

Context-Aware RSI

Integrating Global M2 Money Supply and News Sentiment

A Design Proposal for Adaptive Technical Analysis

Martin Russmann

Independent Researcher · ORCID 0009-0006-4079-1216

mrussmann@proton.me · martinrussmann.com

ABSTRACT

This design proposal presents a framework for enhancing the Relative Strength Index (RSI) by incorporating macroeconomic liquidity conditions and market sentiment. The traditional RSI, while widely used, suffers from well-documented limitations including false signals during trending markets and fixed thresholds that fail to adapt to changing market conditions. The core design principle is to **preserve the canonical RSI scale** $[0, 100]$ while adapting interpretation through dynamic thresholds and conditional normalization, maintaining interpretability for practitioners familiar with traditional RSI while systematically accounting for the external factors — global liquidity and news sentiment — that influence asset prices.

KEYWORDS Relative Strength Index; Global M2 Money Supply; News Sentiment Analysis; Adaptive Technical Indicators; Liquidity-Aware Trading; Technical Analysis Enhancement

1 Executive Summary



This design proposal presents a comprehensive framework for enhancing the Relative Strength Index (RSI) by incorporating macroeconomic liquidity conditions and market sentiment. The traditional RSI, while widely used, suffers from well-documented limitations including false signals during trending markets and fixed thresholds that fail to adapt to changing market conditions.

Our core design principle is to **preserve the canonical RSI scale** $[0, 100]$ while adapting interpretation through dynamic thresholds and conditional normalization. This approach maintains interpretability for practitioners familiar with traditional RSI while systematically accounting for external factors that influence asset prices.

Key contributions of this proposal:

- A canonical design that preserves RSI's bounded nature while enabling context-awareness
- An explicit two-convention definition of global M2 (level aggregation vs. GDP-weighted growth aggregation), with a corrected currency treatment and a recommendation for the growth convention in regime classification
- A distribution-free conditional-percentile signal alongside the Gaussian Z-score, addressing RSI's bounded support
- A regularized calibration framework with overfitting safeguards and a mandatory ablation design
- Operational specifications for real-time implementation, including point-in-time discipline for all conditional statistics
- A working prototype for practitioner validation

Keywords: Relative Strength Index, Global M2 Money Supply, News Sentiment Analysis, Adaptive Technical Indicators, Liquidity-Aware Trading, Technical Analysis Enhancement

2 Introduction



2.1 Background and Motivation

The Relative Strength Index (RSI), introduced by J. Welles Wilder in his seminal 1978 work *New Concepts in Technical Trading Systems* [18], has become one of the most widely used momentum oscillators in technical analysis. The indicator measures the speed and magnitude of price movements on a bounded scale from 0 to 100, with readings above 70 traditionally interpreted as “overbought” and readings below 30 as “oversold.”

Despite its enduring popularity among traders and analysts, the RSI suffers from several well-documented limitations:

1. **False signals in trending markets:** During sustained bull markets, RSI frequently signals overbought conditions while prices continue rising, leading to premature exits or missed opportunities [2].
2. **Static threshold assumptions:** The 70/30 boundaries assume stationary market conditions, yet market dynamics vary substantially across liquidity regimes and sentiment environments.
3. **Price-only calculation:** RSI ignores fundamental drivers of price movements, including macroeconomic liquidity conditions and market psychology.
4. **Arbitrary parameterization:** The standard 14-period lookback lacks theoretical justification and may not be optimal across different assets and timeframes.

These limitations motivate the central question of this proposal: *Can we enhance RSI to account for external market conditions while preserving its interpretability and bounded nature?*

2.2 The Case for M2 Money Supply

Global money supply, particularly the broad M2 aggregate, serves as a fundamental driver of asset prices across multiple channels:

Portfolio rebalancing channel: When central banks expand money supply, the additional liquidity flows into financial assets as investors seek returns, pushing up prices [17].

Discount rate channel: Monetary expansion typically accompanies lower interest rates, reducing the discount rate applied to future cash flows and increasing present values [3].

Risk appetite channel: Abundant liquidity tends to compress risk premiums and encourage risk-taking behavior, amplifying price movements [15].

Empirical research consistently documents significant relationships between M2 growth and equity returns. Bernanke and Kuttner [4] find that a surprise 25-basis-point cut in the federal funds rate is associated with approximately a 1% increase in broad stock indices. Thorbecke [17] demonstrates that expansionary monetary policy reliably increases stock returns across multiple identification strategies.

Why M2 over other macro aggregates? We choose M2 over alternatives for several reasons:

- **Breadth:** M2 captures a wider range of liquid assets than M1, including savings deposits and money market funds
- **Observability:** Unlike credit conditions or real rates, M2 is directly measured and consistently reported across major economies

- **Leading indicator properties:** M2 changes often precede asset price movements, providing predictive value
- **Cross-country comparability:** Central banks report M2 using broadly similar—though not identical—definitions; remaining definitional heterogeneity is discussed in Section 5

2.3 The Role of Market Sentiment

News sentiment captures the collective psychological state of market participants and serves as a distinct information source from price data alone. Tetlock’s foundational work [16] demonstrates that media pessimism predicts downward pressure on market prices, while Engelberg and Parsons [6] provide causal evidence that local media coverage directly influences local trading activity.

Sentiment affects markets through several mechanisms:

Attention channel: News coverage directs investor attention to specific assets or themes, influencing trading activity [6].

Belief updating channel: Investors update their expectations based on news content, affecting demand and supply [16].

Coordination channel: Public information serves as a focal point for coordinating beliefs and actions across dispersed market participants.

By incorporating sentiment, we can potentially anticipate shifts in market direction before they fully manifest in price-based indicators like RSI. An important caveat—that sentiment is itself partly endogenous to recent returns—is addressed in Section 6.4.

2.4 Design Philosophy

A critical challenge in enhancing RSI is maintaining its bounded, interpretable nature. Many proposed modifications—including simple multiplicative adjustments—break the fundamental property that RSI lies in $[0, 100]$. An “enhanced RSI” of 85 or 120 loses intuitive meaning and comparability with historical readings.

We therefore adopt the following canonical design principle:

Design Choice: The base RSI calculation remains unchanged, preserving the $[0, 100]$ scale. External factors (M2, sentiment, volatility) influence interpretation exclusively through *dynamic thresholds* and *conditional normalization*, never by direct transformation of the RSI value itself.

