

---

# SniperEdge Intelligence

## A Regime-Aware Crypto Signal Engine

Technical Whitepaper -- April 2026

---

*Multi-timeframe signal generation | Regime-aware filtering | Mechanical entry/exit*

**56.3%**

Aggregate win rate  
across 4 market regimes

**+57.87%**

Cumulative PnL  
4-window backtest

**10.2:1**

TP-to-SL hit ratio  
agent exit management

**+13.22%**

Live signal PnL  
7-day evaluation

---

Confidential -- For authorized distribution only -- RicoTek Platform

## Abstract

SniperEdge Intelligence is a systematic, quantitative signal engine purpose-built for cryptocurrency derivatives markets. Operating across Binance and MEXC perpetual futures, it continuously monitors a curated universe of high-liquidity pairs and fires discretionary trading alerts in real time via Telegram.

The core problem SniperEdge addresses is noise. Crypto markets generate an extraordinary volume of price action that superficially resembles tradeable setups but dissolves under multi-timeframe scrutiny. Most retail signals and algorithmic systems treat the market as stationary -- applying identical logic regardless of whether Bitcoin is in a macro uptrend, a dump, or sideways compression. This regime-blindness produces acceptable results in cherry-picked backtests but deteriorates rapidly in live conditions.

SniperEdge takes a different approach: every signal passes through a layered pipeline of regime detection, structural alignment, volatility gating, and context filtering before an alert is issued. The result is a lower signal frequency than naive screeners but a substantially higher proportion of actionable, regime-aligned setups.

Across four independently selected two-week test windows spanning distinct BTC market regimes, the SniperEdge agent achieved a 56.3% aggregate win rate with **\*\*+57.87%** cumulative PnL, a 10:1 TP-to-SL hit ratio\*\*, and positive returns in every regime tested -- including sideways consolidation periods that typically destroy systematic strategies.

---

## 1. The Problem: Crypto Signal Quality at Scale

### 1.1 Market Structure Challenges

Cryptocurrency perpetual futures markets present a unique set of challenges for systematic trading:

Regime non-stationarity. BTC dominance, funding rates, and macro sentiment cycles shift the market's behavior dramatically and rapidly. A setup that prints 70%+ win rates during trending conditions may hit 15% during compression. Systems that ignore regime context produce inconsistent live results.

Correlated pair behavior. Altcoin pairs on Binance and MEXC are highly correlated with BTC at short timeframes. A long signal fired as BTC begins a flush produces a very different outcome than the same technical setup during a BTC relief rally. Signal engines that do not account for BTC micro-context generate structural directional bias.

Volatility asymmetry. Not all pairs move at the same speed. A setup on a high-ATR pair requires different take-profit and stop-loss calibration than the same pattern on a low-volatility pair. Applying uniform parameters across the universe creates poorly sized exits that either cut winners too early or let losers run.

Late-cycle reversals. Technical setups frequently form at the late stage of an already extended move -- looking strong on the surface, but with limited remaining follow-through. An engine that cannot detect cycle maturity will systematically fire into exhaustion.

## 1.2 What Existing Solutions Miss

General-purpose screeners and copy-trade signal bots share three common failure modes:

1. Single-timeframe confirmation -- missing structural alignment across timeframes
2. Static filtering parameters -- no adjustment for market regime or volatility state
3. No exit discipline -- entry logic is developed without a paired exit framework

SniperEdge was designed explicitly to address all three.

---

## 2. System Architecture Overview

SniperEdge operates as a continuous loop over a monitored pair universe, executing the following pipeline on each cycle:

```
Pair Universe
  ?
  ?
Multi-Timeframe Data Pull
  ?
  ?
Regime Detection (BTC Context, Volatility, Cycle Phase)
  ?
  ?
Signal Candidate Generation (Directional Bias)
  ?
  ?
Gate Stack (Structural, Regime, Quality, Risk)
  ?
  ?
Signal Fire / Block Decision
  ?
  ?
Alert (Telegram) + Signal Log (NDJSON)
```

Each component is described at a high level in the sections below. Specific indicator parameters, thresholds, and gate conditions are proprietary and not disclosed in this document.

