	1.9	5.9	11.0	20.0	19.8	26.1	26.7
N	458	458	458	458	458	423	423	458	458	458	458	458	423	423

Table B.9: The beta HML portfolio against the rivals. Specifications, controls, the (vi)′ column and reporting exactly as in Table B.8.

	Spot							Total						
	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vi)′	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vi)′
State	−12.82	−12.63	−11.90	−12.15	−12.20	−12.55	−12.82	−13.58	−13.67	−12.96	−13.29	−13.34	−13.00	−13.03
$NW_{12}[p]$	0.001	0.000	0.000	0.000	0.000	0.001	0.000	0.000	0.000	0.000	0.000	0.000	0.001	0.000
WCB$[p]$	0.006	0.004	0.004	0.002	0.002	0.003	0.002	0.001	0.000	0.001	0.001	0.001	0.002	0.002
Financial Conditions														
NFCI		+2.40	+2.50	+7.42	+7.26	−3.11	−4.80		+3.15	+3.27	+8.42	+8.29	−2.09	−3.85
$NW_{12}[p]$		0.566	0.534	0.082	0.096	0.394	0.238		0.474	0.441	0.066	0.077	0.557	0.346
WCB$[p]$		0.413	0.422	0.132	0.155	0.390	0.279		0.283	0.307	0.091	0.111	0.522	0.373
ΔNFCI		−37.20	−31.77	−27.27	−28.36	−38.07	−37.37		−35.92	−30.39	−25.87	−26.79	−36.57	−35.88
$NW_{12}[p]$		0.024	0.030	0.046	0.043	0.032	0.036		0.032	0.045	0.067	0.066	0.044	0.050
WCB$[p]$		0.215	0.192	0.161	0.123	0.172	0.197		0.235	0.217	0.192	0.155	0.198	0.233
Dollar and Commodity Returns														
DOL			+2.65	+2.38	+2.39	+2.44	+2.48			+2.57	+2.30	+2.31	+2.43	+2.48
$NW_{12}[p]$			0.014	0.020	0.020	0.014	0.014			0.018	0.026	0.026	0.014	0.014
WCB$[p]$			0.251	0.329	0.327	0.304	0.293			0.274	0.358	0.357	0.305	0.293
Commodity return			−0.01	+0.02	+0.03	+0.03	+0.05			+0.01	+0.03	+0.04	+0.05	+0.07
$NW_{12}[p]$			0.970	0.873	0.803	0.790	0.682			0.953	0.801	0.739	0.692	0.595
WCB$[p]$			0.974	0.894	0.845	0.787	0.683			0.963	0.840	0.795	0.690	0.607
FX Volatility														
ΔFX vol. (t)				−47.79	−47.66	−54.85	−53.93				−47.68	−47.56	−54.98	−54.14
$NW_{12}[p]$				0.002	0.002	0.001	0.001				0.002	0.002	0.001	0.001
WCB$[p]$				0.021	0.021	0.016	0.018				0.020	0.020	0.014	0.015
FX volatility				−32.05	−32.15	−27.90	−25.00				−33.33	−33.41	−27.83	−24.83
$NW_{12}[p]$				0.039	0.038	0.083	0.130				0.042	0.040	0.081	0.132
WCB$[p]$				0.084	0.082	0.153	0.243				0.071	0.068	0.123	0.223
ΔFX vol. $(t-1)$				+9.74	+9.94	+1.85	−0.11				+10.12	+10.29	+1.18	−0.87
$NW_{12}[p]$				0.452	0.440	0.902	0.994				0.448	0.438	0.938	0.955
WCB$[p]$				0.385	0.371	0.880	0.993				0.369	0.356	0.924	0.949
Leading Indicator														
ΔCLI (1m)					−1.91	−2.67	−3.15					−1.60	−2.83	−3.30
$NW_{12}[p]$					0.721	0.599	0.523					0.769	0.577	0.503
WCB$[p]$					0.931	0.867	0.832					0.943	0.848	0.815
Equity Volatility														
VIX						+0.77	+0.82						+0.73	+0.78
$NW_{12}[p]$						0.001	0.000						0.002	0.001
WCB$[p]$						0.026	0.015						0.030	0.019
ΔVIX						+0.32	+0.34						+0.34	+0.36
$NW_{12}[p]$						0.458	0.421						0.422	0.385
WCB$[p]$						0.463	0.425						0.429	0.389
$\bar{R}^2$ (%)	2.7	4.7	8.1	10.5	10.3	14.4	14.1	3.0	4.9	8.1	10.4	10.3	14.1	13.7
N	458	458	458	458	458	423	423	458	458	458	458	458	423	423