This design ensures that:

- RSI values remain directly comparable across time periods and assets

- Practitioners retain intuitive understanding of the indicator
- Context-awareness is achieved without breaking mathematical bounds
- The enhancement is additive to existing workflows rather than disruptive

3 Literature Review



3.1 RSI Limitations and Extensions

The academic and practitioner literature has long recognized RSI's limitations. Cardwell pioneered the concept of "positive and negative reversals," demonstrating that RSI behavior differs systematically in bull versus bear markets; the method is documented in detail by Baeyens [2]. This work implicitly suggests that interpretation should be conditioned on market regime—an insight we formalize in this proposal.

Murphy's comprehensive treatment [13] notes that RSI tends to oscillate between 40 and 80 during bull markets, rarely reaching the traditional 30 oversold level. Conversely, during bear markets, RSI typically ranges between 20 and 60. This regime-dependent behavior motivates our conditional normalization approach.

3.2 Monetary Policy and Asset Prices

The relationship between monetary policy and asset prices has been extensively studied. Key findings relevant to our framework include:

Thorbecke (1997) [17]: Using multiple identification strategies (VAR innovations, narrative indicators, event studies), this seminal paper establishes that expansionary monetary policy reliably increases stock returns. The effect operates through both expected cash flows and discount rates.

Rigobon and Sack (2004) [15]: Employing heteroskedasticity-based identification, this study finds that a 25-basis-point increase in short-term rates leads to a decline in stock prices of approximately 1.7%. The methodology addresses endogeneity concerns that plague earlier studies.

Bernanke and Kuttner (2005) [4]: This paper decomposes the stock market's response to monetary policy into effects on expected dividends, real interest rates, and equity premiums. Most of the response operates through the equity premium channel, suggesting that liquidity conditions affect risk appetites.

Pícha (2017) [14]: Examining the long-run relationship between money supply and stock indices across multiple countries, this study confirms positive cointegration for several major markets including the US, UK, and emerging economies.

3.3 Sentiment Analysis in Finance

The application of sentiment analysis to financial markets has evolved rapidly:

Lexicon-based approaches: Early work relied on word lists such as the Harvard General Inquirer or finance-specific dictionaries. Loughran and McDonald [11] demonstrate that general-purpose sentiment dictionaries perform poorly in financial contexts, motivating domain-specific lexicons.

VADER: Hutto and Gilbert [8] develop VADER (Valence Aware Dictionary and sEntiment Reasoner), which incorporates grammatical and syntactical conventions for improved accuracy. VADER provides compound sentiment scores in $[-1, 1]$ and is computationally efficient for real-time applications.

Transformer models: Recent advances in NLP enable more nuanced sentiment analysis. FinBERT [1] fine-tunes BERT on financial text, achieving state-of-the-art performance on financial sentiment classification tasks.

Commercial solutions: Providers like RavenPack offer institutional-grade sentiment feeds aggregated from thousands of news sources with low latency, suitable for systematic trading applications.

3.4 Adaptive Technical Indicators

Recent work has explored conditioning technical indicators on external factors:

Mostafavi and Hooman [12] examine which technical indicators provide the most predictive value when combined with machine learning models, finding that adaptive indicators outperform static versions.

Latif et al. [9] integrate macroeconomic indicators with technical signals using data-driven approaches, demonstrating improved forecasting accuracy for stock market direction.

Gerunov [7] applies machine learning to combine macroeconomic, technical, and sentiment features for market forecasting, providing a template for multi-factor integration.

However, few studies specifically address enhancing RSI with macroeconomic and sentiment conditioning—a gap this proposal aims to fill.

4 Standard RSI: Calculation and Limitations

4.1 Traditional RSI Formula

The RSI compares the magnitude of recent gains to recent losses over a specified lookback period n (typically 14 periods). The calculation proceeds as follows:



Step 1: Compute price changes

$$\delta_t = \text{Close}_t - \text{Close}_{t-1} \quad (1)$$

Step 2: Separate gains and losses

$$G_t = \max(\delta_t, 0) \quad (2)$$

$$L_t = \max(-\delta_t, 0) \quad (3)$$

Step 3: Compute smoothed averages using Wilder's exponential moving average

$$\bar{G}_t = \frac{\bar{G}_{t-1} \times (n-1) + G_t}{n} \quad (4)$$

$$\bar{L}_t = \frac{\bar{L}_{t-1} \times (n-1) + L_t}{n} \quad (5)$$

Initialization: The recursion is seeded with simple averages over the first n observations, $\bar{G}_n = \frac{1}{n} \sum_{t=1}^n G_t$ and $\bar{L}_n = \frac{1}{n} \sum_{t=1}^n L_t$, following Wilder's original specification. Implementations that seed differently (e.g., with the first observation) produce transient discrepancies that decay geometrically but matter for short backtests.

Step 4: Calculate Relative Strength

$$RS_t = \frac{\bar{G}_t}{\bar{L}_t} \quad (6)$$

Step 5: Normalize to [0, 100]

$$RSI_t = 100 - \frac{100}{1 + RS_t} = \frac{100 \cdot RS_t}{1 + RS_t} \quad (7)$$

The RSI is bounded by construction: when $\bar{L}_t \rightarrow 0$ (all gains), $RS_t \rightarrow \infty$ and $RSI_t \rightarrow 100$; when $\bar{G}_t \rightarrow 0$ (all losses), $RS_t \rightarrow 0$ and $RSI_t \rightarrow 0$.

4.2 Detailed Analysis of Limitations

4.2.1 False Signals in Trending Markets

Consider a persistent uptrend accompanied by monetary expansion. As prices rise consistently, RSI will naturally trend toward overbought territory. Under traditional interpretation, this signals a sell—yet the underlying liquidity conditions support continued price appreciation.

Illustrative scenario: During 2020-2021, global M2 expanded by approximately 25% as central banks responded to the pandemic. Equity markets rose substantially, with the S&P 500 gaining over 40% from March 2020 lows. Throughout this period, RSI frequently exceeded 70, generating repeated overbought signals that would have resulted in premature exits from a profitable trend.

The fundamental issue is that traditional RSI interpretation assumes mean-reverting price behavior, but liquidity-driven trends can persist far longer than price-only analysis suggests.

4.2.2 Static Threshold Problem

The 70/30 thresholds are essentially arbitrary conventions without theoretical foundation. Empirically, optimal thresholds vary across:

- **Asset classes:** Highly volatile assets like cryptocurrencies may require wider thresholds (e.g., 80/20)
- **Market regimes:** Bull markets exhibit higher average RSI than bear markets
- **Liquidity conditions:** Expansionary environments support higher sustainable RSI levels
- **Volatility regimes:** High volatility produces more extreme RSI readings

4.2.3 Lack of External Context

RSI's exclusive reliance on price data means it cannot distinguish between:

- Price rises driven by fundamental liquidity expansion (sustainable)
- Price rises driven by speculative excess without fundamental support (unsustainable)
- Price declines during temporary sentiment shocks (buying opportunity)
- Price declines reflecting fundamental deterioration (continued risk)

By incorporating M2 and sentiment, we aim to provide this missing context.