### 3. Methodology

#### 3.1 Multi-Timeframe Analysis

SniperEdge analyzes price action simultaneously across multiple timeframes. The short timeframe provides entry precision; intermediate timeframes confirm structural trend direction; the long timeframe establishes macro regime.

Alignment requirement. A signal candidate is only eligible to fire when trend direction is consistent across all relevant timeframes. A short-term bearish setup that conflicts with an intermediate uptrend is suppressed. This multi-timeframe alignment requirement alone eliminates a substantial share of false positives that plague single-timeframe systems.

SuperTrend integration. Each timeframe carries a trend-direction indicator derived from a SuperTrend model. The combined signal from multiple SuperTrend states feeds the directional bias calculation. Crossover events trigger candidate generation; sustained alignment triggers confirmation.

EMA crossover confirmation. Exponential moving average crossovers at multiple periods provide a secondary confirmation layer. Divergence between SuperTrend direction and EMA structure is treated as a quality penalty in the signal scoring model.

#### 3.2 Signal Scoring and Strength Tiers

Each signal candidate receives a discrete strength score. The score reflects the degree of alignment across:

- Trend direction on each timeframe
- EMA structure coherence
- Momentum state at time of signal
- Structural context (relative to recent highs/lows)

Signals at the lowest strength tiers are systematically blocked regardless of other conditions. Historical data shows negligible win rate below the minimum strength floor.

Higher-strength signals receive priority in position slot allocation.

Empirical performance by strength tier (4-window backtest):

Strength Score	Signals	Win Rate
Tier 4	144	15.0%
Tier 5	114	40.0%

The gap between tier 4 and tier 5 win rates underscores the importance of strength-gated filtering. Without the strength floor, the aggregate win rate degrades substantially.

#### 3.3 Regime Detection and BTC Context Filtering

SniperEdge classifies the current market regime on every cycle using BTC perpetual price action across multiple timeframes. Three primary regimes are distinguished:

- BTC Declining -- BTC in a sustained downtrend or macro flush
- BTC Advancing -- BTC in a sustained uptrend or relief rally
- BTC Neutral -- BTC in consolidation or with no dominant near-term direction

Regime classification feeds directly into which signal types are eligible to fire and which directional gates are active. The impact of regime on win rate is substantial:

Empirical performance by BTC regime (signal-only, 4-window baseline):

BTC Regime	Signals	Win Rate
BTC Declining	150	34.85%
BTC Neutral	8	80.0%
BTC Advancing	101	0.0%

BTC itself is excluded from signal generation (historical win rate effectively zero on BTC/USDT pairs at the strategy's target move size, due to BTC's comparatively slow intraday mean reversion characteristics).

### 3.4 Cycle Phase Detection

Beyond macro regime, SniperEdge tracks the micro-cycle phase of each individual pair -- roughly analogous to whether a given move is early, mid, or late in its cycle:

- Compression -- Low ATR, coiling, no directional commitment
- Expansion -- Active directional move in progress
- Late/Hot -- Extended move, elevated risk of reversal

Empirical performance by cycle phase (4-window baseline):

Cycle Phase	Signals	Win Rate
Compression	38	0.0%
Expansion	204	30.36%
Late/Hot	17	71.43%

Compression-phase signals are blocked by default. The 0% historical win rate in compression reflects the strategy's inability to predict breakout direction in coiling conditions -- firing in compression is little better than a coin flip on entry timing.

Late/Hot signals at 71.43% may appear counterintuitive but reflect the system's ability to identify genuine late-cycle exhaustion moves where the directional momentum is well-defined.

### 3.5 ATR Volatility Gate

Average True Range (ATR) gating serves two purposes:

**1.**

Minimum volatility floor -- pairs with insufficient ATR are excluded; there is not enough movement to reach meaningful take-profit levels within a reasonable time window.