Table B.10: Own inversion test for each currency against the dollar. Each state column is the regression $\Delta s_{i,t+1} = \alpha + \gamma\,\text{State}_t + e_{i,t+1}$ on spot returns (%/yr), estimated over all months for that column's state; 1988:01–2026:02. ON counts the months the currency's own signal is on and No Cl. those of them with the IYC signal off, both dated as the signal is observed, and a dash marks a currency with fewer than four own-no-signal months. The state columns are the currency's own signal, that signal in the months the IYC signal is off, and the IYC signal, and the IYC signal, Ctrl. column repeats the signal regression holding the currency's own slope, its three-month change, its slope relative to the U.S., and the global average slope and its one-month change fixed. The relative-slope term is the raw counterpart of the relative Nelson–Siegel slope factor that Chen and Tsang (2013) use to predict exchange rates. Each coefficient reports its Newey–West p and the wild cluster bootstrap p ($B = 4{,}999$; one cluster per calendar signal block, shared across currencies in the pooled rows, and per own-signal spell otherwise). The pooled rows add currency fixed effects and a Driscoll–Kraay (12 lag) p, because the nine spot returns share the dollar factor within a month.

Currency	ON	No Cl.	Own	Own, No Signal	IYC signal	IYC signal, Ctrl.
AU	44	3	−5.14	—	−11.53	−13.08
$NW_{12}[p]$			0.267	—	0.056	0.084
WCB[p]			0.305	—	0.056	0.156
CA	41	10	−1.81	−1.98	−3.68	−4.02
$NW_{12}[p]$			0.621	0.294	0.260	0.353
WCB[p]			0.650	0.360	0.342	0.502
CH	12	0	−11.91	—	+1.45	+1.93
$NW_{12}[p]$			0.195	—	0.748	0.733
WCB[p]			0.416	—	0.738	0.736
EU	28	0	+4.65	—	+1.45	+1.31
$NW_{12}[p]$			0.562	—	0.775	0.836
WCB[p]			0.596	—	0.778	0.843
GB	68	21	−6.15	−8.13	−2.79	−3.75
$NW_{12}[p]$			0.219	0.109	0.554	0.523
WCB[p]			0.228	0.142	0.593	0.597
JP	22	2	−1.74	—	−0.39	+2.42
$NW_{12}[p]$			0.694	—	0.920	0.535
WCB[p]			0.797	—	0.926	0.568
NO	70	27	−4.58	+18.09	−4.77	−8.94
$NW_{12}[p]$			0.450	0.001	0.395	0.218
WCB[p]			0.549	0.013	0.445	0.293
NZ	54	12	−1.89	+4.00	−13.71	−12.59
$NW_{12}[p]$			0.738	0.689	0.006	0.052
WCB[p]			0.785	0.618	0.008	0.112
SE	24	0	−10.75	—	−3.15	−5.84
$NW_{12}[p]$			0.115	—	0.545	0.373
WCB[p]			0.175	—	0.577	0.469
Pooled (FE)			−4.00	+2.90	−4.13	−4.82
$NW_{12}[p]$			0.055	0.474	0.013	0.022
WCB[p]			0.159	0.503	0.335	0.411
DK[p]			0.136	0.479	0.284	0.323

Table B.11: Leave one episode out. Each row removes one signal episode's months from the sample entirely, not merely from the state, and re-estimates the headline regression $r_{t+1} = \alpha + \gamma \, \mathrm{IYC}_t + e_{t+1}$ (annualized, %/yr) on what is left; the first row is the full sample. Mo. is the length of the episode dropped, and the wild cluster bootstrap p on the remaining episodes is in brackets (B = 9,999); 1988:01–2026:02.