5 Global M2 Money Supply: Definition and Construction



5.1 Two Aggregation Conventions

Aggregating national money stocks into a global measure admits two conventions, and it is important not to conflate them.

(a) Level aggregate. The natural level measure of global broad money is the *unweighted* sum of national aggregates converted to a common currency:

$$M2_t^{global,USD} = \sum_{i \in \mathcal{I}} \frac{M2_t^{(i)}}{FX_t^{(i)}} \quad (8)$$

where $M2_t^{(i)}$ is denominated in local currency and $FX_t^{(i)}$ denotes units of currency i per USD (so division converts to USD). No explicit weights appear: each economy's contribution is its actual money stock, and economic size is therefore reflected intrinsically. Applying GDP weights to *levels* would double-count size and yield a quantity without a clean economic interpretation. Note that in this convention China contributes the largest single share of the global stock in USD terms.

(b) Growth aggregate (recommended for regime classification). For the regime signal we recommend the GDP-weighted average of *local-currency* year-over-year growth rates:

$$\Delta M2_t^{global} = \sum_{i \in \mathcal{C}} w_i \Delta M2_t^{(i)}, \quad w_i = \frac{GDP_i}{\sum_{j \in \mathcal{C}} GDP_j}, \quad \Delta M2_t^{(i)} = \frac{M2_t^{(i)} - M2_{t-12}^{(i)}}{M2_{t-12}^{(i)}} \quad (9)$$

computed on monthly data (replace 12 with 252 if a daily interpolated series is differenced instead).

The growth convention has a decisive advantage for our purposes: because each growth rate is computed in local currency *before* aggregation, exchange-rate fluctuations are eliminated from the measure entirely. Under the level convention, a broad USD rally mechanically shrinks the USD value of foreign money stocks—a USD-level global M2 can contract even while every constituent central bank is easing. Since the regime classifier is meant to capture the global liquidity *impulse* rather than the dollar value of the stock, the weighted-growth aggregate is the appropriate input. The level aggregate (Eq. 8) remains useful for descriptive and charting purposes.

The constituent set $\mathcal{C} = \{\text{US, Eurozone, China, Japan, UK}\}$ comprises the five largest contributors to global broad money and represents roughly 70–75% of global economic output.

Economy	GDP Weight w_i	Update Frequency	Data Source
United States	0.35	Monthly	FRED (M2SL)
Eurozone	0.22	Monthly	ECB Statistical Data Warehouse
China	0.25	Monthly	People's Bank of China
Japan	0.10	Monthly	Bank of Japan
United Kingdom	0.08	Monthly	Bank of England (M4)

Table 1: Global M2 composition. GDP weights apply to the growth aggregate (Eq. 9) only; the level aggregate is unweighted. Weights are approximate and should be updated annually based on IMF World Economic Outlook GDP data.

Definitional heterogeneity. National “M2” measures are not perfectly harmonized: the UK does not publish an M2 series and M4 (broader) is used as proxy; the ECB's headline aggregate is M3, with M2 available as a component; China's M2 includes corporate deposits at a scale uncommon elsewhere. These differences affect levels far more than growth rates, which is a further argument for the growth convention. The composition and any proxy substitutions must be documented and held fixed within a backtest.

5.2 Regime Classification

Regime classification applies rolling percentiles over a 5-year lookback to the growth aggregate, establishing context-relative thresholds:

$$R_t^{M2} = \text{Percentile}(\Delta M2_t^{global}, \text{lookback} = 1260 \text{ days}) \quad (10)$$

We partition the percentile distribution into four regimes:

Regime	Percentile Range	Interpretation
Contraction	$R_t^{M2} < 25$	Tightening liquidity; defensive positioning warranted
Below-average	$25 \leq R_t^{M2} < 50$	Moderate liquidity; neutral conditions
Above-average	$50 \leq R_t^{M2} < 75$	Supportive liquidity; bullish lean appropriate
Expansion	$R_t^{M2} \geq 75$	Strong liquidity tailwind; extended over-bought conditions sustainable

Table 2: M2 regime definitions with interpretive guidance

Minimum sample size requirement: Each regime bin must contain at least 63 observations (approximately one quarter of trading days) for stable statistical estimation. If this threshold is not met, adjacent bins should be merged.

5.3 US-Centric Proxies and Their Limitations

For real-time applications, official M2 data arrives with 2-4 week delays. We can estimate current US M2 using high-frequency proxies:

Proxy Variable	Frequency	Rationale
Fed Balance Sheet	Weekly	Direct measure of central bank liquidity provision
Treasury General Account	Daily	Government cash holdings affect money in circulation
Overnight Reverse Repo	Daily	Measures excess liquidity parked at the Fed
Bank Reserve Balances	Weekly	Indicates banking system liquidity

Table 3: High-frequency proxies for US M2 nowcasting

Important limitation: These proxies are US-centric and may not capture global liquidity dynamics. For non-US assets, particularly emerging market equities or local currency bonds, consider:

- Using lagged official global M2 data (accepting the delay)
- Weighting local M2 alongside global M2
- Using asset-class-specific liquidity measures (e.g., FX reserves for EM)

6 Sentiment Quantification



6.1 Sentiment Score Construction

We aggregate news sentiment using a recency-weighted average that gives greater influence to recent articles:

$$S_t = \frac{\sum_{i=1}^N e^{-\lambda(t-t_i)} \cdot s_i}{\sum_{i=1}^N e^{-\lambda(t-t_i)}} \quad (11)$$

where:

- $s_i \in [-1, 1]$ is the sentiment score of article i
- t_i is the publication timestamp of article i
- N is the number of articles in the lookback window
- λ controls recency decay (higher λ = faster decay)

Recommended parameter: $\lambda = 0.1$ per day, implying a half-life of approximately 7 days.

6.2 Sentiment Scoring Methods

Method	Description	Advantages	Disadvantages
VADER [8]	Rule-based lexicon with grammatical heuristics	Fast, no training required	Limited financial domain knowledge
FinBERT [1]	BERT fine-tuned on financial text	High accuracy, context-aware	Computationally intensive
RavenPack	Commercial aggregated feed	Low latency, broad coverage	Costly, black-box

Table 4: Sentiment scoring methods with tradeoffs

6.3 Sentiment Regime Classification

Analogous to M2, we classify sentiment into regimes using rolling percentiles:

$$R_t^S = \text{Percentile}(S_t, \text{lookback} = 252 \text{ days}) \quad (12)$$

The shorter lookback (1 year vs. 5 years for M2) reflects sentiment's higher frequency dynamics.