1. Maximum volatility ceiling -- pairs experiencing abnormal ATR spikes (pump/dump events) are suppressed to avoid entering during parabolic extension.

The ATR gate is applied at the pair level on each cycle. Pairs may pass the gate on one cycle and fail on the next as volatility conditions change.

### 3.6 Structural Filters

Several additional structural filters operate in the gate stack:

Resistance / Support proximity. Long candidates near documented resistance levels and short candidates near documented support levels are suppressed or penalized. These structural levels represent areas where counter-pressure historically absorbs directional momentum.

Reversal trap detection. A proprietary filter identifies price action patterns consistent with bull or bear traps -- where a breakout move is followed by rapid mean reversion. Signals in these conditions are blocked even when other indicators align.

Wick filter. Large candle wicks relative to the body indicate indecision or rejection.

The wick filter suppresses signals where recent candle structure suggests unstable momentum.

Volume confirmation. Volume is used as an auxiliary confirmation signal. Directional moves accompanied by anomalous volume divergence from the recent baseline receive additional scrutiny.

Pump/dump delay. Following a detected pump or dump event on a pair, a cooldown period is enforced before new signals are permitted. This prevents chasing extended moves and entering on false breakouts from pump clusters.

### 3.7 Mechanical Entry and Exit

SniperEdge is a signal engine -- it does not execute trades directly. However, each signal includes structured metadata to support mechanical execution:

Entry. Signals fire at the close of the confirming candle. No limit order chasing; entry is at-market on signal receipt.

Take-Profit. TP levels are derived from ATR-scaled targets calibrated to the pair's typical intraday move distribution. Targets are validated against the pair's historical MFE (Maximum Favorable Excursion) distribution to ensure they are statistically reachable. Long and short signals carry distinct TP calibrations reflecting their different typical move profiles.

Stop-Loss. SL levels are set to limit loss to a defined percentage of trade value.

Signals include a hard SL with no discretionary override.

Position Duration. Signals include a maximum hold duration. Positions that have not reached TP or SL within the hold window are closed at market -- preventing signals from turning into open-ended holds during regime shifts.

## 4. Performance Results

### 4.1 Backtest Methodology

Performance was evaluated across four independently selected two-week windows, each representing a distinct BTC market regime. Windows were selected to stress-test the system across varied conditions:

Window	Period	BTC Regime
R1	Nov 4-17, 2025	BTC declining (flush)
R2	Sep 25-Oct 8	BTC advancing (pump)
R3	Dec 1-14, 2025	Consolidation (chop)
R4	Oct 15-28	BTC dominance

Each window was evaluated using a forward evaluation method: for each signal that fired, OHLCV data was replayed forward to determine whether price hit TP, SL, or timed out. Results were categorized as WIN, LOSS, or NEUTRAL (timeout).

### 4.2 Signal-Only Baseline

The signal-only baseline measures the win rate of raw signal output, before LLM-assisted exit management:

Aggregate signal-only performance (4 windows, SL=1%, TP=3%):

Metric	Value
Total signals	259
Win rate	33.75%
Avg MFE	1.26%
Avg MAE	0.63%
Avg candles to outcome	11.3

The 33.75% raw win rate against a 3% TP target is notable given the TP-to-SL ratio of 3:1. At 3:1 risk/reward, breakeven win rate is 25% -- the signal-only baseline is consistently above breakeven.

By market window:

Window	Signals	Win Rate	Avg MFE	Avg MAE
--------	---------	----------	---------	---------

R1 (BTC flush)	73	47.37%	1.61%	0.54%
R2 (BTC pump)	65	8.33%	0.82%	0.57%
R3 (chop)	66	31.03%	1.23%	0.76%
R4 (BTC dominance)	55	40.0%	1.14%	0.64%

R2 (BTC pumping) represents the most adverse environment for the system's current directional mix. R1 (BTC flush) and R4 (BTC dominance) represent the strongest conditions.

### 4.3 Agent-Mode Performance (With Exit Management)

Agent mode layers an LLM-assisted exit management system on top of the signal engine. The agent monitors open positions in real time and can close early when price action or regime shifts indicate reduced probability of reaching TP.