Episode Dropped	Mo.	Carry Spot	Carry Total	Beta HML Spot	Beta HML Total
None (baseline)	—	−11.63 [0.001]	−9.86 [0.007]	−12.82 [0.006]	−13.58 [0.001]
1988:08–1990:09	26	−13.39 [0.003]	−12.67 [0.008]	−17.55 [0.002]	−17.91 [0.001]
1990:12–1991:06	7	−12.30 [0.002]	−10.42 [0.008]	−11.13 [0.012]	−11.74 [0.002]
1991:08–1991:08	1	−11.67 [0.002]	−9.87 [0.008]	−12.79 [0.009]	−13.54 [0.001]
1995:01–1995:02	2	−11.18 [0.002]	−9.39 [0.010]	−12.69 [0.010]	−13.47 [0.002]
1997:08–1998:05	10	−11.90 [0.003]	−10.09 [0.012]	−11.92 [0.011]	−12.83 [0.001]
1998:08–1998:10	3	−10.37 [0.003]	−8.59 [0.016]	−12.96 [0.008]	−13.66 [0.001]
2000:04–2000:10	7	−11.33 [0.003]	−9.46 [0.014]	−12.60 [0.012]	−13.50 [0.001]
2000:12–2001:01	2	−12.33 [0.001]	−10.53 [0.005]	−13.41 [0.005]	−14.21 [0.000]
2005:03–2005:03	1	−11.83 [0.001]	−10.05 [0.008]	−12.66 [0.006]	−13.40 [0.001]
2006:01–2006:12	12	−13.06 [0.001]	−11.19 [0.004]	−13.48 [0.011]	−14.09 [0.001]
2007:08–2007:09	2	−11.56 [0.001]	−9.77 [0.009]	−13.20 [0.006]	−14.01 [0.001]
2008:06–2008:10	5	−8.99 [0.004]	−7.23 [0.023]	−10.98 [0.013]	−11.86 [0.002]
2020:02–2020:03	2	−10.65 [0.002]	−8.80 [0.013]	−11.82 [0.008]	−12.59 [0.000]
2022:08–2023:08	13	−13.06 [0.001]	−10.95 [0.008]	−13.42 [0.012]	−14.38 [0.001]

for the carry portfolio ($p = 0.006$) and −17.3 for the beta HML portfolio ($p = 0.003$). Equal weighting across episodes strengthens the estimate, moving the carry spot coefficient from −11.63 to −16.82, because the longest episodes are the mildest, which again points to the fresh inversions as the ones that carry the risk.

Timing. Every regressor is observed at the end of month t and predicts the return of month $t + 1$. A truncation audit rebuilds the signal from data ending at each of 435 month-ends and finds that no month's classification changes when later data arrives, and a second audit finds every conditioning series dated T identical to its full-sample value. Every conditioning variable is market-priced at month-end except the NFCI, and lagging the signal one *additional* month decays the coefficients as an ex-ante state variable should.

Table B.12: The headline in the two halves of the sample, split by the downcurve regime. State regressions $r_{t+1} = \alpha + \gamma\,\mathrm{State}_t + e_{t+1}$ (%/yr; no controls) estimated separately on 1988:01–2004:12 and 2005:01–2026:02, for the signal and for its non-downcurve and downcurve months, each against all other months of that half. Each coefficient reports its Newey–West (12 lag) p and the wild cluster bootstrap p on the state's episodes within the half ($B = 9{,}999$). The second half's downcurve cell has seven months (2023–24), so its rows are descriptive. The footer reports each half's months, signal months and their downcurve share, episodes, and how many of its three worst carry months fall inside the lagged state.

