The Journal of High School Research

Macro Climate Regimes and U.S. Sector Rotation

Sasha Gooch*

Submitted: 15 September 2026 **Accepted:** 10 October 2026 **Publication date:** 10 October 2026

DOI: 10.70671/g015p139

Abstract: Practitioners have long believed that the economic "climate" determines which equity sectors lead the market, yet the link is usually drawn through a single growth-and-inflation grid, such as the well-known investment clock. This paper asks whether a richer, machine-learned description of the macroeconomic environment can do better. The answer, for the most part, is that it cannot, although the analysis identifies the one application where it genuinely earns its keep. We apply k-means clustering, a simple and transparent unsupervised-learning algorithm, to eleven monthly U.S. macroeconomic indicators spanning growth, inflation, monetary policy, the yield curve, credit spreads, financial conditions, and household sentiment, sorting every month from January 1986 through June 2026 into five recurring "economic climates." All inputs are public, drawn from the Federal Reserve Economic Data (FRED) service, the Kenneth French data library, and exchange-traded fund (ETF) prices, and every result is reproducible from the accompanying replication package. The five climates align closely with the credit cycle and with familiar episodes: the 2008 and 2020 crises fall in a credit-stress contraction, the 2021–2023 inflation surge in late-cycle overheating, and the early-1990s and 2001 downturns in recessionary easing. Each climate carries a distinct sector-rotation signature: technology leads in mid-cycle expansions, energy and defensive sectors lead in overheating, and rate-sensitive defensive sectors lead in easing. The dispersion between the best and worst sectors also widens dramatically in the adverse climates. Two disciplined tests, however, temper the machine-learning promise. First, in a complexity-matched comparison, the algorithmic climates do not describe forward sector returns any better than the investment clock: a five-state clock built from only growth and inflation matches the eleven-variable clustering in sample, and out of sample the simple clock is more accurate. Second, a strictly real-time strategy that tilts toward each climate's historically leading sectors fails to beat an equal-weight sector portfolio net of transaction costs. The durable application is risk control rather than sector selection: holding the equal-weight sector portfolio in the three benign climates and moving to cash in the two adverse ones roughly halves the maximum drawdown over 2001–2026, to 23 percent against roughly 50 percent for the buy-and-hold alternatives, while matching the market's return and improving the Sharpe ratio from 0.55 to 0.64. This capital-preservation benefit survives an untouched 2011–2026 holdout, a grid of alternative exposure settings, and, critically, the reconstruction of the entire macro dataset from real-time ALFRED vintages exactly as the data appeared on each historical decision date. It is competitive with, though not clearly superior to, the best simple alternatives, such as trend following, and, unlike those alternatives, it comes with an economic diagnosis that an investor can inspect, debate, and override. As of June 30, 2026, the framework classifies the U.S. economy as a mid-cycle expansion. The broader lesson is methodological: honest out-of-sample design, complexity-matched benchmarks, and real-time data vintages are what separate a genuine empirical finding from an artifact of hindsight.

Author keywords: Sector rotation; regime detection; business cycle; asset allocation; machine learning

Introduction

Most investors carry a mental model of sector rotation. When growth is accelerating, the conventional account favors cyclical sectors: industrials, materials, and consumer discretionary. When the economy overheats and inflation rises, investors rotate toward energy and other real-asset sectors. When the cycle turns down and the Federal Reserve begins to cut rates, investors shift to staples, utilities, and health care. The intuition is old, sensible, and widely taught. What it usually lacks is a disciplined, reproducible way to say *which climate we are in right now* and *how confident the historical sector pattern really is*.

This article supplies one. We take the position that the relevant "climate" is not a single variable but a configuration of the whole macro backdrop (growth momentum, inflation, the policy rate, the shape of the yield curve, credit spreads, and financial conditions), and we let an algorithm sort history into a small number of recurring climates rather than imposing a grid by hand. The approach is deliberately transparent: standard monthly data from FRED, a standard clustering method, and an out-of-sample test that any practitioner can reproduce.

*Corresponding Author: Sasha Gooch.
Email: sasha_gooch@asl.org
The American School in London, One Waverley Place, London NW8 0NP, United Kingdom

Three findings follow. First, the data organize cleanly into five climates that map onto the credit cycle and onto familiar episodes: the 2008 and 2020 crises land in a "credit-stress contraction," the 2021–2023 inflation surge lands in "late-cycle overheating," and the early-1990s and 2001 in "recessionary easing." Second, and contrary to our initial expectation, the algorithmic climates do not reliably out-describe the investment clock: their apparent in-sample edge disappears once the clock is granted the same number of states, and out of sample, the clock is at least as accurate. Third, while a real-time bet on each climate's leading sectors adds little net of cost, a climate-aware rule that simply *reduces risk* in the two adverse climates preserves capital materially, matching the market's return over our evaluation period with roughly half the drawdown and a higher Sharpe ratio.

The idea that the macro environment governs style and sector returns is not new, and we build on a substantial body of practitioner and academic work. Stovall[1] codified the conventional sector-rotation playbook around the business cycle. Conover et al.[2] writing in The Journal of Investing, showed that monetary conditions, meaning whether the Fed is in an expansive or restrictive stance, materially shape sector returns, with cyclical sectors rewarded when policy is easy. The investment clock of Greetham and Hartnett[3] distilled the macro backdrop into a four-quadrant growth-and-inflation grid that remains the reference point for tactical allocators. More recent evidence is decidedly mixed: Chava et al.[4] document business-cycle-dependent industry returns and a profitable out-of-sample rotation, and Chong and Phillips[5] report gains from tilting sector portfolios based on macroeconomic factors, whereas Molchanov and Stangl[6] find little systematic support for conventional sector rotation once transaction costs and cycle-timing errors are accounted for, and Sauer[7] reports, as we do, that regime nowcasting does not produce consistent sector-selection outperformance but can reduce drawdowns. On the methodological side, Hamilton[8] introduced the regime-switching framework that underpins much of the formal literature on identifying macro states, and Ang[9] and Ilmanen[10] survey how factor and asset-class premia vary across those states. Our contribution is to bring these threads together in a deliberately simple, reproducible way: rather than fixing the regimes by economic assumption or estimating a latent-state model, we let an unsupervised algorithm sort the realized macro record into climates, characterize them after the fact, and then test, out of sample and against the incumbent clock, whether the extra structure earns its keep.

The remainder of the article proceeds as follows. The next section describes the data and how we construct a point-in-time macro state. We then explain the sorting algorithm and characterize the five climates. We introduce the investment-clock benchmark, map each climate to its sector-rotation signature, and test head-to-head whether the algorithm beats the clock. We then build and evaluate the tradable overlay, report the classifier's reading as of mid-2026, discuss guidelines for application and directions for future work, and conclude. Methodological detail is deferred to the Appendix (Sections A1–A13) so that the main text stays readable.

Data and the Macro Climate State

We use three public sources. Macroeconomic and market series come from FRED. The long-history sector returns used to estimate rotation patterns are the value-weighted industry portfolios from the Kenneth French data library,[11] which extend back to the 1920s and give us statistical power that exchange-traded funds cannot. The tradable overlay is then validated on the nine original Select Sector SPDR exchange-traded funds (ETFs), which have traded since December 1998. All data are frozen as of June 30, 2026.

