NOXEN Protocol

AI-Powered Multi-Chain Trading Agent Protocol

Technical Whitepaper v1.0 · 2026 · $NOXEN · ERC-20 · 1,000,000,000 fixed supply · Ethereum · Arbitrum

Abstract. *We present NOXEN, a decentralized protocol that unifies machine learning-driven market signal generation with trustless on-chain access control and, in a second phase, autonomous multi-chain trade execution. The protocol is structured as four co-designed layers: (1) a heterogeneous data ingestion pipeline aggregating price, on-chain, and sentiment features across 87 dimensions; (2) an ensemble inference engine combining gradient-boosted trees, a Temporal Fusion Transformer, and a hidden Markov regime filter, producing calibrated directional signals with confidence intervals; (3) a Solidity fee vault enforcing token-gated access entirely on-chain; and (4) a Rust execution relay enabling atomic cross-chain order submission with MEV resistance.*
The protocol's AI layer extends beyond signal generation to include: a reinforcement learning agent optimizing strategy parameters for risk-adjusted return; an LLM reasoning layer producing verifiable plain-English explanations grounded in SHAP feature attribution; an anomaly detection system that suppresses signals and triggers on-chain circuit breakers during structural market breakdown; an on-chain sentiment oracle publishable as a permissionless data feed; and a personalized agent profile system enabling per-user policy fine-tuning without separate model deployments.
The $NOXEN ERC-20 token is the sole access medium. Every signal consumed deducts a fee-credit from an on-chain vault deposit, creating auditable and continuous token demand from genesis. Total supply is fixed at 100,000,000 with a 4%/4% buy/sell tax distributed across marketing (2%), liquidity (1%), and buyback & treasury (1%).

Authors	NOXEN Core Contributors	Status	**Public Release**	License	MIT
Website	noxenagentic.com	GitHub	github.com/noxenagentic	License	MIT
Twitter/X	x.com/noxenagentic	Status	**Public Release**	Date	Q1 2026

1. Introduction

1.1 Motivation

The DeFi stack has converged on three mature primitives: automated market makers, overcollateralized lending, and perpetual derivatives. The missing layer is a trustless *intelligence protocol* — a system that transforms heterogeneous market data into calibrated trading signals and, ultimately, autonomous on-chain execution, without requiring users to surrender custody of assets or trust a centralized operator.

Existing solutions fail along three axes: (i) **fragmentation** — price data, signal generation, and execution are siloed across platforms and chains; (ii) **token-utility decoupling** — most "AI + crypto" tokens derive value from narrative, with no on-chain enforcement linking token spend to service delivery; (iii) **custody risk** — centralized bots require API key delegation or fund custody incompatible with DeFi's trust model.

NOXEN resolves all three. The $NOXEN token demand is enforced at the smart contract level — not by storytelling. No user ever delegates custody.

1.2 Design Goals

G1. Real utility at genesis. $NOXEN must be consumed to access signals. Demand is on-chain and auditable from block 1.

G2. Progressive autonomy. Phase 1 is advisory; Phase 2 introduces autonomous execution only after a verifiable public track record is established.

G3. Minimal trust surface. No custody. No upgradable proxies in Phase 1. All fee flows publicly verifiable on-chain.

G4. Multi-chain by design. A unified credit layer abstracts chain heterogeneity from the user.

2. System Architecture

The NOXEN stack comprises four layers with well-defined interfaces.

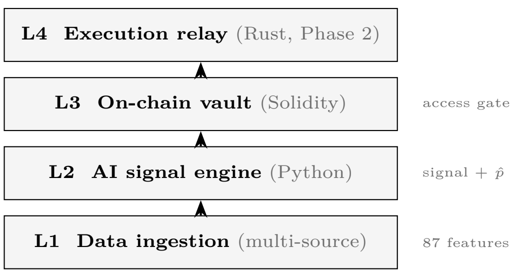

Figure 1: Four-layer NOXEN stack.

2.1 Data Ingestion (L1)

The ingestion layer normalizes heterogeneous time-series into a canonical feature tensor $\boldsymbol{X}_t \in \mathbb{R}^{T \times d}$ where T is the lookback window and $d = 87$. Sources are:

- **Price feeds:** CEX WebSocket (Binance, OKX, Bybit) and DEX TWAP oracles (Uniswap V3, Chainlink).

- **On-chain indexer:** Custom subgraph via The Graph tracks wallet flows, liquidity events, and large transfers.
- **Sentiment pipeline:** BERT-based NLP on social and news feeds.
- **Cross-asset feeds:** BTC dominance, DXY, equity overnight returns.