6.4 Endogeneity of Sentiment and Orthogonalization

News sentiment is not exogenous to prices: coverage and tone respond to recent returns, which are precisely the inputs to RSI. Conditioning RSI interpretation on raw sentiment therefore risks conditioning the indicator on a noisy re-measurement of its own momentum, inflating apparent explanatory power while adding little independent information.

To isolate the incremental content of sentiment, we orthogonalize S_t against trailing returns before regime classification:

$$S_t = a + \sum_{k \in \{1, 5, 20\}} b_k r_{t-k:t} + u_t, \quad S_t^\perp := \hat{u}_t \quad (13)$$

where $r_{t-k:t}$ denotes the cumulative log return over the trailing k days and the regression is estimated on an expanding window using only data available at time t . The regime percentile R_t^S is then computed from S_t^\perp . The validation protocol (Section 12) requires reporting results with both raw and orthogonalized sentiment; if the framework's value disappears under orthogonalization, the sentiment channel is redundant with price momentum and should be dropped.

6.5 Event-Time vs. Clock-Time Aggregation

Two approaches exist for sentiment aggregation:

Clock-time: Aggregate all articles within fixed time windows (e.g., daily). Simple but may miss intraday sentiment shifts around major events.

Event-time: Weight articles by proximity to significant events (earnings, Fed announcements, economic releases). More responsive but requires event calendar integration.

For most applications, daily clock-time aggregation with recency weighting provides a reasonable balance of simplicity and responsiveness.

7 The Canonical Design: Dynamic Thresholds with Conditional Normalization

7.1 Design Rationale

We deliberately reject approaches that directly transform RSI values (e.g., multiplicative scaling) because they:

1. Break the $[0, 100]$ bound, destroying comparability
2. Can produce negative or > 100 values without clear interpretation
3. Mix conceptually distinct adjustments in a single number

Instead, we adapt RSI *interpretation* through two complementary mechanisms:



- **Method 1:** Adjust the thresholds that define overbought/oversold
- **Method 2:** Normalize RSI relative to regime-conditional distributions

Both methods preserve the original RSI value while providing context-aware signals.

7.2 Method 1: Regime-Conditional Dynamic Thresholds

The overbought and oversold thresholds adjust based on the joint regime state:

$$\theta_t^{OB} = \theta_{base}^{OB} + \alpha_{M2} \cdot (R_t^{M2} - 50) + \alpha_S \cdot (R_t^S - 50) + \alpha_\sigma \cdot (R_t^\sigma - 50) \quad (14)$$

$$\theta_t^{OS} = \theta_{base}^{OS} - \alpha_{M2} \cdot (R_t^{M2} - 50) - \alpha_S \cdot (R_t^S - 50) - \alpha_\sigma \cdot (R_t^\sigma - 50) \quad (15)$$

where:

- $\theta_{base}^{OB} = 70$, $\theta_{base}^{OS} = 30$ (standard thresholds)
- $R_t^{M2}, R_t^S, R_t^\sigma \in [0, 100]$ are regime percentiles
- $\alpha_{M2}, \alpha_S, \alpha_\sigma > 0$ are sensitivity parameters (threshold points per percentile deviation from median)

Admissibility constraints: To guarantee well-defined signals under any parameterization (including future recalibrations), the adjusted thresholds are clipped to admissible ranges,

$$\theta_t^{OB} \leftarrow \min(\max(\theta_t^{OB}, 60), 85), \quad \theta_t^{OS} \leftarrow \min(\max(\theta_t^{OS}, 15), 40), \quad (16)$$

and a minimum separation $\theta_t^{OB} - \theta_t^{OS} \geq 20$ is enforced by symmetric shrinkage toward the midpoint $(\theta_t^{OB} + \theta_t^{OS})/2$ whenever it is violated. Without such constraints, nothing in the additive specification prevents the thresholds from converging or drifting into uninterpretable territory under extreme joint regimes.

Example: Suppose $R_t^{M2} = 80$ (strong expansion), $R_t^S = 60$ (mildly positive sentiment), $R_t^\sigma = 40$ (below-average volatility), and $\alpha_{M2} = \alpha_S = \alpha_\sigma = 0.1$.

$$\begin{aligned} \theta_t^{OB} &= 70 + 0.1 \cdot (80 - 50) + 0.1 \cdot (60 - 50) + 0.1 \cdot (40 - 50) \\ &= 70 + 3 + 1 - 1 = 73 \end{aligned}$$

The overbought threshold rises from 70 to 73, reflecting that higher RSI values are sustainable given supportive liquidity and sentiment conditions.

7.3 Method 2: Conditional Normalization (Recommended Primary Method)

This method normalizes RSI readings relative to the historical distribution conditional on the current regime:

$$Z_t^{RSI} = \frac{RSI_t - \mu_{RSI|R}}{\sigma_{RSI|R}} \quad (17)$$

where $\mu_{RSI|R}$ and $\sigma_{RSI|R}$ are the mean and standard deviation of RSI conditional on regime $R = (R^{M2}, R^S, R^\sigma)$.

Practical implementation: Rather than conditioning on the full joint distribution (which may have sparse cells), we use a simplified additive adjustment:

$$\mu_{RSI|R} = \mu_{RSI} + \beta_{M2} \cdot (R_t^{M2} - 50) + \beta_S \cdot (R_t^S - 50) + \beta_\sigma \cdot (R_t^\sigma - 50) \quad (18)$$

$$\sigma_{RSI|R} = \sigma_{RSI} \cdot \left(1 + \gamma \cdot \frac{R_t^\sigma - 50}{100} \right) \quad (19)$$

where μ_{RSI} and σ_{RSI} are unconditional RSI statistics, and $\beta_{M2}, \beta_S, \beta_\sigma, \gamma$ are estimated coefficients.

Point-in-time discipline: All conditional moments (μ_{RSI} , σ_{RSI} , and the β/γ coefficients) must be estimated on expanding windows using only data available at time t . The same release-lag discipline applied to M2 data (Section 9) applies to the indicator's own statistics: a Z-score computed against full-sample moments embeds lookahead bias even if the underlying inputs are point-in-time clean.

7.3.1 Bounded-Support Caveat and Distribution-Free Alternative

RSI has compact support $[0, 100]$, and in strong trends its distribution becomes markedly skewed with mass accumulating near the bounds. The Gaussian Z-score is therefore an approximation that degrades precisely where signals matter most: in the tails. Two remedies are available.