The macro state is a monthly vector of eleven indicators grouped into seven economically meaningful blocks: growth (industrial production and payroll growth, and the change in the unemployment rate), inflation (headline consumer prices and core personal-consumption expenditures), the policy stance (the twelve-month change in the federal funds rate), the yield curve (the ten-year-minus-two-year and ten-year-minus-three-month spreads), credit (the Baa-minus-ten-year corporate spread), financial conditions (the Chicago Fed's National Financial Conditions Index), and household sentiment (the twelve-month change in the University of Michigan consumer-sentiment index). Two practical choices keep the exercise honest. We lag the slow data releases by one month so that the state dated to the end of a month uses only information an investor could actually have known by then; market series such as yields and spreads are used contemporaneously because they are observable in real time. We then standardize each indicator so that variables measured in different units (percent, basis points, index points) can be compared on a common scale, and so that no single high-variance series dominates the clustering. For the descriptive taxonomy, we standardize over the full sample, which is appropriate for characterizing the historical record; for the live overlay, we standardize using only data available up to each point in time, which is essential for an honest out-of-sample test. The distinction matters: a framework that quietly uses full-sample knowledge to classify the past will flatter itself, and much of the disappointment with regime strategies traces to exactly this kind of look-ahead. The common, full-feature window runs from January 1986, the start of the credit-spread series, through June 2026, covering the 1987 crash, the early-1990s recession, the dot-com boom and bust, the global financial crisis, the COVID-19 shock, and the 2022 inflation episode. It contains 485 monthly observations rather than the full 486: the November 2025 state cannot be formed because the October 2025 consumer-price and household-survey releases were canceled during that autumn's federal government shutdown, a real-world reminder that the framework only sees what the statistical agencies publish.

Sorting History into Climates

The "sorting algorithm" is k-means clustering applied to the standardized macro state. K-means is about as transparent as a machine-learning method gets: it finds a chosen number of group centers and assigns each month to its

nearest center, so each climate is simply a region of macro space, summarized by its average conditions. We choose five climates. The silhouette score, a standard measure of how cleanly the groups separate, is essentially flat between four and seven climates (around 0.24–0.26), so the choice is governed by economic interpretability rather than a sharp statistical optimum; five is the smallest number that yields distinct, nameable regimes without redundancy. We confirm in Appendix A3 that hierarchical (Ward) clustering and a Gaussian mixture produce very similar partitions.

We name each climate based on its average conditions using a one-to-one assignment to five economic archetypes (the procedure is in Appendix A2). Fig. 1 plots the resulting climate of every month from 1986 to 2026 as a colored band, with the Baa credit spread overlaid and NBER recessions hatched.

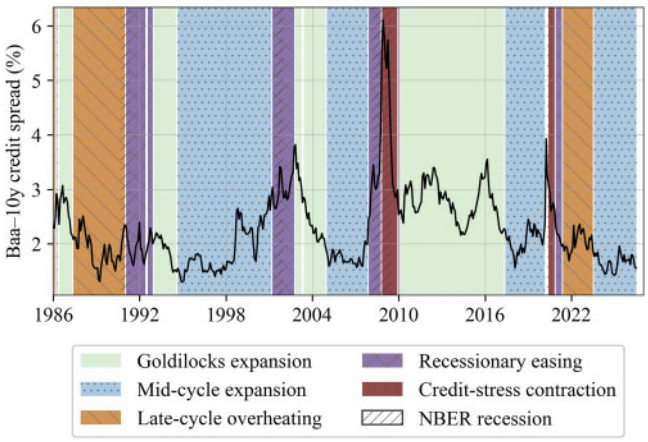


Figure 1. U.S. macro climate timeline and credit stress, 1986–2026. Source: FRED (Baa–10y corporate spread), NBER recession dates; author calculations. Climate bands are the full-sample taxonomy labels.

The figure is the first validation of the approach: the clustering, which never sees a recession date or any subsequent sector return, places the 2008 and 2020 crises squarely in the red "credit-stress contraction" climate, exactly where the spread spikes, brackets them with the purple "recessionary easing" climate, paints 2021–2023 in the orange for "late-cycle overheating" climate as inflation surges, and fills the long, quiet expansions of the mid-1990s and mid-2010s with the green and blue benign-growth climates.

Table 1 reports the average macro reading in each climate, which is how the names are earned. To read it, scan down a column: the credit-stress contraction climate shows industrial production falling at an eleven-percent annual rate, unemployment rising by more than four points, and a credit spread near four percent. That is a crisis, and indeed fifty-five percent of those months overlap an official recession. Late-cycle overheating pairs strong growth with five-percent inflation and a flat curve. Recessionary easing shows falling growth and the Fed cutting aggressively. The two expansion climates differ mainly in inflation and the curve: goldilocks expansion has low inflation and a steep curve, while mid-cycle expansion has moderate inflation and a flat curve.

It is worth walking through the five climates one by one, because their economic personalities are what give the rotation results their meaning. *Goldilocks expansion* (146 of the 485 months) is the most favorable state: growth running a little above two percent, inflation below two percent, and a comfortably positive yield curve, as in the disinflationary booms of the mid-to-late 1990s and stretches of the 2010s. *Mid-cycle expansion* (185 months) is the most common state: still-solid growth, but inflation nearer three percent and a flatter curve as the Fed has been tightening; this is the maturing phase of an upswing. *Late-cycle overheating* (72 months) is where the economy runs hot: growth is strong, but inflation is above five percent and the curve is flat to inverted, as in the classic prelude to trouble seen in the late 1980s and again in 2022–2023. *Recessionary easing* (62 months) is the downturn proper: growth is negative, unemployment is rising by more than a point, and the Fed is cutting rates by more than two points over the year. Finally, *credit-stress contraction* (20 months) is the tail: the acute phase of the 2008–2009 and 2020 crises, when output collapses and credit spreads blow out. In the full-sample taxonomy, these twenty months form just two episodes, October 2008 to November 2009 and May to October 2020 (the real-time classifier of the later sections also placed May 2001 to March 2002 in this climate). These months are few, but they account for the largest losses and sector dislocations, and their small number is the central caveat to everything the paper says about this climate.

The Investment-Clock Benchmark

To judge whether the algorithm adds anything, we need a benchmark that practitioners already use. The natural choice is the investment clock popularized by Merrill Lynch,[3] which divides the world into four quadrants by whether growth is above or below trend and whether inflation is rising or falling: recovery, overheat, stagflation, and reflation. We implement the clock faithfully on the same data, using a growth composite relative to its trailing median and the six-month direction of an inflation composite, and assign each month to a quadrant. Fig. 2 shows the months in growth-inflation space, colored by quadrant.

The clock is a sensible, interpretable scheme, and we treat it not as a straw man but as the incumbent the algorithm has to beat.

How Climates Drive Sector Rotation

We measure rotation as relative performance. Each month we compute every industry's return in excess of the equal-weighted average across industries, which strips out the overall market move and isolates leadership. We then ask: conditional on the climate prevailing at the end of a month, what is the average excess return of each sector over the following quarter? We use a quarter, rather than a single month, because rotation plays out over months, not weeks. Fig. 3 presents the resulting map: sectors across the top,

Table 1. Climate characterization: average macro conditions, 1986–2026

	Goldilocks expansion	Mid-cycle expansion	Late-cycle overheating	Recessionary easing	Credit-stress contraction
Months	146	185	72	62	20
Recession overlap	0%	0%	6%	34%	55%
Ind. production, YoY %	2.0	3.0	3.2	−1.1	−11.3
Payrolls, YoY %	1.3	1.8	3.3	−1.0	−5.3
Unempl. rate, 12-m chg (pp)	−0.5	−0.3	−1.1	1.2	4.4
CPI, YoY %	1.9	2.7	5.3	3.1	0.3
Core PCE, YoY %	1.9	2.1	4.3	2.4	1.1
10y–2y curve (pp)	1.8	0.3	0.4	1.6	1.7
10y–3m curve (pp)	2.4	0.5	1.0	2.4	2.3
Fed funds, 12-m chg (pp)	−0.2	0.6	0.9	−2.2	−2.5
Baa–10y spread (pp)	2.6	1.9	2.0	2.6	3.9
Financial conditions (NFCI)	−0.6	−0.5	0.0	−0.3	0.8
Consumer sentiment, 12-m chg	3.9	0.9	−4.7	−10.6	−9.7

Note. Raw-unit centroid averages over the 1986–2026 window; recession overlap is the share of each climate's months inside an NBER recession. Source: FRED, NBER; author calculations.