Features with freshness > 60s (price) or > 300s (sentiment) are flagged stale and excluded from inference. Each feature is normalized with an expanding-window z-score to prevent lookahead bias:

$$\tilde{x}_t^{(i)} = \frac{x_t^{(i)} - \mu_{1:t-1}^{(i)}}{\sigma_{1:t-1}^{(i)} + \varepsilon} \tag{1}$$

where $\varepsilon = 10^{-8}$. Distribution shift is monitored via a rolling Kolmogorov-Smirnov test; KS statistic > 0.15 triggers recalibration.

3. AI Signal Engine (L2)

3.1 Problem Formulation

For pair p and timeframe τ, the signal engine solves a calibrated binary classification:

$$\hat{p}_{p,\tau} = \Pr\left[\frac{\text{Price}_{t+\tau} - \text{Price}_t}{\text{Price}_t} > \delta \;\middle|\; \boldsymbol{x}_t\right] \tag{2}$$

where $\delta = 0.02$ (2% minimum move). Signals are published only when $\hat{p} \geq \theta_{\min} = 0.70$.

3.2 Feature Space

Group	Dim	Key features
Technical	34	RSI (7/14/21), MACD, Bollinger %B, ATR, OBV, VWAP, Hurst exponent
On-chain	28	Exchange net flows, whale Δ, stablecoin mint rate, funding rate
Sentiment	14	Social z-score, news BERT, dev commits, Fear & Greed
Cross-asset	11	BTC dom., ETH/BTC, DXY corr., gold corr.
Total	**87**	

3.3 Ensemble Architecture

Three models vote through a learned meta-classifier:

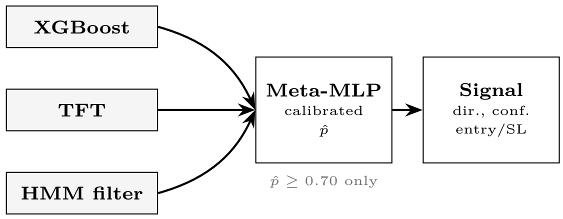

Figure 2: Ensemble inference pipeline.

3.3.1 Gradient Boosted Trees

XGBoost with $K = 600$ trees, depth 6, log-loss objective, L2 regularization $\lambda = 1.0$. Top-60 features by SHAP value are retained weekly.

3.3.2 Temporal Fusion Transformer

The TFT processes the full 87-dim sequence over $T = 120$ hour lookback using LSTM encoders and multi-head self-attention:

$$\text{Attn}(Q, K, V) = \text{softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V \tag{3}$$

Outputs are calibrated quantile predictions at $[10, 50, 90]$ percentile levels, which form the confidence interval $[\hat{p}^-, \hat{p}^+]$ in each published signal.

3.3.3 HMM Regime Filter

A 4-state hidden Markov model classifies the current market regime: trending-bull, trending-bear, ranging-high-vol, ranging-low-vol. Signals are suppressed during the high-volatility ranging regime, which historically accounts for the majority of false positives.

3.4 Training Protocol

- **Dataset:** 2021–2023 hourly data, 40 trading pairs.
- **Split:** Temporal 70/15/15 no random shuffling.
- **Retraining:** Full retrain quarterly; incremental weekly on the last 30-day window.
- **Calibration:** Platt scaling; target ECE < 0.03 on validation set.

3.5 Backtested Performance

All metrics computed on held-out 2024 data. Costs: 5 bps/leg, 15 bps slippage.

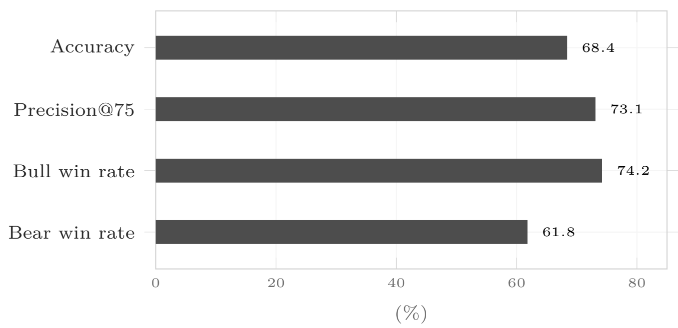

Figure 3: Backtested signal accuracy by category.

