The Signal — Live

Methodology, Governance & Calibration

Mehrzad Mahdavi, PhD · Physicist · 30 years quantitative finance experience

Founder, Digital Hub Insights · Version 1.0 · April 2026

Most retail signal services sell opaque picks. Customers do not see the model, the data discipline, the leakage controls, or whether this week's numbers survive a basic audit. The Signal — Live is built on the opposite premise: a signal service with published methodology and weekly calibration against a known reference is more valuable than one that hides its work.

How It Works

Weekly Up/Down equity signals across a 100-ticker liquidity-filtered U.S. equity universe, Long/Cash (default) or optional Long/Short, delivered Monday 07:00 ET. A single disciplined walk-forward pipeline produces both the historical backtest and each week's live picks — the same machine, applied at different times:

1. **Feature snapshot.** At Friday close, point-in-time features (momentum, volatility, cross-sectional) are computed, tagged by an immutable snapshot ID.
2. **Walk-forward retrain.** A Gradient Boosting classifier is re-fit on the most recent 156 weeks ending one embargo week before the prediction date.
3. **Leakage gates.** Four automated tests verify the feature matrix contains no future information and no target-adjacent columns before training proceeds.
4. **Rank and publish.** Model scores rank the universe; top-k are published as Up, bottom-k as Down (optional), remainder Cash. Each delivery carries an integrity block: model version, training window, embargo, gates passed.

Model class and hyperparameters are selected once, offline, via TimeSeriesSplit cross-validation. They are locked for production; the walk-forward re-fits a fresh instance of the locked specification every week.

Governance Principles

1. **Real data only** — no simulation shortcuts. Every pick is produced from real, adjusted market data.
2. **Audit before optimization** — the four leakage gates run before every weekly fit, not after. Impressive numbers that fail a gate never ship.
3. **Timing safety is non-negotiable** — the embargo between training and prediction is enforced programmatically, not by convention.
4. **Every decision is reversible** — every signal traces to an exact model version, snapshot ID, and universe. Anything can be forensically re-run.
5. **Calibrated weekly** — the runner must reproduce canonical benchmark numbers on the frozen 2019–2024 reference data before any week's picks go out. If it fails, the week is blocked.

Calibration: The Weekly Integrity Check

Every Monday run automatically reproduces our canonical backtest on the frozen 2019–2024 dataset. The values below are asserted within tight tolerance. If any metric drifts, the week's picks are not published and the founder is paged.

Strategy	CAGR	Sharpe	Volatility	Max Drawdown
SPY (Benchmark)	0.149	0.594	0.200	0.322
Gradient Boosting — L/C	0.154	0.717	0.174	0.226

Universe: 30 liquid U.S. equities (canonical reference). **Validation:** Walk-forward with one-week embargo. **Frequency:** Weekly. **Model:** Gradient Boosting.

These numbers are public. They are our weekly acceptance test — not a promise of future returns, but a commitment that today's pipeline is operating within the same integrity envelope as the audited research.

What Makes This Different

Typical Signal Service	The Signal — Live
Black-box recommendations	Published methodology, leakage gates, walkable code
Headline numbers only	Weekly calibration asserts canonical benchmark reproduces
No governance discipline	Four leakage gates run before every publication
Unversioned picks	Every signal traces to exact model + snapshot + universe
Editorial, opaque, or silent on drift	First-principles research; drift surfaced, not hidden
No education layer	Paired with a course that teaches every design decision

Code Template

Two of the four leakage gates that run as hard assertions before the model is fit each week. The full reference implementation is the Python template in the free Research Kit.

```
# Run BEFORE model.fit() on every weekly rebalance.

def gate_no_future_in_features(X, pred_date):
    """No feature column may reference data at or after pred_date."""
    latest = X.index.get_level_values('date').max()
    assert latest < pred_date, \
        f"Leakage: feature date {latest} >= pred_date {pred_date}"

def gate_time_split_integrity(train_end, pred_date, embargo_weeks=1):
    """Prediction date must be strictly after train_end + embargo."""
    gap_ok = (pred_date - train_end).days >= embargo_weeks * 7
    assert gap_ok, f"Embargo violated: {pred_date} vs {train_end}"
```

Availability

Pilot opens Monday, April 27, 2026. **First live picks** Monday, May 4, 2026. Weekly delivery, Monday 07:00 ET. **Subscribe:** cybersafeid.com/signal-live · **Course:** digitalhub-labs.com

We do not promise better returns. We promise we will show you the work.