(a) Logit transform. Apply conditional standardization to the logit of the normalized RSI,

$$\ell_t = \ln \left(\frac{RSI_t}{100 - RSI_t} \right), \quad (20)$$

which maps $[0, 100]$ to \mathbb{R} and substantially improves the adequacy of the Gaussian approximation near the bounds. Conditional moments are then estimated for ℓ_t rather than RSI_t .

(b) Conditional empirical percentile (most robust). Replace the Z-score entirely with the regime-conditional empirical CDF:

$$Q_t^{RSI} = \widehat{F}_{RSI|R}(RSI_t), \quad (21)$$

estimated point-in-time within the active regime bin, with signals at $Q_t^{RSI} > 0.95$ (overbought) and $Q_t^{RSI} < 0.05$ (oversold). This is distribution-free, requires no moment assumptions, and admits a conformal-prediction interpretation: under exchangeability of observations within a regime, the signal thresholds carry finite-sample marginal coverage guarantees. Its cost is data hunger—the minimum-sample requirement of 63 observations per regime bin (Table 2) is binding here, and bins should be merged when violated.

We recommend the conditional percentile (b) as the primary signal channel where regime bins are well populated, with the Z-score (optionally logit-transformed) as the fallback for sparse regimes.

Signal generation (Z-score channel):

- $Z_t^{RSI} > 2.0$: Strongly overbought relative to current conditions
- $Z_t^{RSI} > 1.0$: Moderately overbought
- $-1.0 < Z_t^{RSI} < 1.0$: Normal for current regime
- $Z_t^{RSI} < -1.0$: Moderately oversold
- $Z_t^{RSI} < -2.0$: Strongly oversold relative to current conditions

Why we recommend conditional normalization: The approach:

1. Preserves the original RSI for direct inspection
2. Provides a statistically grounded measure of extremity
3. Automatically adapts to regime-specific RSI distributions
4. Produces standardized signals comparable across assets

7.4 Volatility Integration

Market volatility affects RSI reliability: high volatility produces more extreme readings that may not represent true overbought/oversold conditions. We compute annualized realized volatility:

$$\sigma_t = \sqrt{\frac{252}{n} \sum_{i=1}^n (r_{t-i} - \bar{r})^2} \quad (22)$$

where $r_t = \ln(\text{Close}_t / \text{Close}_{t-1})$ are log returns and n is the lookback (typically 20 days).

The volatility regime percentile is computed over a rolling 1-year window,

$$R_t^\sigma = \text{Percentile}(\sigma_t, \text{lookback} = 252 \text{ days}), \quad (23)$$

and enters both the threshold adjustment and conditional distribution estimation.

7.5 Combined Signal Generation

The two channels (conditional normalization and dynamic thresholds) can disagree: in a strong-liquidity regime the dynamic threshold rises while the conditional distribution shifts, and a raw RSI of 72 may be “normal” by one channel yet a threshold crossing by the other. Rather than an undifferentiated OR-rule—which lets the laxer channel dominate—we adopt a tiered scheme: a **strong** signal requires agreement of both channels, a **weak** signal fires when only one does. Position sizing and confirmation requirements should differentiate accordingly.

Algorithm 1 Context-Aware RSI Signal Generation

Require: Price series, M2 data, sentiment data

- 1: Compute standard RSI_t using Wilder’s formula
 - 2: Compute regime percentiles: R_t^{M2} , R_t^S , R_t^σ (point-in-time)
 - 3: Calculate conditional statistics: $\mu_{RSI|R}$, $\sigma_{RSI|R}$ (expanding window)
 - 4: Compute extremity measure: Z_t^{RSI} or, preferably, Q_t^{RSI} (Eq. 21)
 - 5: Calculate dynamic thresholds θ_t^{OB} , θ_t^{OS} and apply clipping (Eq. 16)
 - 6: **Output:** RSI_t , Z_t^{RSI} (or Q_t^{RSI}), θ_t^{OB} , θ_t^{OS}
 - 7: Let $C_1 \leftarrow (Z_t^{RSI} > 2 \text{ or } Q_t^{RSI} > 0.95)$; $C_2 \leftarrow (RSI_t > \theta_t^{OB})$
 - 8: Let $D_1 \leftarrow (Z_t^{RSI} < -2 \text{ or } Q_t^{RSI} < 0.05)$; $D_2 \leftarrow (RSI_t < \theta_t^{OS})$
 - 9: **if** C_1 **and** C_2 **then**
 - 10: Strong overbought signal
 - 11: **else if** C_1 **or** C_2 **then**
 - 12: Weak overbought signal
 - 13: **else if** D_1 **and** D_2 **then**
 - 14: Strong oversold signal
 - 15: **else if** D_1 **or** D_2 **then**
 - 16: Weak oversold signal
 - 17: **else**
 - 18: Neutral
 - 19: **end if**
-

8 Calibration and Regularization

8.1 Parameter Space

The framework requires estimation of the following parameters:

8.2 Objective Function

Parameters are optimized to maximize regularized risk-adjusted returns:



Parameter	Description	Typical Range
α_{M2}	M2 threshold sensitivity	[0.05, 0.20]
α_S	Sentiment threshold sensitivity	[0.05, 0.15]
α_σ	Volatility threshold sensitivity	[0.05, 0.15]
β_{M2}	M2 effect on conditional mean	[0.05, 0.20]
β_S	Sentiment effect on conditional mean	[0.05, 0.15]
β_σ	Volatility effect on conditional mean	[-0.10, 0.10]
γ	Volatility effect on conditional std	[0.0, 0.5]
λ	Sentiment recency decay	[0.05, 0.20]

Table 5: Calibration parameters and typical ranges

$$\hat{\theta} = \arg \max_{\theta} \left[\frac{E[R_p(\theta)] - R_f}{\sigma_p(\theta)} - \lambda_{reg} \|\theta\|_2^2 \right] \quad (24)$$

where:

- $R_p(\theta)$ is the strategy return under parameter vector θ
- R_f is the risk-free rate
- $\sigma_p(\theta)$ is strategy volatility
- λ_{reg} is the regularization strength

Recommended: $\lambda_{reg} = 0.01$, which penalizes extreme parameter values without overly constraining the optimization.

8.3 Overfitting Mitigation Strategies

Given the multi-dimensional parameter space, overfitting is a significant concern. We employ several safeguards. The two evaluation procedures below serve distinct roles and should not be conflated: nested cross-validation is the *design-stage* tool for honest comparison of specifications, while walk-forward optimization is the *deployment simulation* that estimates realistically attainable performance.

