

— GMSI INSIGHT · GLOBAL MACRO SENTIMENT INDICES

From research to live: putting macro sentiment to work

Turning macro narrative into a point-in-time feature layer for systematic research, validation and production.

Narrative moves before the data does

Macro narratives move markets before official data confirms them. The difficulty is what to do with the concession — because **narrative arrives as language, and language does not back-test.**

The challenge for systematic investors is turning that narrative into a point-in-time feature that can be tested, validated and reproduced live.

Before language can sit in a model, it needs

- ◆ **A country attached**
Where the signal applies
- ◆ **A driver attached**
Which economic channel is moving
- ◆ **A magnitude**
How far it moved, on what weight of coverage
- ◆ **Enough history**
To show whether it ever carried information

Global headline flow, one schema

GMSI reads global headline flow and returns **hourly, point-in-time signals** — from economic data and monetary policy to fiscal stance, geopolitics, trade and market stress.

95

Markets covered

70

Macro topics across 8 pillars

11 yrs

Point-in-time history

250k+

Curated publishers

100+

Countries, any language

Hourly

Historical & live, one schema

Non-English reporting is detected, translated and run through the same classifier — a headline filed in Tokyo or São Paulo lands in the same framework as one from New York or London.



A narrative object, not a mood score

The point of GMSI is to make macro narrative **behave like market data**.

A quant does not simply learn that US inflation sentiment rose. They learn how far, on what weight of coverage, how much of it disagreed, and whether it came from inside the country or outside.

Country

Topic

Sentiment type

Index type

Timestamp

Fields carried beside every reading

HEADLINE COUNT

How many headlines fed the reading

NET SENTIMENT

The direction and weight of the flow

SIGNAL MAGNITUDE

The total strength behind the move

DISPERSION

How much the coverage disagreed



Kept apart, each driver stays readable

Folded into one score, drivers cancel and blur. Three separations carry most of the weight.

01

THE DRIVER

The economic channel you are trading. A rates model wants policy outlook and inflation pressure; an FX model wants relative policy and external vulnerability.

02

DIRECTION VS TONE

Directional sentiment tracks the topic itself; semantic sentiment tracks the register it is told in. Direction belongs near the forecast, tone near the risk.

03

SOURCE PERSPECTIVE

Domestic, international or combined. Where local and foreign perception stop agreeing tends to be where a repricing starts.



Eight choices are the signal

You start with a market expression — a point on the curve, a currency pair, a policy-rate path — then **fix the feature before you let the data judge it.**

1

Country

Where the signal applies

2

Topic

Which driver is moving

3

Sentiment type

Direction or tone

4

Index type

The series used

5

Coverage view

Domestic, international or combined

6

Lookback

The aggregation window

7

Normalisation

Fixed or expanding z-score

8

Cadence

How often it trades

The data is hourly, but few strategies trade hourly. A medium-horizon book might care about pressure built over 28 or 84 days; a live rates desk only about the hour a policy statement lands — those eight choices are not settings around the signal, they are the signal.



The same feature, both sides of the line

VALIDATION DISCIPLINE

The bar is not the best in-sample fit — it is a mechanism you can explain and a result that survives the training window.

- Walk-forward & lead-lag tests
- A benchmark to beat, a stable sign
- Broken out by country, topic & sentiment type

LIVE REPRODUCTION

The part that quietly sinks most sentiment projects. Both endpoints hand over the same raw structure — neither hands over a finished feature.

- Same aggregation window & normalisation
- Same coverage view & lag
- Same treatment of sparse observations




US policy outlook and the front end

- Feature — US Monetary Policy · Policy Outlook · combined directional index

One driver mapped to the stretch of curve where its channel is strongest — **built client-side**, then run unchanged on the live feed.

- **Raw endpoint input**
Hourly GMSI directional sentiment
- **Client-side transform**
84-day cumulative pressure, z-score, lagged one day
- **Market target**
US 2Y yield — the front-end series tested



A +1 σ rise in Policy Outlook

3-MONTH YIELD

+18.6bp

Over the next 60 trading days

2-YEAR YIELD

+15.9bp


Over the same 60-day window

2S10S CURVE

Flatter

Repricing lands at the front, not spread evenly

A +1 σ rise in Policy Outlook — repricing concentrated at the **policy-sensitive front end**.



One feature, several jobs

REGIME CLASSIFIER

Rolling features flag when policy, inflation, growth or geopolitical pressure is building or fading ahead of the official print

CROSS-SECTIONAL FACTOR

A single topic ranked across economies — policy sentiment or inflation breadth, comparable from developed to frontier

OVERLAY

Wave a trade through, hold it back when the narrative disagrees, or warn when a position runs on story rather than data

RISK CONTROL

Rising semantic stress or a widening domestic–international gap can bring exposure down before volatility shows in price

None of that turns on a particular instrument. It turns on the lag between a narrative forming and the data confirming it — the window a point-in-time sentiment feature is built to trade.



An unbroken line from research to production — **the same raw data, structured the same way, on history and live alike.**

95+

Economies

70

Macro topics

11 yrs

Point-in-time history

Hourly

Directional & semantic series

You still build the feature. What GMSI provides is the input beneath it — **and the transformation in between stays identical on both sides.**