	1988:01–2004:12	2005:01–2026:02
Carry, spot	−12.21	−12.64
$\mathrm{NW}_{12}[p]$	0.002	0.143
WCB$[p]$	0.018	0.057
non-downcurve	−16.99	−17.94
$\mathrm{NW}_{12}[p]$	0.002	0.079
WCB$[p]$	0.029	0.027
downcurve	−3.77	+9.62
$\mathrm{NW}_{12}[p]$	0.237	0.005
WCB$[p]$	0.067	0.113
Carry, total	−10.94	−11.27
$\mathrm{NW}_{12}[p]$	0.008	0.185
WCB$[p]$	0.025	0.231
non-downcurve	−17.10	−16.47
$\mathrm{NW}_{12}[p]$	0.001	0.105
WCB$[p]$	0.031	0.110
downcurve	−1.61	+10.33
$\mathrm{NW}_{12}[p]$	0.621	0.003
WCB$[p]$	0.219	0.104
Beta HML, spot	−13.56	−14.57
$\mathrm{NW}_{12}[p]$	0.008	0.020
WCB$[p]$	0.018	0.009
non-downcurve	−16.05	−20.21
$\mathrm{NW}_{12}[p]$	0.000	0.007
WCB$[p]$	0.041	0.001
downcurve	−6.85	+9.38
$\mathrm{NW}_{12}[p]$	0.356	0.381
WCB$[p]$	0.680	0.296
Beta HML, total	−14.31	−15.02
$\mathrm{NW}_{12}[p]$	0.005	0.013
WCB$[p]$	0.011	0.007
non-downcurve	−15.40	−20.94
$\mathrm{NW}_{12}[p]$	0.000	0.003
WCB$[p]$	0.044	0.001
downcurve	−8.68	+10.08
$\mathrm{NW}_{12}[p]$	0.245	0.352
WCB$[p]$	0.370	0.272
N	204	254
Signal months (downcurve)	58 (30)	35 (7)
Episodes	8	6
Three worst carry months inside	1 of 3	3 of 3

B.5 Beta Definitions

Table B.13 rebuilds the beta HML portfolio under three alternatives to the baseline daily estimator of (4), a monthly 24-month rolling estimator against the same S&P 500 and daily and monthly estimators against the MSCI World index. Every definition preserves the loss. The daily MSCI cells are the weakest, which may reflect no more than a different exposure.

Table B.13: The beta HML portfolio under alternative currency-beta definitions. The baseline daily 90-trading-day estimator of (4) against the S&P 500, the same window against the MSCI World index, and monthly 24-month rolling estimators of the same regression on monthly appreciation against each index, each sort formed on the previous month's value. Cells are state regressions $r_{t+1} = \alpha + \gamma\,\text{State}_t + e_{t+1}$ of the resulting portfolio's spot and total returns on the signal and on the filtered signal, with the wild bootstrap p on signal episodes. All four definitions cover all nine currencies in every month of the sample.

Beta Definition	N	Signal, Spot	Signal, Total	Excl. DC, Spot	Excl. DC, Total
Daily 90-day rolling, S&P 500 (baseline)	458	−12.82	−13.58	−17.93	−18.03
$\text{NW}_{12}[p]$		0.001	0.000	0.000	0.000
WCB$[p]$		0.006	0.001	0.001	0.001
Monthly 24-month rolling, S&P 500	458	−9.79	−9.62	−17.62	−18.38
$\text{NW}_{12}[p]$		0.022	0.025	0.000	0.000
WCB$[p]$		0.031	0.022	0.002	0.000
Daily 90-day rolling, MSCI World	458	−7.73	−10.24	−10.19	−10.60
$\text{NW}_{12}[p]$		0.061	0.009	0.078	0.062
WCB$[p]$		0.103	0.004	0.086	0.062
Monthly 24-month rolling, MSCI World	458	−10.11	−10.45	−15.13	−15.69
$\text{NW}_{12}[p]$		0.007	0.004	0.003	0.001
WCB$[p]$		0.013	0.013	0.014	0.001