8.3.1 Nested Cross-Validation (Design Stage)

1. **Outer loop:** 5-fold time-series cross-validation for performance estimation
2. **Inner loop:** 3-fold cross-validation within each training fold for hyperparameter selection

This ensures that reported performance is estimated on data not used for parameter selection, and is the appropriate procedure for comparing design alternatives (e.g., Z-score vs. conditional percentile, raw vs. orthogonalized sentiment).

8.3.2 Walk-Forward Optimization (Deployment Simulation)

- **Training window:** Rolling 3-year (756 trading days)
- **Test window:** 1-year (252 trading days)
- **Re-estimation frequency:** Quarterly (63 trading days)

Parameters are re-optimized at each step using only past data, simulating realistic implementation conditions. Walk-forward results—not nested-CV results—are the headline numbers to report.

8.3.3 Parameter Count Constraint

To limit model complexity, we impose a maximum of 5 free parameters in any single specification. A recommended reduction: tie the threshold sensitivities ($\alpha_{M2} = \alpha_S = \alpha_\sigma =: \alpha$), fix $\lambda = 0.1$ at its theoretically motivated value, leaving $\{\alpha, \beta_{M2}, \beta_S, \beta_\sigma, \gamma\}$ free.

8.3.4 Multiple Testing Correction

When evaluating across multiple assets, apply Bonferroni or Benjamini-Hochberg FDR correction to p-values. For n assets tested, the adjusted significance threshold becomes α/n (Bonferroni) or uses the BH procedure for FDR control.

8.4 Re-estimation Protocol

Component	Update Frequency	Rationale
Regime statistics ($\mu_{RSI R}, \sigma_{RSI R}$)	Monthly	Balance stability and adaptiveness
Sensitivity parameters (α, β)	Quarterly	Allow adaptation to changing dynamics
M2 weights (w_i)	Annually	GDP shares change slowly
Sentiment decay (λ)	Semi-annually	Information dynamics relatively stable

Table 6: Recommended re-estimation schedule

9 M2 Nowcasting for Real-Time Applications

9.1 The Data Lag Problem

Official M2 data from central banks is typically released with a 2-4 week lag, creating challenges for real-time trading applications. For example, US M2 for month t is usually released in the third week of month $t + 1$.



9.2 Nowcasting Model

We estimate current US M2 using observable high-frequency proxies:

$$\widehat{M2}_t^{US} = \beta_0 + \sum_{i=1}^k \beta_i X_{i,t} + \varepsilon_t \quad (25)$$

Estimation approach:

- **Static OLS:** Estimate coefficients on a rolling 2-year window. Simple but may miss structural changes.
- **Kalman filter:** Allow dynamic coefficient evolution. More adaptive but requires state-space specification.

Model validation: Assess out-of-sample nowcast accuracy using RMSE and directional accuracy metrics. The nowcast should correctly predict M2 growth direction at least 70% of the time to add value.

9.3 Interpolation for Daily Values

Since M2 is reported monthly, we need daily values for indicator calculation. Options include:

- **Step function:** Hold monthly value constant until next release (simplest)
- **Linear interpolation:** Smooth transition between monthly values
- **Spline interpolation:** Smoother curves but risk of overfitting

For regime classification, the choice matters less than for level estimation, as percentile ranks are relatively robust to interpolation method. Note that linear and spline interpolation between a release and the *subsequent* release implicitly use future data; in backtests, interpolation must only span already-released values (i.e., interpolate between releases $t - 2$ and $t - 1$ while holding the segment after $t - 1$ as a step until release t arrives).

10 Implementation Architecture

10.1 Data Pipeline

10.2 Data Requirements

- **Price data:** Daily OHLC, minimum 252-day history for regime estimation, 1260 days preferred for full M2 regime lookback



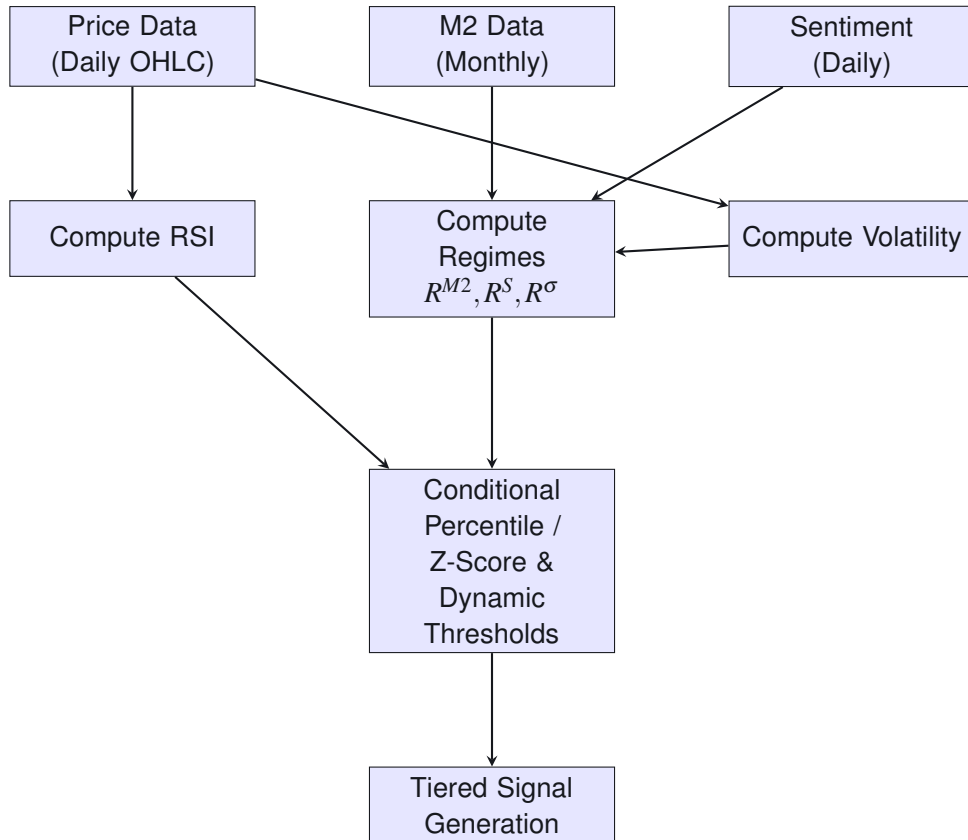


Figure 1: Data pipeline architecture

- **M2 data:** Monthly releases from FRED (US), ECB (Eurozone), PBoC (China), BoJ (Japan), BoE (UK)
- **Sentiment data:** Daily aggregates from news APIs (NewsAPI, Polygon) or commercial providers (RavenPack, Bloomberg)
- **Exchange rates:** Daily rates for level-aggregate conversion (available from FRED or central banks); not required for the growth aggregate (Eq. 9)

10.3 Computational Considerations

The indicator calculation is computationally lightweight:

- RSI calculation: $O(n)$ where n is the lookback period
- Regime percentile: $O(T \log T)$ for sorting T historical values (can be maintained incrementally)
- Conditional statistics: $O(1)$ lookup after initial computation
- Total per-bar computation: Sub-millisecond on modern hardware

The main computational cost is in the initial historical calculation and parameter optimization, not real-time signal generation.