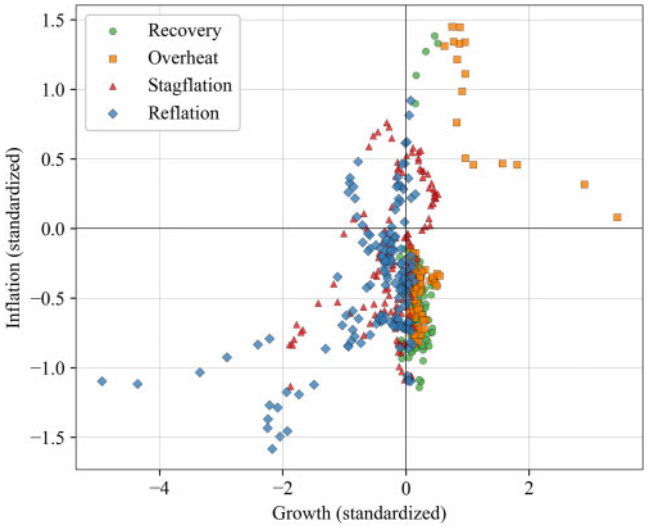


Figure 2. The investment-clock benchmark. Source: FRED; author calculations. The growth and inflation composites are standardized over the full sample; quadrants follow Greetham and Hartnett.[3]

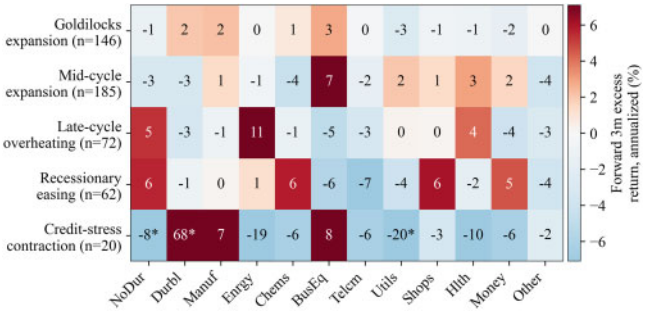


Figure 3. Sector rotation map: forward excess return by climate. Source: Kenneth French data library (12-industry portfolios); author calculations. An asterisk marks the cells significant at 5% after a Benjamini–Hochberg correction (Appendix A6); Table 2 and Appendix A13 give bootstrap intervals. Climate labels carry the number of months behind each row.

climates down the side, with leadership in red and lagging in blue.

The pattern is textbook. In late-cycle overheating, energy leads by roughly eleven points annualized, with consumer staples and health care also positive, while technology and financials lag, which is exactly the real-asset, defensive tilt the overheating narrative predicts. In recessionary easing, the leaders are defensive sectors and rate-sensitive sectors: staples, discretionary staples (shops), chemicals, and, as the curve re-steepens, financials; communications and technology lag. In mid-cycle expansion, technology leads. In the credit-stress contraction, the cross-section is violent and dominated by a small number of crisis months: deep-cyclical durables rebound hard off the bottom, while energy and utilities are punished. We flag with an asterisk the handful of cells that survive a multiple-testing correction; most do not, which is the honest truth about monthly sector returns: they are noisy, and the economic content lies in the broad pattern rather than in any single cell.

How much weight can any one cell bear? Table 2 answers for the climate where the question matters most. For the credit-stress contraction, it reports each sector's average forward excess return, together with a 90-percent block-bootstrap interval and the number of months behind the estimate. The intervals are wide: their median width is fifteen

Table 2. Uncertainty in the crisis climate: forward sector excess returns in the credit-stress contraction with 90% block-bootstrap intervals (20 months, two episodes)

Sector (Fama-French industry)	Mean (%)	90% interval (%)	Excludes zero
Consumer durables	+67.5	[+44.2, +91.8]	yes
Business equipment (technology)	+8.4	[+1.2, +16.0]	yes
Manufacturing	+7.1	[+2.0, +12.4]	yes
Other	−2.4	[−9.6, +4.1]	no
Shops (retail, wholesale)	−3.5	[−11.2, +4.4]	no
Finance	−6.3	[−19.3, +7.1]	no
Telecommunications	−6.3	[−10.7, −2.0]	yes
Chemicals	−6.3	[−14.4, +1.4]	no
Consumer nondurables	−8.5	[−12.2, −4.6]	yes
Health care	−10.3	[−18.2, −2.1]	yes
Energy	−19.2	[−36.3, −1.0]	yes
Utilities	−20.3	[−27.2, −13.9]	yes

Note. Annualized forward 3-month excess return over the equal-weight sector average, conditional on the credit-stress contraction climate (20 months in the 1986–2026 taxonomy, forming two episodes: October 2008 to November 2009 and May to October 2020). Intervals are the 5th and 95th percentiles of a stationary block bootstrap (mean block length of six months, 5,000 resamples) that resamples within these 20 months; 8 of 12 intervals exclude zero, but the intervals capture noise within the two episodes, not how a different credit crisis might behave. Intervals for all sixty cells in Fig. 3 are in Appendix A13. Source: Kenneth French data library; author calculations.

points, against six or seven points in the two expansion climates, and even the headline durable figure of 68 points spans roughly 44 to 92. Eight of the twelve intervals exclude zero, but this reflects how extreme and internally consistent the two episodes were, not a large sample: the bootstrap resamples within the twenty months drawn from two crises, so it measures month-to-month noise within those episodes and cannot say how different the next credit crisis might look. Intervals for all sixty cells in the map are in Appendix A13. We therefore read the credit-stress row as showing a direction, deep cyclicals rebounding off the bottom while energy and utilities lag, rather than as a set of magnitudes, and nothing in the risk-dial results that follow depends on any sector number from this climate.

Two features of the map deserve emphasis because they drive everything that follows. First, dispersion is not constant across climates: in the benign goldilocks state, sectors barely separate (about five annualized points between the best and worst sectors), whereas the spread widens to thirteen points or more in the easing and overheating climates and approaches ninety points in the twenty credit-stress months. Rotation, in other words, is most worth acting on exactly when the climate turns hostile, a point the static investment clock, which treats all of its quadrants symmetrically, tends to obscure. Second, the leadership rotates in an economically sensible sequence as one moves down the map from expansion to overheating to easing to crisis: from growth-sensitive technology, through real-asset and defensive energy and staples, to rate-sensitive and countercyclical defensive sectors, and finally to the violent, mean-reverting moves of a crisis. The map is not a collection of statistical curiosities; it is the cross-sectional shadow of the credit cycle.

A Fair Test Against the Clock

The map is intuitive, but intuition is cheap, and a fair comparison against the incumbent clock is harder than it first looks. The natural statistic is the fraction of the variance in forward three-month sector excess returns explained by knowing the regime, pooled across sectors and months and computed identically for the climates and for the clock. On that raw measure, the five climates explain 5.2 percent of sector dispersion against the 3.7 percent for the four-quadrant clock, an apparently decisive forty-percent edge. But two corrections are needed before that number can be trusted, and together, as Fig. 4 shows, they dissolve most of the apparent edge.

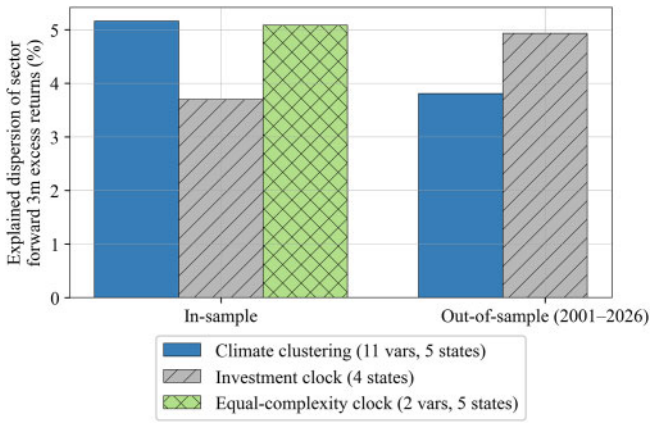


Figure 4. A fair comparison: climate taxonomy versus the investment clock. Source: FRED, Kenneth French data library; author calculations. Out-of-sample bars use strictly past-only labels, 2001–2026 (Appendix A4).