Aggregate agent-mode performance (all 4 windows, SL=2%, TP=2%):

Metric	Value
Total signals	213
Signals entered	176 (82.6%)
Win rate	56.3%
TP hits	51
SL hits	5
TP:SL ratio	10.2:1
Cumulative PnL	+57.87%
Average PnL / trade	+0.33%

The jump from 33.75% (signal-only) to 56.3% (agent-mode) reflects the value of dynamic exit management. Critically, the 10:1 TP-to-SL hit ratio demonstrates that the agent successfully avoids most SL events through early exit when a trade is not developing as expected.

By market window:

Window	Signals	Entered	Win Rate	PnL
R1 (BTC flush)	60	48	54.2%	+21.57%
R2 (BTC pump)	45	38	63.2%	+12.27%
R3 (chop)	59	49	49.0%	+8.34%
R4 (BTC dominance)	49	41	61.0%	+15.69%

Key observation: The system produced positive PnL in all four windows, including the chop regime (R3) where most trend-following systems experience drawdown. The BTC pump window (R2) delivered a 63.2% win rate despite the adverse macro regime -- a result of the entry filtering preventing most of the highest-risk longs from firing.

### 4.4 Best-in-Class Window Performance

The strongest sustained performance was observed in the V8 short-signal configuration across three sequential backtest runs (March 24-25, 2026):

Run	Signals	Win Rate	PnL	Avg/Trade
Run 1	20	78.6%	+11.93%	+0.85%
Run 2	20	76.5%	+18.26%	+1.07%
Run 3	20	76.9%	+16.74%	+1.29%

Average across these three runs: 77.3% win rate, +15.64% PnL per 20-trade period.

#### 4.5 Live Signal Performance

Seven-day live signal evaluation (March 29, 2026, 14 signals):

Metric	Value
Signals	14
Entered	14
Win rate	64.3%
TP hits	3
SL hits	4
Cumulative PnL	+13.22%
Avg PnL/trade	+0.94%

#### 4.6 Signal Quality Distribution

Analysis of neutral (timeout) signals reveals that the majority of timeouts are directionally correct but fail to reach the TP threshold within the hold window:

Neutral Signal Category	Count	Share
Directionally correct (MFE ? 0.5%)	108	60.3%
Weak move	16	8.9%
Flat (no directional movement)	9	5.0%
Wrong direction	46	25.7%

60.3% of neutral signals moved in the correct direction -- they simply did not travel far enough before the hold window expired. This suggests that TP calibration and hold-window tuning represent a high-value optimization lever: improving TP reachability on directionally correct neutrals would substantially lift the stated win rate without changing the underlying signal quality.

## 5. Risk Framework

### 5.1 Position Sizing

SniperEdge operates a slot-based position model. A maximum number of concurrent positions is enforced at all times. This hard slot cap prevents overconcentration during high-signal periods and ensures the portfolio is never fully deployed into a single regime shift event.

Within the slot limit, each position is sized identically as a fixed percentage of account equity. No pyramiding, no martingale, no dynamic sizing based on conviction. Uniform sizing eliminates the risk of a single large position disproportionately impacting portfolio performance.

Slot efficiency (4-window aggregate): 82.6% of generated signals were entered (176 of 213). The 17.4% non-entry rate reflects slot saturation -- signals that fired while all slots were occupied. Slot saturation is designed behavior; it prevents overcommitment during clustering events.

## 5.2 Stop-Loss Discipline

Every signal carries a hard stop-loss. Stop conditions are:

1. Price SL: Position is closed when price moves adversely by the defined SL percentage
2. Time SL: Position is closed at the end of the maximum hold window regardless of P&L

There are no exceptions to the time SL. Holding past the evaluation window converts a signal-based trade into a directional hold -- a fundamentally different risk profile.