11 Practical Trading Applications



11.1 Signal Integration

The context-aware RSI can be integrated into trading systems in several ways:

Standalone signals: Use strong-tier signals (both channels agreeing; Algorithm 1) as direct entry/exit signals, with weak-tier signals requiring confirmation from an independent indicator.

Filter for other strategies: Require RSI confirmation (e.g., not overbought per dynamic threshold) before taking signals from other indicators.

Position sizing: Scale position size inversely with $|Z_t^{RSI}|$ (or with $|Q_t^{RSI} - 0.5|$)—larger positions when RSI is in normal range, smaller at extremes.

Regime overlay: Use M2/sentiment regime classification to select among multiple strategy variants.

11.2 Risk Management Considerations

- **Stop losses:** Consider wider stops during high-volatility regimes when RSI readings are naturally more extreme
- **Position limits:** Reduce exposure when M2 is contracting, as liquidity-driven drawdowns can be severe
- **Correlation awareness:** M2 affects multiple assets simultaneously; diversification benefits may be reduced during liquidity-driven moves

11.3 Backtesting Considerations

When backtesting this framework:

1. **Avoid lookahead bias:** Use only data available at each historical point—for M2 releases, for interpolated segments, *and* for the indicator's own conditional statistics (Section 7.3.1)
2. **Account for transaction costs:** RSI signals can generate frequent trades; include realistic costs
3. **Test across regimes:** Ensure backtest period includes both M2 expansion and contraction phases
4. **Out-of-sample validation:** Reserve most recent data for final validation

12 Empirical Validation Framework



12.1 Prototype Tool

The prototype at <https://macropoiesis.com> implements this framework and allows users to:

1. Configure all sensitivity parameters (α , β , γ , λ)
2. Select from multiple assets and asset classes
3. Define custom backtesting periods
4. Compare enhanced RSI signals against baseline RSI
5. Export performance metrics and signal history

12.2 Recommended Validation Protocol

For rigorous evaluation, we recommend:

Asset universe:

- US Equity: S&P 500 ETF (SPY)
- Cryptocurrency: Bitcoin (BTC/USD)
- Commodity: Gold ETF (GLD)
- Emerging Markets: EM Equity ETF (EEM)
- Foreign Exchange: EUR/USD

Time period: 2015-2024, which includes:

- Multiple M2 expansion phases (2015-2016, 2020-2021)
- M2 contraction (2022-2023)
- Various volatility regimes

Parameterization: Use a single, pre-specified parameter set across all assets to avoid in-sample optimization bias.

Benchmarks and ablations: Performance must be reported against the following ladder, in order:

1. Buy-and-hold (passive baseline)
2. Classical RSI with static 70/30 thresholds

3. Classical RSI with static wide thresholds (80/20)
4. RSI with *unconditional* Z-score normalization (no regime conditioning)
5. Full framework, volatility conditioning only
6. Full framework, volatility + M2
7. Full framework, volatility + M2 + sentiment (raw and orthogonalized per Eq. 13)

Ablation step 4 is essential: without it, one cannot distinguish the value of regime *conditioning* from the value of mere standardization. Steps 5–7 attribute incremental value to each factor.

12.3 Performance Metrics

Metric	Formula / Description
Sharpe Ratio	$(E[R_p] - R_f) / \sigma_p$
Sortino Ratio	$(E[R_p] - R_f) / \sigma_{downside}$
Maximum Drawdown	$\max_t (\max_{s \leq t} P_s - P_t) / \max_{s \leq t} P_s$
Calmar Ratio	Annualized Return / Maximum Drawdown
Win Rate	Number of profitable trades / Total trades
Profit Factor	Gross profits / Gross losses
Average Trade	Mean P&L per trade

Table 7: Performance metrics for strategy evaluation

Statistical testing: Report Sharpe ratio differences with bootstrap confidence intervals (1000 replications). Apply Bonferroni correction when comparing across multiple assets.

13 Testable Hypotheses

Based on this design, we propose the following hypotheses for empirical validation. Operational definitions: an overbought (oversold) signal at time t is classified as **false** if the cumulative return over the subsequent 20 trading days is $\geq +2\%$ ($\leq -2\%$), i.e., the price continues in the signal-contradicting direction by more than a round-trip transaction-cost threshold. A **whipsaw** is a round-trip trade closed at a loss within 5 trading days of entry.

1. **H1 (False signal reduction):** Regime-conditional signals (percentile or Z-score channel) reduce false overbought signals during M2 expansion periods by at least 20% compared to the fixed 70 threshold, without proportionally increasing missed signals (defined as 20-day declines of $\leq -2\%$ not preceded by an overbought signal within the prior 10 days).
2. **H2 (Sentiment timing):** Incorporating sentiment improves signal timing around major news events (earnings releases, Fed announcements), as measured by average trade



P&L in event windows—and the improvement survives orthogonalization of sentiment against trailing returns (Eq. 13).

3. **H3 (Volatility adaptation):** Dynamic volatility thresholds reduce whipsaw trades by at least 15% during high-volatility regimes ($R_t^\sigma \geq 75$).
4. **H4 (Cross-asset heterogeneity):** Liquidity-sensitive assets (crypto, EM) show greater improvement in risk-adjusted returns from M2 conditioning than traditional developed market equities.
5. **H5 (Robustness):** Performance improvements persist out-of-sample and across different sub-periods, ruling out pure data mining; in particular, improvements over ablation step 4 (unconditional standardization) remain statistically significant after multiple-testing correction.