The first correction is for complexity. Five regimes fit the cross-section more flexibly than four simply because there are more of them, and the pooled statistic rewards extra categories mechanically. So we build an equal-complexity benchmark: a five-state clock formed by clustering on only the two clock inputs, growth and inflation. It explains 5.1 percent, essentially identical to the eleven-variable climate map. Almost all of the apparent advantage, in other words, comes from using five states rather than four, not from the nine extra macro variables the climates add. A block bootstrap confirms the point: the gap between the climates and the standard four-quadrant clock has a 95 percent confidence interval that comfortably straddles zero, so it is not statistically distinguishable from no difference.

The second correction is for out-of-sample honesty. Using the strictly past-only labels described in the next section, the climates explain 3.8 percent of forward sector dispersion over 2001–2026, while the clock explains 4.9 percent. Out of sample, the simple clock is the better description, not the climates. The taxonomy itself is stable: re-estimating it with hierarchical and Gaussian mixture clustering recovers essentially the same five climates (adjusted Rand indices 0.80 and 0.74), and the same five-climate structure recurs over a longer, credit-free sample extending back to 1971. What is not supported is the stronger claim we initially set out to make. As a forecaster of sector rotation, the algorithmic climate map is no better than the growth-and-inflation grid practitioners already use.

That is a useful thing to have established, even though it is a negative result, and it reshapes the rest of the article. If the climates do not beat the clock at sector selection, two honest possibilities remain: the taxonomy may earn its keep as an interpretable read on where we are in the cycle, and its one durable economic application may be risk control rather than rotation. We take up both next.

From Selection to Risk Control

Can an investor monetize the climates at all? We test the obvious idea first and report it plainly. At each month-end, we read the prevailing climate using only past data, then hold an equal weight of that climate's historically leading sectors, rebalancing the monthly net of ten basis points in costs. The classifier is strictly past-only: each year, it re-clusters the expanding history and re-derives the climate names from that window's own data, never referring to full-sample information, and the sector ranking uses only return windows that have closed by the decision date. The evaluation window runs from 2001, after a decade of burn-in, through June 2026.

The candid finding comes first: a pure bet on each climate's leading sectors does *not* beat an equal-weight portfolio of the same sectors. The rotation overlay earns 7.6 percent annually against the 9.4 percent for the equal-weight benchmark, with a lower Sharpe ratio and no better drawdown; the ETF implementation tells the same story. Out of sample, the relative-sector signal is too weak to overcome trading costs and estimation noise, consistent with the fair-test result above, and with a literature littered with rotation strategies that look wonderful in sample but evaporate afterward.[6,7] The dispersion in Fig. 3 is real, but much of it sits in a handful of crisis months that are hard to harvest through a long-only sector tilt.

The climates make a second prediction, though, that does not require picking the right sectors: two of them, credit-stress contraction and recessionary easing, are simply bad states of the world, and they are identifiable in real time. So instead of rotating, we de-risk. We hold the equal-weight sector book in the benign climates and step entirely to cash in the two adverse ones, with the climate label fixed on economic grounds, not on returns. The exposure rule is economically motivated rather than optimized: the two adverse climates are simply the recession- and credit-stress-shaped states, the all-in-or-all-out exposure is the plainest choice available rather than a tuned parameter, and the benefit is stable across reasonable exposure settings (Appendix A7).

Because a risk-off rule is only as honest as the choice of its risk-off states, that choice deserves a direct account. The decision to test de-risking at all came after the sector-selection result, as the order of this section makes plain. The choice of which climates count as adverse, however, did not involve returns. The two states are defined by their archetype names, and the names come from the fixed rule in Appendix A2, which scores each cluster's macro centroid on growth, inflation, financial stress, policy easing, and curve steepness, as well as its recession overlap. No sector return, market return, or strategy statistic enters that rule; the walk-forward classifier re-derives the names each year from past-only data; and the recession-shaped and credit-stress-shaped states are adverse by construction. Would a search over the alternatives have produced the same pair? No. Appendix A7 evaluates every possible cash set of climates, all thirty of them, under the identical convention. The pre-specified pair is not the return-maximizing choice: exiting the credit-stress climate alone earns a higher return and Sharpe ratio with the same drawdown, so a rule tuned to the results would have dropped recessionary easing, and the subsets that score best on the Sharpe ratio hold cash through long benign expansions and make no economic sense. What the sweep does confirm is the mechanism. Every rule that exits the credit-stress climate reaches the 23 percent drawdown, and no rule that stays invested through it does better than 46 percent. Fig. 5 shows the result for the pre-specified rule.

The climate-dialed book earns 9.4 percent annually, matching the equal-weight benchmark and ahead of the S&P 500′s 9.0 percent, but with much lower volatility (12.5 versus 15.2 percent) and a maximum drawdown of 23 percent against roughly 50 percent for both the equal-weight sectors and the index. The Sharpe ratio rises from 0.55 to 0.64 and the Sortino from 0.75 to 0.86. The chart makes the mechanism plain: the strategy simply does not participate in the 2001–2003 and 2008–2009 collapses, then compounds from a higher base; the shaded band traces the 5–95 percent block-bootstrap range of cumulative paths. Turnover is negligible (the dial trades only when the economy crosses a climate boundary, a few dozen times in twenty-five years), so the result does not hinge on the cost assumption.

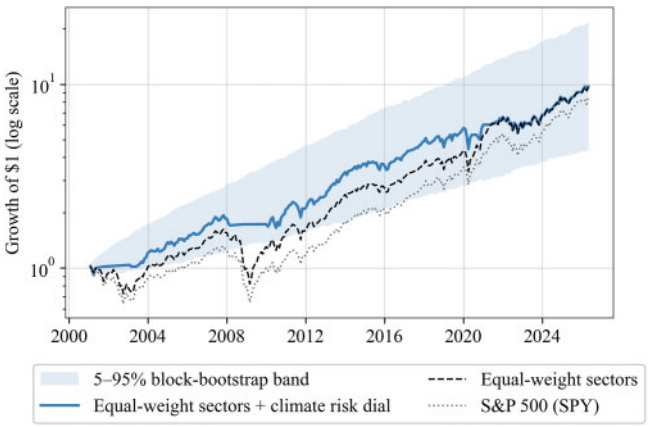


Figure 5. Capital preservation from climate de-risking, 2001–2026. Source: Kenneth French data library, Yahoo Finance (SPY), FRED (federal funds rate); author calculations. All strategies are long-only, unlevered, net of 10 bps per unit of turnover.

We put the de-risking result through the tests a skeptical reader would demand; Table 3 and Appendix A7 collect them. It survives an untouched holdout: splitting the sample, the dial improves risk-adjusted returns in both 2001–2010 and the held-out 2011–2026 period. It is robust across a pre-specified grid of exposure choices rather than depending on the exact weights, and the benefit is driven mainly by exiting the credit-stress climate. Nor does the dial smuggle recession dates in through the climate names: the naming rule scores clusters partly on their past NBER overlap, but re-running the entire walk-forward classifier with that term removed leaves the result unchanged, with an identical 23 percent maximum drawdown (Appendix A5). And the dial is not obviously beaten by simpler rules: a credit-spread threshold and an unemployment-trend rule both do worse. In the interest of candor, however, a plain ten-month trend-following rule achieves a comparable, even slightly smaller, drawdown with a similar Sharpe ratio. The honest conclusion is therefore measured: climate-based de-risking is a genuine, economically interpretable capital-preservation tool that is competitive with the best simple alternatives, not a singular source of edge. Why, then, prefer the climate model to a trend rule? Because everything around the signal is different. The climate dial comes with an economic diagnosis that an investment committee can debate, document, and override; it maps macro scenarios directly to positioning; and it is built from macro quantities rather than price, so it de-risks for a different reason than a trend rule does, which makes the two natural complements in one process rather than competitors for the same slot.