B.6 Portfolio Composition by State

Table B.14 reports how often each currency occupies each leg of the carry portfolio, by state. The membership is stable across states, with the same currencies anchoring both legs and only the marginal positions rotating. The one visible shift is in the third long-leg slot of the downcurve column, where the GBP displaces the NOK, which follows from the era composition of those 37 months. The top-two anchor of each leg is unchanged, so the downcurve gap of Section 3.5 is priced on the same core portfolio.

Table B.14: Portfolio composition by state, measured as the share of months (%) each currency occupies the long (top-3) and short (bottom-3) leg of the G10 carry portfolio, in the full sample, the IYC signal months, and the signal's non-downcurve (Excl. DC) and downcurve (DC) months; 1988:01–2026:02, all states dated as the signal is. The sorting variable is the previous month's money-market rate.

	Long Leg (% of Months)				Short Leg (% of Months)			
	Full	IYC signal	Excl. DC	DC	Full	IYC signal	Excl. DC	DC
NZ	85	92	95	89	2	0	0	0
AU	71	71	71	70	4	0	0	0
NO	62	30	48	3	3	4	7	0
GB	43	74	59	97	4	0	0	0
CA	21	19	21	16	23	12	18	3
SE	15	13	5	24	33	35	48	16
EU	4	0	0	0	51	48	27	81
JP	0	0	0	0	83	100	100	100
CH	0	0	0	0	96	100	100	100

Months per state: full 458,
IYC signal 93, excl. downcurve 56, downcurve 37.

B.7 Panel Interactions outside the G10

Table B.15 estimates (13) and (14) on the seven EM currencies of Section 5.1 and on all sixteen of Section 5.2. Neither panel reproduces the G10 gradient of Table 5, and the live-beta interaction is itself indistinguishable from zero.

B.8 The EM-Native Curve

Table B.16 rebuilds the conditioning using EM curves, from an inversion count over the seven EM curves and the shape of the EM average curve, with the G10+EM carry portfolio of Section 5.2 as the outcome (2003:07–2026, when all seven EM curves are available). The downcurve result carries over from the G10, though on only seven downcurve months. Outside that shape the EM signal is weaker than the G10 signal is on the same portfolio, at −6.43 in total returns and insignificant on either standard.

Table B.15: Panel interactions outside the G10. Each column estimates (13) or (14) with the annualized spot appreciation (%/yr) as the outcome and currency fixed effects throughout, on the seven EM currencies in Panel A and on all sixteen currencies of Section 5.2 in Panel B; the specification, the exposure measures, and the two p-values are those of Table 5. In columns (iii) and (v) the coefficient on the signal is the effect at zero exposure, which lies outside this sample, since every average IRD and every average beta here is positive. Sample 1996:02–2026:02.