14 Limitations and Future Work



14.1 Current Limitations

1. **M2 data latency:** 2-4 week reporting delay limits real-time accuracy; nowcasting partially addresses this but introduces model error.
2. **Sentiment language bias:** English-only analysis may miss sentiment shifts in non-English markets (China, Japan, emerging markets).
3. **Sentiment endogeneity:** Even after orthogonalization (Eq. 13), residual sentiment may correlate with omitted price-derived variables; the channel's incremental value must be demonstrated, not assumed.
4. **Bounded-support approximation:** The Gaussian Z-score is an approximation on compact support; the conditional-percentile alternative mitigates this but is data-hungry in sparse regimes.
5. **Definitional heterogeneity in M2:** National aggregates differ in coverage (UK M4, ECB definitions, PBoC conventions); the growth convention reduces but does not eliminate the resulting measurement error.
6. **Parameter instability:** Optimal parameters may shift across market regimes; our re-estimation protocol mitigates but does not eliminate this concern.
7. **Overfitting risk:** Despite regularization and cross-validation, the multi-parameter nature of the framework creates overfitting potential.
8. **No empirical results:** This proposal presents the design specification; rigorous empirical validation is left to practitioners using the prototype tool.
9. **Single-indicator focus:** The framework enhances RSI specifically; integration with other indicators (MACD, Bollinger Bands) is not addressed.

14.2 Future Extensions

- **Separate liquidity factors:** Distinguish “US liquidity” from “global liquidity” and “local liquidity” for non-US assets
- **Credit conditions:** Incorporate high-yield spreads, bank lending standards, and financial conditions indices alongside M2
- **Multi-language sentiment:** Extend sentiment analysis to Chinese, Japanese, and European language sources
- **Regime-switching models:** Develop formal regime-switching specifications with automatic parameter adaptation
- **Conformal calibration:** Develop the conformal interpretation of the conditional-percentile signal (Section 7.3.1) into formal finite-sample validity statements, including adaptive regime-conditional coverage
- **Machine learning integration:** Use gradient boosting or neural networks to learn optimal threshold functions from data
- **Alternative indicators:** Apply the conditional normalization framework to other oscillators (Stochastic, CCI, Williams %R)

15 Conclusion



This design proposal presents a principled approach to enhancing the Relative Strength Index with macroeconomic and sentiment context. The key contributions are:

1. **Canonical design:** By preserving RSI in $[0, 100]$ and applying context through thresholds and normalization, we maintain interpretability while enabling adaptation.
2. **Explicit M2 construction:** The two-convention treatment—unweighted level aggregate for description, GDP-weighted growth aggregate for regime classification—provides a reproducible liquidity measure free of exchange-rate artifacts.
3. **Distribution-aware signals:** The conditional-percentile channel addresses RSI’s bounded support without parametric assumptions, with the (optionally logit-transformed) Z-score as fallback.
4. **Regularized calibration:** L2 penalization, nested cross-validation, walk-forward optimization, and a mandatory ablation ladder mitigate overfitting and attribute value to each component honestly.
5. **Operational clarity:** Specified re-estimation frequencies, data sources, point-in-time discipline, and validation protocols enable practical implementation.

6. **Working prototype:** The tool at <https://macropoiesis.com> allows immediate experimentation and validation.

The framework addresses a genuine limitation of traditional RSI—its ignorance of external market conditions—while respecting the indicator’s core strengths of simplicity and interpretability. We invite practitioners and researchers to validate the hypotheses using the prototype tool and welcome feedback for future refinements.

Acknowledgments

The author thanks reviewers for detailed feedback that substantially improved this proposal, particularly regarding the canonical design choice, the M2 aggregation conventions, and the regularization framework.

References

- [1] Araci, D. (2019). FinBERT: Financial Sentiment Analysis with Pre-trained Language Models. arXiv:1908.10063.
- [2] Baeyens, W. J. (2007). *RSI: Logic, Signals & Time Frame Correlation*. Greenville, SC: Traders Press.
- [3] Bernanke, B. S., & Blinder, A. S. (1992). The Federal Funds Rate and the Channels of Monetary Transmission. *American Economic Review*, 82(4), 901-921.
- [4] Bernanke, B. S., & Kuttner, K. N. (2005). What Explains the Stock Market’s Reaction to Federal Reserve Policy? *Journal of Finance*, 60(3), 1221-1257.
- [5] Du, K., Xing, F., Mao, R., & Cambria, E. (2024). Financial Sentiment Analysis: Techniques and Applications. *ACM Computing Surveys*, 56(9).
- [6] Engelberg, J., & Parsons, C. (2011). The Causal Impact of Media in Financial Markets. *Journal of Finance*, 66(1), 67-97.
- [7] Gerunov, A. (2025). Using Machine Learning on Macroeconomic, Technical, and Sentiment Indicators for Stock Market Forecasting. *Information*, 16(7), 584.
- [8] Hutto, C. J., & Gilbert, E. (2014). VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text. *Proceedings of the International AAAI Conference on Web and Social Media*, 8(1), 216-225.



- [9] Latif, S., Aslam, F., Ferreira, P., & Iqbal, S. (2025). Integrating Macroeconomic and Technical Indicators into Forecasting the Stock Market: A Data-Driven Approach. *Economies*, 13(1), 6.
- [10] Liu, H., Lin, Z., & Rojas, R. R. (2025). Enhancing Trading Performance Through Sentiment Analysis with Large Language Models. arXiv:2507.09739.
- [11] Loughran, T., & McDonald, B. (2011). When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks. *Journal of Finance*, 66(1), 35-65.
- [12] Mostafavi, S. M., & Hooman, A. R. (2025). Key Technical Indicators for Stock Market Prediction. *Machine Learning with Applications*, 20, 100631.
- [13] Murphy, J. J. (1999). *Technical Analysis of the Financial Markets*. New York: New York Institute of Finance.
- [14] Pícha, V. (2017). Effect of Money Supply on the Stock Market. *Acta Universitatis Agriculturae et Silviculturae Mendelianae Brunensis*, 65(2), 465-472.
- [15] Rigobon, R., & Sack, B. (2004). The Impact of Monetary Policy on Asset Prices. *Journal of Monetary Economics*, 51(8), 1553-1575.
- [16] Tetlock, P. C. (2007). Giving Content to Investor Sentiment: The Role of Media in the Stock Market. *Journal of Finance*, 62(3), 1139-1168.
- [17] Thorbecke, W. (1997). On Stock Market Returns and Monetary Policy. *Journal of Finance*, 52(2), 635-654.
- [18] Wilder, J. W. (1978). *New Concepts in Technical Trading Systems*. Greensboro, NC: Trend Research.
- [19] Yadav, A., Jha, C. K., Sharan, A., & Vaish, V. (2020). Sentiment Analysis of Financial News Using Unsupervised Approach. *Procedia Computer Science*, 167, 589-598.

Citation. Russmann, M. (2026) 'Context-aware RSI: integrating global M2 money supply and news sentiment'. Design proposal, version 2.0, June 2026.

Status. Design proposal circulated for discussion. Comments and corrections are welcome. · © 2026 Martin Russmann. All rights reserved.