One further test matters most: real-time data. Our macro state, like most, is built from data as it looks today, after years of revisions, whereas an investor saw earlier prints. We rebuild it in two ways using ALFRED vintages. The faithful version uses each series exactly as it stood on each decision date; here, the result is reassuring: the maximum drawdown is unchanged at 23 percent, and the Sharpe ratio barely moves, to 0.63. A stricter version uses only the first published print of each series, discarding even revisions an investor would have seen; that is more punishing, trimming the Sharpe ratio to 0.50 and the return by about two points. Across all three vintages, the drawdown protection is intact; the risk-adjusted-return edge holds under realistic real-time data and weakens only under the most conservative first-release assumption. Table 3 reports the as-of-date figures, and Appendix A9 reports all three. A K-sweep and a leave-2020-out re-clustering (Appendix A10) leave the picture intact: the drawdown benefit is stable for five or more climates and does not depend on the COVID shock.

Discussions and Guidelines for Applications

The classifier's most recent reading illustrates how the framework is meant to be used. As of June 30, 2026, the live, past-only classifier assigns the U.S. economy to mid-cycle expansion: solid growth, moderate inflation, a flat-to-modest curve, easy financial conditions, and contained credit spreads. The timing of that reading is exact. The

Table 3. Performance and risk, 2001–June 2026 (net of costs; long-only, unlevered)

Strategy	CAGR	Vol	Sharpe	Sortino	Max DD
Equal-weight sectors + climate risk dial	9.4%	12.5%	0.64	0.86	−23.2%
Equal-weight sectors (buy and hold)	9.4%	15.2%	0.55	0.75	−49.5%
S&P 500 (SPY)	9.0%	15.1%	0.52	0.72	−50.8%
Sector-rotation overlay (selection only)	7.6%	16.1%	0.42	0.57	−49.9%
. . . same, on real-time (as-of-date) data	9.0%	11.9%	0.63	n/a	−23.2%
Simple trend de-risking	8.3%	10.5%	0.63	n/a	−17.9%
Credit-spread de-risking	3.9%	7.6%	0.30	n/a	−29.8%
Unemployment de-risking	6.9%	10.9%	0.49	n/a	−23.2%

Note. All strategies are long-only and unlevered, net of 10 bps per unit of turnover. The climate risk dial moves the equal-weight sector book to cash in the two adverse climates (credit-stress contraction and recessionary easing), using strictly past-only climate labels; cash earns the effective federal funds rate. Sharpe and Sortino ratios use the federal funds rate as the risk-free rate; maximum drawdown is computed from monthly returns. Sources: FRED, Kenneth French data library; author calculations.

state dated June 30, 2026 uses the May 2026 values of the seven slow releases (industrial production, payrolls, unemployment, CPI, core PCE, the twelve-month change in the federal funds rate, and consumer sentiment), all published by mid-June, and the June 30 values of the four market-observable series (the two curve spreads, the Baa–10y credit spread, and the NFCI). Each indicator is standardized on an expanding window ending in June 2026, and the classifier's centroids and names were last re-estimated at its January 2026 annual refit on data through December 2025, so nothing in the model draws on information unavailable on the reading date; the sensitivity of the results to data vintage is quantified in Table 3 and Appendix A9. Table 4 shows the historical sector tilt associated with this climate. As the earlier sections established, that tilt is a description of past leadership rather than a dependable forecast; the actionable element is the risk dial, which is invested because mid-cycle expansion is a benign climate. The reading is a starting point for judgment rather than a black box: a reader who disagrees with the classification can see exactly which macro readings produced it, and because mid-cycle expansion borders both late-cycle overheating and recessionary easing, the indicators that would move the classification (inflation and the curve on one side, growth and the policy rate on the other) are identifiable in advance.

Table 4. Current climate reading, June 2026: mid-cycle expansion

Sector (Fama-French industry)	Historical forward 3-m excess return, annualized (%)
BusEq	+7.5
Hlth	+3.0
Utils	+2.0
Money	+1.6
Shops	+1.4
Manuf	+1.4
Enrgy	−0.6
Telcm	−1.6
NoDur	−3.2
Durbl	−3.2
Other	−3.9
Chems	−4.3

Note. Historical average forward 3-month excess return of each sector in the climate read by the strictly past-only classifier as of June 30, 2026; a description of past leadership, not a forecast. Source: Kenneth French data library; author calculations.

Three caveats temper the results and shape how they should be applied. First, the crisis climate rests on only twenty months in two episodes; its sector statistics are illustrative rather than precise (Table 2), and the *direction* of the risk dial in that state is the finding, not any specific sector number. Second, classification is necessarily a coincident-to-slightly-lagging exercise: because slow macro data are released with a lag, the framework recognizes a climate shortly after it begins rather than in anticipation. In practice, this means the risk dial protects against the bulk of a drawdown but not its first few weeks; the 2008 and 2020 experiences in Fig. 5 show exactly this shape. Relatedly, the headline backtest uses revised data; on genuinely real-time vintages (Table 3), the drawdown protection holds but the return advantage does not, so we lean on the former. Third, the strategy deliberately uses the simplest possible rules (equal weights and a binary in-or-out risk dial, rebalanced monthly) so that the economic content is transparent; an implementation would naturally smooth the transitions, blend the climate signal with valuation, trend, or momentum, and size positions to a volatility target. We see the framework less as a finished strategy than as a disciplined lens: a reproducible way to state which climate prevails and when the macro backdrop argues for caution.

Recommendations for Future Work

Several extensions would strengthen or broaden the framework. The most natural is to soften the machinery at its edges: replacing the hard climate assignments with class probabilities would let exposure scale down gradually as the economy approaches an adverse climate rather than switching all at once, and would likely reduce the whipsaw around climate boundaries. A second direction is combination: because the climate dial de-risks on macro information and a trend rule de-risks on price, a study of the joint rule (de-risk when either signal turns cautious, or size by both) could quantify how complementary the two really are, and our preliminary volatility-targeted variant (Appendix A12) suggests the blend is promising. Third, the framework should travel: the same eleven-indicator state can be built for other developed economies from OECD and national-bank data, and finding the same five climates, or instructively different ones, outside the United States would be a genuine test of the taxonomy rather than a recycling of one country's history. Fourth, the sector side could be enriched, from the twelve Fama-French industries to the eleven GICS sectors or finer industry groups, and from long-only tilts to long-short implementations that might harvest crisis-month dispersion more effectively than a long-only book can. Finally, the honest-comparison methodology itself (complexity-matched benchmarks, strictly past-only labels, and real-time data vintages) deserves wider use: many published regime results would look different under these controls, and replicating them under this discipline would be a service to the literature.

Conclusions

The intuition behind sector rotation is sound, but the way it is usually framed, as a single growth-and-inflation grid, invites more confidence than the evidence supports. Letting a simple, transparent algorithm sort the full macro backdrop into five recurring climates produces a map that aligns closely with the credit cycle and gives an investor an

interpretable read on where the economy stands. What that map does not do, we find, is forecast sector rotation better than the investment clock: once the clock is granted equal complexity, the advantage disappears, and out of sample the clock is the better description. Nor can the climates be monetized through sector selection, which proves fragile net of cost. Their one durable economic application is risk: recognizing in real time the two adverse climates, credit-stress contraction and recessionary easing, and stepping out of equities roughly halved the maximum drawdown over 2001–2026, a downside benefit that is robust to an out-of-sample holdout and to real-time data vintages, and it is competitive with the best simple de-risking rules. The practical lesson is therefore simple: use the climates as a risk dial, not as a sector-selection oracle. The whole exercise rests on public data and standard methods, so any reader can reproduce it, argue with the climate definitions, and fold the signal into their own process.

Acknowledgements

The author developed the analysis code with assistance from an artificial-intelligence assistant (Claude, by Anthropic), used under the author's direction; the research questions, methodological choices, manuscript, verification of results, and final content are the author's responsibility. No external funding was received, and the author declares no conflicts of interest.