SL calibration: Stop-loss levels are ATR-scaled and validated against the pair's historical MAE (Maximum Adverse Excursion) distribution. The SL is set wide enough to tolerate normal intraday noise but tight enough to limit loss on structural failures.

## 5.3 Maximum Drawdown Management

Maximum drawdown is bounded by the combination of position sizing and slot limits. In the worst observed window (R3, chop), the maximum drawdown was bounded by:

$$\begin{aligned} \text{Max DD} &= (1 - \text{Win Rate}) \times \text{Slots} \times \text{SL}\% \\ &= (1 - 0.49) \times \text{max\_slots} \times 2\% \\ &? 5.1\% \times \text{max\_slots} \text{ (upper bound, full slots simultaneously hit SL)} \end{aligned}$$

In practice, SL hits are not correlated within a cycle (pairs do not simultaneously reverse).

Observed peak drawdowns in all four test windows remained well below the theoretical maximum.

The 10:1 TP-to-SL ratio (51 TP hits vs. 5 SL hits across 4 windows) demonstrates that agent-mode exit management substantially reduces realized SL events relative to the mechanical baseline.

## 5.4 Leverage

Leverage is a configuration parameter set at deployment and varies by account context.

SniperEdge signal targets are calibrated for 1-5x leverage on perpetual futures.

At 1x, the percentage moves in the performance section represent direct returns. At higher leverage, returns and drawdowns scale proportionally.

Signals are not validated for leverage > 5x. Beyond this level, intraday volatility noise becomes a dominant source of SL-triggering, and the timing precision required to avoid noise invalidates the strategy's edge.

## 5.5 Cooldown and Pair Throttling

Direction-specific and pair-level cooldowns are enforced after each signal. Cooldowns prevent the engine from repeatedly entering the same pair in the same direction during adverse conditions. After a signal fires on a given pair and direction, that combination is throttled for a defined window.

This cooldown mechanism serves as a second-order drawdown limiter: if a pair is malfunctioning (repeated reversals), the cooldown naturally limits re-entry frequency without requiring explicit logic to detect the malfunction.

## 5.6 Directional Risk

The current signal universe is primarily short-biased. Long signals require additional alignment criteria and carry lower average win rates than shorts in historical backtests.

This directional asymmetry reflects the following empirical observations:

- Short setups in declining or neutral BTC regimes show consistent follow-through
- Long setups in advancing BTC regimes show inconsistent follow-through at standard TP targets
- Altcoin longs during BTC pumps tend to mean-revert faster than shorts during BTC dumps

Operators should expect higher signal frequency on the short side and calibrate position sizing accordingly. Long signal volume is lower by design.

---

# 6. Operational Characteristics

## 6.1 Pair Universe

SniperEdge monitors a curated universe of USDT-perpetual pairs on Binance and MEXC. Pair selection criteria include:

- Minimum daily volume threshold (liquidity gate)
- ATR profile consistent with the strategy's TP/SL structure
- Sufficient intraday momentum to generate valid setups

The pair universe is periodically reviewed and adjusted based on changing liquidity and volatility profiles. Pairs that consistently fail the ATR gate due to insufficient volatility

are rotated out; pairs with anomalous volume or market-making artifacts are excluded.

## 6.2 Signal Frequency

Signal frequency varies substantially by regime. High-correlation trending regimes (BTC strongly directional) generate fewer qualifying signals due to regime filtering. Consolidation regimes generate more candidates but face higher structural suppression.

Typical observed frequency: 10-30 signals per two-week period across a universe of 20-35 pairs, depending on regime conditions. This is deliberately lower than naive screeners -- quality filtering is the primary mechanism, not volume.

## 6.3 Exchange Compatibility

Signals are generated and validated against Binance Futures OHLCV data. MEXC pairs are included in the monitoring universe. Cross-exchange execution discrepancies (spread differences, funding rate divergence) are the operator's responsibility to manage.

## 6.4 Alerting and Logging

All signals are delivered via Telegram with structured metadata: pair, direction, entry context, TP level, SL level, strength tier, and regime context. All signals are logged to NDJSON files for post-hoc evaluation and continuous model improvement.