	(i)	(ii)	(iii)	(iv)	(v)
Panel A: the seven EM currencies					
IYC_t	−13.87	−14.62	−18.94	−12.08	−16.09
$\mathrm{DK}_{12}[p]$	0.082	0.137	0.196	0.154	0.034
$\mathrm{WCB}[p]$	0.175	0.182	0.471	0.238	0.018
$\hat{\beta}_{i,t}$		−7.86			
$\mathrm{DK}_{12}[p]$		0.475			
$\mathrm{WCB}[p]$		0.532			
$\mathrm{IYC}_t \times \hat{\beta}_{i,t}$		+3.32			
$\mathrm{DK}_{12}[p]$		0.886			
$\mathrm{WCB}[p]$		0.830			
$\mathrm{IYC}_t \times \bar{\beta}_i$			+27.13		
$\mathrm{DK}_{12}[p]$			0.590		
$\mathrm{WCB}[p]$			0.887		
$\mathrm{IRD}_{i,t}$				−0.05	
$\mathrm{DK}_{12}[p]$				0.865	
$\mathrm{WCB}[p]$				0.849	
$\mathrm{IYC}_t \times \mathrm{IRD}_{i,t}$				−0.24	
$\mathrm{DK}_{12}[p]$				0.827	
$\mathrm{WCB}[p]$				0.696	
$\mathrm{IYC}_t \times \overline{\mathrm{IRD}}_i$					+0.37
$\mathrm{DK}_{12}[p]$					0.560
$\mathrm{WCB}[p]$					0.549
N	2527	2527	2527	2527	2527
Panel B: all sixteen currencies					
IYC_t	−9.55	−9.73	−5.55	−7.53	−6.93
$\mathrm{DK}_{12}[p]$	0.112	0.129	0.304	0.137	0.170
$\mathrm{WCB}[p]$	0.189	0.208	0.453	0.179	0.241
$\hat{\beta}_{i,t}$		+1.14			
$\mathrm{DK}_{12}[p]$		0.868			
$\mathrm{WCB}[p]$		0.851			
$\mathrm{IYC}_t \times \hat{\beta}_{i,t}$		+2.05			
$\mathrm{DK}_{12}[p]$		0.915			
$\mathrm{WCB}[p]$		0.920			
$\mathrm{IYC}_t \times \bar{\beta}_i$			−30.83		
$\mathrm{DK}_{12}[p]$			0.123		
$\mathrm{WCB}[p]$			0.187		
$\mathrm{IRD}_{i,t}$				+0.16	
$\mathrm{DK}_{12}[p]$				0.555	
$\mathrm{WCB}[p]$				0.665	
$\mathrm{IYC}_t \times \mathrm{IRD}_{i,t}$				−0.77	
$\mathrm{DK}_{12}[p]$				0.387	
$\mathrm{WCB}[p]$				0.502	
$\mathrm{IYC}_t \times \overline{\mathrm{IRD}}_i$					−1.01
$\mathrm{DK}_{12}[p]$					0.127
$\mathrm{WCB}[p]$					0.256
N	5776	5776	5776	5776	5776

The sixteen-currency beta HML portfolio, on-signal spot effect (%/yr)					
re-sorted monthly on the live beta	−24.29	$\mathrm{DK}_{12}[p]$ 0.023	$\mathrm{WCB}[p]$ 0.019		
frozen sample-average ranking	−12.34	$\mathrm{DK}_{12}[p]$ 0.050	$\mathrm{WCB}[p]$ 0.071		

Table B.16: EM-native conditioning, G10+EM portfolio returns (annualized, %/yr) by the EM own-signal state and the EM average curve's shape, 2003:07–2026. Gap regression $r_{t+1} = \alpha + \gamma_1 \, \mathrm{IYC}_t + \gamma_2 \, \mathrm{IYC}_t \times D_t + e_{t+1}$ on EM-signal episodes. The average EM curve in the downcurve months runs 8.38 at the money market, 7.94 at two years and 7.81 at ten (%).

Cell	N	Spot	Income	Total
EM signal on, EM curve downward-sloping	7	+17.41	+10.31	+27.71
EM signal on, other shapes	71	−7.18	+9.85	+2.67
EM signal off	194	+0.84	+8.26	+9.10
Gap regression, total returns				
EM signal in other shapes (γ_1)				−6.43
$\mathrm{NW}_{12}[p]$				0.209
WCB[p]				0.257
Signal × downward-sloping (γ_2)				+25.04
$\mathrm{NW}_{12}[p]$				0.004
WCB[p]				0.000

B.9 The Common Term Premium

The rival to the expectations interpretation is a common global term premium, proxied by the U.S. ACM slope premium of Adrian et al. (2013). Entering it as a control, in level and in changes, makes no difference to the headline. Table B.17 reports the more demanding version, which strips each country's slope of its fitted exposure and rebuilds the rule on the stripped slopes. Of the 1,122 inverted curve-months it leaves 937 inverted, ends 185, and creates 110 elsewhere. The rebuilt state is nonetheless wider, 124 months against 93, because the rule counts inversions that arrive together and not how many there are. On that broader state the beta HML portfolio loses much as on the baseline, −8.00%/yr in spot and −8.39 in total, both significant. The carry portfolio keeps about a third of its baseline loss and no longer clears significance. The regime split survives in shape, though the downcurve gap is no longer significant.