Supplemental Materials

Three items accompany the article. First, the Appendix at the end of this paper (Sections A1–A13) collects the methodological detail deferred from the main text: the exact indicator definitions and release-lag conventions; the climate-naming procedure; the alternative-clustering and long-sample (1971–2026) robustness checks; the exposure-grid, simpler-rule, and exhaustive cash-set analyses for the risk dial; the leave-2020-out and K-sweep analyses; the volatility-targeted variant; the full real-time (ALFRED) vintage results; and bootstrap intervals for every cell of the rotation map. Second, a data workbook provides the underlying numbers behind every figure and table, one worksheet per display item, with a native spreadsheet chart duplicating each figure. Third, the complete replication code is deposited in a permanent, versioned repository[12]: the archive contains every ingestion, clustering, testing, and figure-generation script, the run-everything driver, and a dated manifest of the frozen June 30, 2026 data snapshot, including a SHA-256 hash of every raw input file, and it is citable by digital object identifier at https://doi.org/10.5281/zenodo.23023444. All raw inputs are public[13–16] and the archive re-downloads them from their source via the included ingestion script; the raw files themselves cannot be redistributed under their source terms, so the restricted inputs are documented by manifest hash and the exact frozen copies are available from the author on request.

References

[1] Stovall S. *Standard & Poor's Sector Investing: How to Buy the Right Stock in the Right Industry at the Right Time*. New York: McGraw-Hill; 1996.

[2] Conover CM, Jensen GR, Johnson RR, Mercer JM. Sector rotation and monetary conditions. *J Invest*. 2008;17(1):34–46. doi:10.3905/joi.2008.701955.

[3] Greetham T, Hartnett M. *The Investment Clock. Special Report #1: Making Money from Macro*. New York: Merrill Lynch, Global Securities Research & Economics Group; 2004.

[4] Chava S, Hsu A, Zeng L. Does history repeat itself? Business cycle and industry returns. *J Monet Econ*. 2020;116(12):201–218. doi:10.1016/j.jmoneco.2019.10.005.

[5] Chong J, Phillips GM. Sector rotation with macroeconomic factors. *J Wealth Manag*. 2015;18(1):54–68. doi:10.3905/jwm.2015.18.1.054.

[6] Molchanov A, Stangl J. The myth of business cycle sector rotation. *Int J Fin Econ*. 2024;29(4):4419–4442. doi:10.1002/ijfe.2882.

[7] Sauer M. Sector rotation through the business cycle: a machine learning regime approach. SSRN working paper; 2019. doi:10.2139/ssrn.3473907.

[8] Hamilton JD. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*. 1989;57(2):357–384. doi:10.2307/1912559.

[9] Ang A. *Asset Management: A Systematic Approach to Factor Investing*. 1st ed. New York: Oxford University Press; 2014.

[10] Ilmanen A. *Expected Returns: An Investor's Guide to Harvesting Market Rewards*. 1st ed. Chichester: John Wiley & Sons Ltd; 2011.

[11] Fama EF, French KR. Industry costs of equity. *J Fin Econ*. 1997;43(2):153–193. doi:10.1016/S0304-405X(96)00896-3.

[12] Gooch S. Replication code for "Macro Climate Regimes and U.S. Sector Rotation" [software]. Zenodo; 2026. doi:10.5281/zenodo.23023444.

[13] Federal Reserve Bank of St. Louis. FRED and ALFRED economic databases; 2026. Accessed June 30, 2026. https://fred.stlouisfed.org.

[14] French KR, Kenneth R. French data library; 2026. Accessed June 30, 2026. https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

[15] National Bureau of Economic Research. US business cycle expansions and contractions; 2026. Accessed June 30, 2026. https://www.nber.org/research/business-cycle-dating.

[16] Finance Y. Historical ETF and index price data; 2026. Accessed June 30, 2026. https://finance.yahoo.com.

Appendix A

Methodological Detail and Robustness Results

This appendix collects the methodological detail and robustness checks referred to in the main text. Section numbers (A1–A13) are those cited there. Script names refer to the replication archive described under Supplemental Materials.

A1. Data and the macro state

Sources. FRED (macro and market series), the Kenneth French data library (value-weighted 12- and 49-industry portfolios), and Yahoo Finance (Select Sector SPDR ETFs and SPY/RSP). Primary full-feature window: January 1986–June 2026 (485 months), bounded by the Baa–10y credit-spread series.

State vector (11 indicators). Growth: industrial production YoY, payrolls YoY, and the 12-month change in unemployment. Inflation: headline CPI YoY and core PCE YoY. Rates/curve: 10y–2y and 10y–3m spreads, and the 12-month change in the fed-funds rate. Credit: Baa–10y spread. Conditions: Chicago Fed NFCI. Sentiment: 12-month change in the University of Michigan consumer-sentiment index. Slow macro releases (including the sentiment survey) are lagged one month for publication realism; market series are used contemporaneously.

Sample gap. The window contains 485 of a possible 486 months: the November 2025 state is undefined because the October 2025 CPI and household-survey releases were canceled during the autumn-2025 federal government shutdown, so its (lagged) CPI and unemployment inputs do not exist.

Timing (point-in-time). FRED returns current-vintage data, which incorporate later revisions. The one-month lag addresses publication delay but not historical revisions. We therefore also run the full exercise on ALFRED first-release vintages of the revision-prone series (A9); the drawdown result is robust to vintage, but the return advantage is not, and we therefore frame the conclusions based on the conservative real-time figures.

Standardization. Full-sample z-scores for the descriptive taxonomy; expanding-window z-scores (minimum 60 months) for the strictly past-only walk-forward labels.

A2. Choosing the number of climates

Silhouette/inertia for K = 3 . . . 7: K = 3 (0.227), K = 4 (0.244), K = 5 (0.242), K = 6 (0.250), K = 7 (0.256). The silhouette is essentially flat; K = 5 is chosen as the smallest number of distinct, economically nameable climates. The robustness of the main conclusions to K is confirmed in A10 (the qualitative taxonomy is stable; see A3).

Naming the climates (archetype assignment). Each cluster is named from its centroid in *raw* (unstandardized) units. We first collapse the centroid into five derived dimensions, each standardized across the K clusters: growth (average of industrial production and payroll growth, less the 12-month unemployment change), inflation (average of CPI and core PCE inflation), financial stress (average of the Baa–10y spread and the NFCI; the credit-free sample of A8 uses the columns available to it), curve steepness (average of the available term-spread columns), and policy easing (the negative of the 12-month fed-funds change). Each of the five economic archetypes is scored as a fixed linear combination of these dimensions plus the cluster's NBER recession overlap; e.g., *credit-stress contraction* = 1.5 × stress + 1.0 × recession overlap − 1.0 × growth; *Goldilocks expansion* = 1.0 × growth + 0.9 × curve − 0.9 × inflation. We then solve the 5 × 5 one-to-one assignment (Hungarian algorithm) that maximizes total affinity, so every archetype name is used exactly once. The weights are fixed constants set ex ante on economic grounds; the walk-forward classifier of A5 re-applies this same fixed rule to each annual refit's past-only centroids, so naming never uses future information (taxonomy.name_climates).

A3. Clustering robustness

Re-estimating with hierarchical (Ward) clustering and a 5-component Gaussian mixture recovers essentially the same partition: agreement with k-means (adjusted Rand index) is 0.80 (Ward) and 0.74 (GMM). All three methods give the same in-sample dispersion to within rounding (k-means 0.0515, Ward 0.0493, GMM 0.0538). The taxonomy is thus a stable description; whether it *beats the clock* is addressed fairly in A4.

A4. A fair comparison with the investment clock (key result)

Pooled R^2 values of forward 3-month sector excess returns on regime dummies, in sample:

The clustering's apparent edge over the 4-state clock is almost entirely a complexity effect: a 5-state clock built on only growth and inflation matches the 11-variable map (0.0508 vs 0.0515). The block-bootstrap 95% CI for the clustering-minus-4-state-clock gap is [−0.017, +0.071] (straddles zero, P(gap > 0) = 0.81); versus the equal-complexity clock-5, the gap is ≈0 (P = 0.49).