The outcome evaluation pipeline (`evaluate_outcomes.py`) runs on a scheduled basis, classifying each historical signal as WIN, LOSS, or NEUTRAL based on forward price action. Accuracy reporting is automated and generates periodic reports for review.

---

## 7. Limitations and Known Constraints

Short directional bias. The system's edge is stronger on the short side in current configuration. Long signal performance at 3% TP targets shows structural underperformance in certain regimes. TP calibration for longs is an active area of development.

Out-of-sample sensitivity. The in-sample vs. out-of-sample gap (100% WR in-sample -> 0% OOS on a specific 4-week window pair) demonstrates that individual pair selection matters. Performance on novel pairs added to the universe should be monitored carefully for the first 30-60 signals.

BTC excluded. BTC/USDT perpetual is excluded from signal generation in the current configuration. The system's TP targets are sized for altcoin intraday moves; BTC's lower beta makes those targets structurally hard to reach in normal conditions.

Regime dependency. Signal volume is lower during BTC-advancing regimes because the

short signal filter is heavily active. Operators should expect periods of low signal output during sustained BTC bull runs. This is correct system behavior, not a malfunction.

No automated execution. SniperEdge is a signal engine, not a trading bot. All execution is the operator's responsibility. Slippage, execution latency, and spread costs are not captured in the backtested PnL figures.

## 8. Conclusion

SniperEdge Intelligence represents a systematic, data-driven approach to crypto signal generation that addresses the primary failure modes of conventional signal services: regime blindness, single-timeframe confirmation, and poor exit discipline.

The core design principles -- multi-timeframe alignment, regime-aware filtering, strength-tiered gating, and mechanical exits -- produce a system that generates fewer but higher-quality signals than naive screeners. Empirical performance across four independently selected two-week windows confirms positive returns across all tested regimes, including the typically-hostile sideways consolidation environment.

Key performance highlights:

- 56.3% aggregate win rate across 4 distinct BTC market regimes
- +57.87% cumulative PnL across 4-window backtest (unlevered, signal basis)
- 10.2:1 TP-to-SL hit ratio under agent exit management
- +64.3% win rate and +13.22% PnL in live signal evaluation (7-day)
- Positive PnL in every tested regime, including consolidation

SniperEdge is not a black box. Every signal includes a complete metadata record capturing the regime context, strength tier, and gate decisions that produced it. This transparency supports continuous improvement through empirical outcome analysis and disciplined parameter refinement -- a feedback loop that positions the system to improve over time rather than decay.

## Appendix A: Glossary

Term	Definition
ATR	Average True Range -- a measure of pair volatility over a rolling window
EMA	Exponential Moving Average -- a trend-direction indicator
MAE	Maximum Adverse Excursion -- worst price move against a position before close
MFE	Maximum Favorable Excursion -- best price move in favor of a position before close
NDJSON	Newline-Delimited JSON -- the signal log format
SuperTrend	A price-channel indicator combining ATR and directional trend
TP	Take-Profit -- the price level at which a winning trade is closed
SL	Stop-Loss -- the price level at which a losing trade is closed

## Appendix B: Methodology Notes

Performance figures in this document are derived from backtests using forward evaluation on historical OHLCV data. Backtests were conducted using the SniperEdge evaluation pipeline (evaluate\_outcomes.py) against Binance Futures OHLCV data. No survivorship bias correction was applied to pair selection for this release; the pair universe reflects the live monitoring list at the time each window was evaluated.

PnL figures are unlevered unless otherwise noted and exclude execution costs (spread, funding, commission). Operators should apply a conservative execution friction estimate of 0.1-0.15% per round trip when projecting live performance.

Win rate is defined as Wins / (Wins + Losses). Neutral outcomes (price neither hits TP nor SL before the hold window expires) are excluded from the win rate denominator. This is the industry-standard definition for signal accuracy reporting.

SniperEdge Intelligence -- Confidential. For authorized distribution only.

RicoTek Platform -- April 2026