Table B.17: The signal rebuilt on term-premium-stripped slopes. Each country's 10Y−2Y slope is stripped of its fitted exposure to the U.S. ACM slope term premium, the IYC rule is rebuilt on the stripped slopes and dated as the signal is, and the rebuilt state is split by the downcurve regime. Cells are state regressions $r_{t+1} = \alpha + \gamma\,\text{State}_t + e_{t+1}$ of the carry and beta HML portfolios' spot and total returns (annualized, %/yr) against all other months; the downcurve − other row is the within-signal regression of the return on the downcurve dummy on the rebuilt state's months. Wild cluster bootstrap p on the state's runs in brackets ($B = 9{,}999$); 1988:01–2026:02.

State	Carry, Spot	Carry, Total	Beta HML, Spot	Beta HML, Total
Baseline signal	−11.63	−9.86	−12.82	−13.58
$\text{NW}_{12}[p]$	0.003	0.013	0.001	0.000
WCB$[p]$	0.001	0.007	0.006	0.001
Stripped signal	−4.10	−2.69	−8.00	−8.39
$\text{NW}_{12}[p]$	0.123	0.309	0.001	0.001
WCB$[p]$	0.127	0.370	0.034	0.020
non-downcurve months	−5.44	−4.71	−8.54	−8.52
$\text{NW}_{12}[p]$	0.111	0.164	0.002	0.002
WCB$[p]$	0.126	0.203	0.009	0.011
downcurve months	+0.27	+2.46	−3.65	−4.69
$\text{NW}_{12}[p]$	0.923	0.301	0.475	0.371
WCB$[p]$	0.902	0.101	0.790	0.826
downcurve − other	+4.66	+6.08	+3.59	+2.62
$\text{NW}_{12}[p]$	0.277	0.145	0.499	0.638
WCB$[p]$	0.221	0.117	0.552	0.694

B.10 Local Projections After a Signal Month

The state is persistent, so a projection on it averages the path after every signal month. Projecting on the onset of each run, or controlling for the state's own lags, avoids that, and Figure B.2 does both. The CLI's decline is concentrated in the first three to six months.

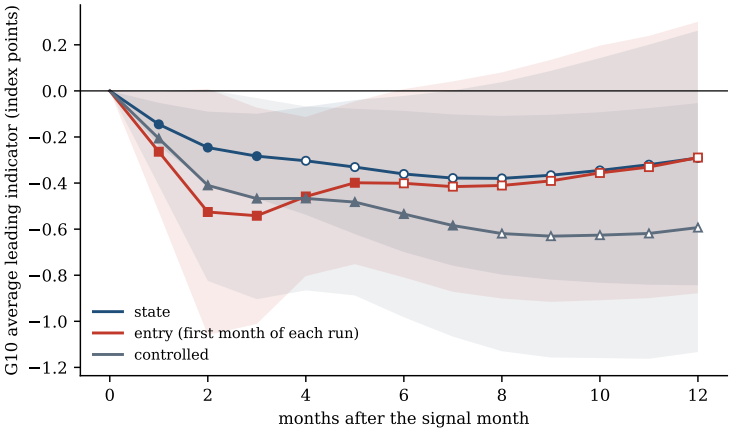


Figure B.2: The CLI after a signal month. Local projections of the nine-country average CLI (the panel of Table 1) on the state dated t, cumulative change from t to $t + h$. Circles mark the baseline state, squares the entry month of each run, and triangles the state with its own lag and three lags of the outcome; bands are 90% Newey–West intervals.