Out-of-sample (strictly past-only walk-forward labels, 2001–2026): clustering R^2 = 0.0381 versus clock R^2 = 0.0493: the clock is the better description out of sample. Conclusion: the climate map does not forecast sector rotation better than the investment clock.

A5. The strictly past-only walk-forward classifier

The label at month-end *t* uses only data through *t*. We burn in 120 months, then each year (a) re-cluster the expanding window with k-means and (b) **name** the clusters by applying the fixed archetype-scoring rule to that window's own past-only raw-feature centroids. Each month is assigned to its nearest

Scheme	States	Variables	R^2	Adjusted R^2
Climate clustering	5	11	0.0515	0.0417
Investment clock	4	2	0.0370	0.0290
Equal-complexity clock-5	5	2 (growth, inflation)	0.0508	0.0410

past-only centroid and labeled with that cluster's name. No full-sample centers or identities are used at any point. (An earlier version aligned refits to full-sample centroids; that introduced look-ahead and has been removed. The headline de-risking result is reported under the corrected, strictly past-only labels.) Live labels run from December 2000 through June 2026. The final refit in the frozen sample occurs in January 2026 and is estimated on data through December 2025; the June 2026 reading in Table 4 of the main article assigns the June state (May 2026 values of the seven one-month-lagged releases, June 30 values of the market series, expanding-window standardization) to those refit centroids.

Naming without NBER overlap (guardrail). The clustering never sees NBER dates, but the archetype-naming rule scores each cluster partly on its (past-only, one-month-lagged) NBER-recession overlap, and the risk dial acts on the two *named* adverse climates. To rule out label leakage through the names, we re-ran the entire walk-forward classifier with the recession-overlap term removed from the naming rule, with names assigned from the macro and financial-stress dimensions only. Results are unchanged: the equal-weight + risk-dial strategy earns a 9.4% CAGR with a 0.65 Sharpe ratio and an identical −23.2% maximum drawdown (baseline: 9.5%, 0.64, −23.2%); monthly labels agree 90.5% of the time and the adverse-versus-benign classification 92.8%. The dial's protection does not depend on recession dates entering the naming step (src/robustness_nber_naming.py).

A6. Significance of the rotation map (Fig. 3)

Per (climate, sector) cell, mean forward 3-month excess return versus zero via stationary block bootstrap (mean block 6 months, 5,000 resamples), with Benjamini–Hochberg FDR correction across the 60 cells. Only three cells survive at 5%, all in the small (20-month) credit-stress state. Fig. 3 is therefore read as an economically coherent *description*, not as a set of individually significant forecasts. Table 2 of the main article reports 90% block-bootstrap intervals for the credit-stress column, and A13 below tabulates them for all sixty cells; bootstrap leadership probabilities are reported in A11, and a leave-2020-out re-clustering in A10.

A7. The climate risk dial: holdout, grid, and simpler benchmarks (key result)

Headline strategy: Equal-weight sector book, moved entirely to cash in the two adverse climates (credit-stress contraction, recessionary easing), past-only labels, cash at the fed funds rate.

Untouched holdout. Development period, 2001–2010: Sharpe ratio 0.57, max DD −15.7%. Evaluation period, 2011–2026 (held out): Sharpe ratio 0.54, max DD −24.5%. The benefit appears in both halves.

Exposure grid (equity weight in credit-stress/recessionary-easing), applied to the sector-rotation book: the ∼−24% drawdown and ∼0.5 Sharpe ratio are stable across (0,0), (0,0.5), (0.25,0.5), and (0,1.0); the worst case is (0.5,0.5). The result is driven mainly by exiting the credit-stress climate, not by the exact weights.

Simpler de-risking rules (equal-weight book, 2001–2026):

The climate dial clearly beats the credit-spread and unemployment rules and lifts Sharpe ratio and drawdown well above buy-and-hold. A plain 10-month trend rule, however, achieves a comparable (slightly smaller) drawdown at a similar Sharpe ratio. We therefore present climate de-risking as competitive with, not dominant over, the best simple alternatives, with the added value of an interpretable macro rationale.

Was the cash set chosen ex ante? An exhaustive sweep. The two risk-off climates are the two whose archetype names are "credit-stress contraction" and "recessionary easing." Those names are assigned by the fixed rule of A2 from each cluster's macro centroid (growth, inflation, financial stress, policy easing, curve steepness) and its NBER-recession overlap; no sector return, market return, or strategy statistic enters the naming rule or the choice of the cash set, and the walk-forward classifier of A5 re-derives the names each year from past-only data. To show that the pair was not selected by inspecting returns, src/robustness_adverse_set.py evaluates the dial for every one of the 30 non-empty proper subsets of the five climates as the cash set, under exactly the headline convention (past-only label at month-end t applied to month $t+1$, cash at the fed funds rate, 2001–June 2026). The single- and two-climate rules are tabulated below; the full sweep is in data/robustness_adverse_set.csv.

Three facts stand out. First, the pre-specified pair is not the return-maximizing choice: holding cash in the credit-stress climate alone earns a higher CAGR (10.9%) and Sharpe ratio (0.71) with the same −23.2% drawdown, so adding recessionary easing costs return; a rule chosen by looking at results would not have included it. Second, every rule that exits the credit-stress climate reaches the −23.2% drawdown or better, and no rule that stays invested through it improves on a −46% drawdown: the protection comes from recognizing credit stress, which is what the paper claims. Third, the subsets that score best on Sharpe ratio (for example, mid-cycle plus credit-stress, 0.78) hold cash through benign, common expansions and are economically incoherent, and the best drawdowns in the sweep (the three- and four-climate sets, down to −16.4% and below) come from being in cash more than half the time at a heavy cost in return. The pre-specified pair ranks second of 31 rules on CAGR, fourth on Sharpe ratio, and shares the −23.2% drawdown with every other rule that includes credit stress. We report it as the economically motivated rule it was, not as the best cell in this table.

Rule	CAGR	Sharpe	Max DD
Buy and hold equal-weight	9.4%	0.55	−49.5%
Climate dial (binary)	9.4%	0.64	−23.2%
Credit-spread threshold	3.9%	0.30	−29.8%
Unemployment-trend rule	6.9%	0.49	−23.2%
10-month trend (MA)	8.3%	0.63	−17.9%
NBER oracle (not real-time)	10.4%	0.67	−26.2%

Cash set (climates held in cash)	Months in cash	CAGR	Sharpe	Max DD
(None: buy and hold)	0	9.4%	0.55	−49.5%
Credit-stress	30	10.9%	0.71	−23.2%
Mid-cycle	112	7.5%	0.51	−46.2%
Goldilocks	99	6.2%	0.37	−49.5%
Late-cycle	40	6.7%	0.39	−49.5%
Rec. easing	22	7.9%	0.47	−49.5%
Mid-cycle + Credit-stress	142	9.0%	0.78	−18.3%
Goldilocks + Credit-stress	129	7.7%	0.54	−23.2%
Late-cycle + Credit-stress	70	8.2%	0.55	−23.2%
Rec. easing + Credit-stress (pre-specified)	52	9.4%	0.64	−23.2%
Mid-cycle + Rec. easing	134	6.1%	0.42	−46.2%
Goldilocks + Mid-cycle	211	4.4%	0.29	−46.2%
Mid-cycle + Late-cycle	152	4.9%	0.32	−46.2%
Goldilocks + Rec. easing	121	4.8%	0.28	−49.5%
Goldilocks + Late-cycle	139	3.6%	0.20	−49.5%
Late-cycle + Rec. easing	62	5.3%	0.31	−49.5%

A8. Long, credit-free sample (1971–2026)

Rebuilding the climates on a credit-free state (no Baa spread/DGS2-3m curves; a 10y-minus-fed-funds curve from the 1960s) extends the sample to 599 monthly observations between March 1971 and May 2026. The count is below the full calendar span because the Michigan sentiment survey is quarterly before 1978, so roughly eight months a year drop out through 1978 (plus the November 2025 gap of A1). Five recognizable climates re-emerge, confirming the taxonomy is not an artifact of the post-1986 window. (We do not claim a dispersion advantage over the clock here; per A4 that advantage does not survive a fair, out-of-sample test.)

A9. Point-in-time (ALFRED) validation (key robustness)

We rebuild the macro state from ALFRED **first-release** vintages of the five revision-prone monthly series (industrial production, payrolls, unemployment, CPI, core PCE; market series are not meaningfully revised) and re-run the strictly past-only walk-forward labels and the headline equal-weight + climate-dial strategy.

Data vintage	CAGR	Vol	Sharpe	Max DD
Current-vintage	9.5%	12.5%	0.64	−23.2%
Real-time, as-of-date (ALFRED)	9.0%	11.9%	0.63	−23.2%
First-release only (strict bound)	7.2%	11.7%	0.50	−23.7%

The drawdown protection is robust across all three vintages. Under the *as-of-date* reconstruction (the data exactly as it stood at each decision date, the most faithful real-time test), the Sharpe advantage is essentially intact (0.63). Only the stricter *first-release* test, which discards even revisions an investor would have seen, trims the Sharpe ratio to 0.50. (As-of-date uses the four deep-vintage series with current-vintage core PCE, whose ALFRED vintages start only in 2000; core PCE is barely revised.) Regenerate the results via src/pit_ingest.py, src/pit_validation.py, and src/pit_asof.py.

A10. K, COVID, and variable-weighting robustness

Number of climates (K). In-sample dispersion and an indicative in-sample dial:

Leave-2020-out. Re-clustering with all of 2020 removed yields virtually the same partition (adjusted Rand index 0.98 on non-2020 months), and 75% of 2020 months are still classified as adverse when assigned to the leave-out centroids; the COVID shock does not manufacture the taxonomy.

Block-composite state. Collapsing the eleven indicators into five block composites (growth, inflation, curve, credit, conditions) gives similar dispersion (R^2 0.048) but weaker drawdown protection (−38%), indicating that the de-risking benefit draws on the finer resolution of the full state. Regenerate the results via src/robustness_batch.py.

A11. Leadership probabilities (Fig. 3 confidence)

Rather than treating the point estimates in Fig. 3 as forecasts, we report the block-bootstrap probability that each sector ranks in the top three by forward 3-month excess within each climate. Sectors reaching ≥60% probability are, for goldilocks, manufacturing and business equipment; for mid-cycle, business equipment; for late-cycle overheating, consumer nondurables and energy; for recessionary easing, nondurables and chemicals; and for credit-stress contraction, durables and manufacturing (the last based

K	Dispersion R^2	Adverse climates	Dial CAGR	Dial Sharpe	Dial Max DD
4	0.029	1	7.5%	0.48	−38.2%
5	0.052	2	9.0%	0.62	−23.2%
6	0.043	3	8.7%	0.60	−23.2%
7	0.048	3	8.7%	0.60	−23.2%

K = 5, 6, 7 give essentially the same drawdown protection; K = 4 under-resolves the adverse states (one adverse cluster, weaker protection), supporting five as a floor.

Sector	Goldilocks expansion (n = 146)	Mid-cycle expansion (n = 180)	Late-cycle overheating (n = 72)	Recessionary easing (n = 62)	Credit-stress contraction (n = 20)
NoDur	−1.3 [−4.5, +1.9]	−3.2 [−6.7, +0.4]	+5.3 [+0.7, +9.9]	+5.6 [+0.4, +11.0]	−8.5 [−12.2, −4.6]
Durbl	+2.0 [−2.7, +6.6]	−3.2 [−8.6, +2.7]	−3.3 [−14.2, +7.9]	−1.0 [−12.2, +9.7]	+67.5 [+44.2, +91.8]
Manuf	+2.2 [+0.1, +4.2]	+1.4 [−1.6, +4.2]	−1.0 [−4.6, +3.0]	+0.3 [−4.6, +5.3]	+7.1 [+2.0, +12.4]
Enrgy	+0.1 [−6.0, +6.2]	−0.6 [−7.2, +5.9]	+10.7 [−1.0, +25.1]	+1.2 [−7.2, +9.6]	−19.2 [−36.3, −1.0]
Chems	+0.8 [−1.3, +3.1]	−4.3 [−8.0, −0.8]	−1.3 [−3.7, +1.0]	+5.8 [+0.3, +11.7]	−6.3 [−14.4, +1.4]
BusEq	+2.8 [−0.4, +6.1]	+7.5 [−0.8, +15.4]	−5.0 [−13.4, +3.9]	−5.8 [−13.4, +1.4]	+8.4 [+1.2, +16.0]
Telcm	−0.5 [−3.8, +2.7]	−1.6 [−7.4, +3.9]	−3.2 [−12.5, +6.2]	−7.1 [−15.7, +1.1]	−6.3 [−10.7, −2.0]
Utils	−2.7 [−7.2, +1.5]	+2.0 [−2.7, +6.8]	+0.3 [−6.5, +7.0]	−4.0 [−8.7, +0.8]	−20.3 [−27.2, −13.9]
Shops	−0.8 [−3.5, +2.2]	+1.4 [−2.1, +4.9]	+0.3 [−4.0, +4.3]	+6.0 [+2.2, +10.0]	−3.5 [−11.2, +4.4]
Hlth	−1.1 [−5.2, +3.0]	+3.0 [−0.5, +6.5]	+4.2 [−1.8, +10.9]	−1.8 [−10.4, +6.8]	−10.3 [−18.2, −2.1]
Money	−1.6 [−5.3, +2.1]	+1.6 [−2.0, +5.1]	−4.0 [−9.1, +0.6]	+4.7 [−3.1, +11.3]	−6.3 [−19.3, +7.1]
Other	0.0 [−1.5, +1.5]	−3.9 [−6.2, −1.6]	−2.9 [−7.6, +1.5]	−4.0 [−7.3, −0.5]	−2.4 [−9.6, +4.1]

on only 20 months). Regenerate the results via src/leadership_prob.py.

A12. Further extensions

As-of-date vintages. Reported in A9 (drawdown robust; Sharpe ratio 0.63 under faithful real-time data). *Confidence band*. Fig. 5 of the main article shades the 5–95% block-bootstrap range of cumulative paths for the headline strategy. *Volatility targeting*. Adding a 12% trailing-volatility target to the climate dial (unlevered, exposure capped at one) gives a CAGR of 7.8%, a Sharpe ratio of 0.61, and a maximum drawdown of −19.9%, smoother still, at a modest cost in return (src/risk_dial_tests.py). Remaining for future work: weekly re-classification and a multi-asset (cross-region) extension.

A13. Bootstrap intervals for every cell of the rotation map

Table 2 of the main article reports 90% stationary-block-bootstrap intervals (mean block six months, 5,000 resamples) for the twelve credit-stress cells in Fig. 3. The table below gives the same intervals for all sixty cells: annualized forward 3-month excess returns in percentage points, with the 5th and 95th bootstrap percentiles in brackets (src/crisis_uncertainty.py, data/rotation_ff_ci.csv).

Interval width tracks sample size: the median width is 6.5 points in goldilocks (146 months) and 7.2 in mid-cycle (180 months), against 12.9 in recessionary easing (62) and 15.2 in credit-stress contraction (20). The credit-stress column has the most intervals that exclude zero (8 of 12), but this reflects how extreme and internally consistent its two episodes were, not a large sample: the block bootstrap resamples within twenty months drawn from two crises, so it captures month-to-month noise within those episodes and cannot capture how different the next credit crisis might be. The main text therefore treats the direction of the credit-stress rotation (deep cyclicals rebound, while energy and utilities lag) as the finding, and the magnitudes as illustrative.

About the Author



Sasha Gooch is a Grade 12 student at The American School in London, where this paper began as an independent research project in quantitative finance. Sasha's interests lie at the intersection of economics, statistics, and software, with a particular focus on how machine-learning methods can be tested honestly against the simpler rules they are meant to replace. Outside this research, Sasha builds educational software, including an artificial-intelligence mathematics tutor used by students at ASL, and plans to study mathematics and economics at university.