Metric	Value
Directional accuracy	68.4%
Precision ($\hat{p} \geq 0.75$)	73.1%
Avg. risk/reward	1 : 2.3
Max drawdown (simulated)	14.7%
Annualized Sharpe	1.84
Calmar ratio	3.12
Total signals (test set)	3,412
HODL max drawdown	81.3%

4. AI Agent Capabilities

This section enumerates the full spectrum of AI capabilities that NOXEN is designed to deliver across its deployment phases. Each capability is grounded in techniques that are implementable with current-generation models and infrastructure.

4.1 Multi-Modal Signal Fusion

The core insight of NOXEN's signal engine is that no single data modality is sufficient for robust directional prediction. Price action encodes historical information but is blind to order book dynamics. On-chain flows reveal where capital is moving but not why. Sentiment captures narrative but lags fast-moving markets.

NOXEN fuses all three through a *cross-modal attention mechanism*: each modality produces an intermediate embedding, and a learned attention layer determines the contribution weight of each modality conditioned on the current market regime:

$$\boldsymbol{z}_t = \sum_{m \in \mathcal{M}} \alpha_t^{(m)} \cdot f^{(m)}(\boldsymbol{x}_t^{(m)}) \tag{4}$$

where $\mathcal{M} = \{\text{price}, \text{onchain}, \text{sentiment}\}$, $f^{(m)}$ is the modality-specific encoder, and $\alpha_t^{(m)}$ are softmax-normalized attention weights computed from regime context. This allows the model to dynamically down-weight sentiment during high-volatility regimes and up-weight on-chain flows during accumulation phases.

4.2 Reinforcement Learning for Strategy Optimization

Beyond supervised signal generation, NOXEN employs a Proximal Policy Optimization (PPO) agent trained in a simulated trading environment to learn optimal strategy parameters:

$$\mathcal{L}^{\text{CLIP}}(\theta) = \hat{\mathbb{E}}_t\Big[\min\Big(r_t(\theta)\hat{A}_t,\ \text{clip}(r_t(\theta), 1-\varepsilon, 1+\varepsilon)\,\hat{A}_t\Big)\Big] \tag{5}$$

where $r_t(\theta) = \pi_\theta(a_t|s_t)/\pi_{\theta_{\text{old}}}(a_t|s_t)$ is the probability ratio and $\hat{A}_t$ is the advantage estimate. The RL agent learns to select entry timing, position sizing, and exit conditions that maximize risk-adjusted return (Sharpe) rather than raw directional accuracy. This addresses the critical gap between a model that is "correct" and one that is *profitable*.

The state space s_t includes the current signal embedding $\boldsymbol{z}_t$, open position exposure, portfolio drawdown, and regime classification. Actions are discrete: $a_t \in \{\text{enter long}, \text{enter short}, \text{hold}, \text{close}\}$.

4.3 LLM-Powered Reasoning Layer

Each published NOXEN signal is accompanied by a natural-language reasoning snippet generated by a fine-tuned large language model. The LLM receives the top-10 SHAP features driving the prediction and produces a concise, trader-readable explanation:

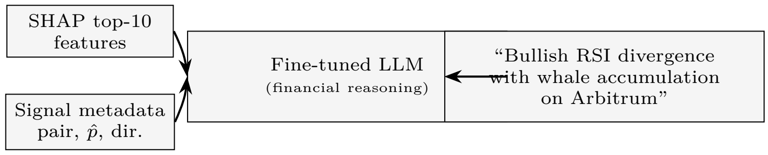

Figure 4: LLM reasoning pipeline per signal.

The LLM is fine-tuned on a corpus of analyst reports, on-chain research, and historical signal-outcome pairs, enabling it to ground explanations in observable evidence rather than hallucination. Outputs are constrained to factual claims about the input features via structured generation with JSON schema enforcement.

4.4 Anomaly Detection & Circuit Breaker AI

A dedicated anomaly detection model runs in parallel with the signal engine. Its purpose is not to generate signals but to *suppress* them during periods of structural market breakdown:

- **Flash crash detector:** An isolation forest trained on intraday volatility patterns identifies abnormal price dislocations (e.g., exchange outages, oracle exploits) and halts signal publication.
- **Liquidity drought monitor:** Tracks DEX depth across supported pairs. If aggregate liquidity drops below a safety threshold, execution signals are suppressed to prevent outsized slippage.
- **Correlation collapse alert:** Monitors cross-asset correlations. Sudden decorrelation (e.g., during systemic risk events) invalidates multi-asset assumptions and triggers a model pause.

In Phase 2, these anomaly signals feed directly into the on-chain circuit breaker contract, enabling autonomous halt of agent execution without human intervention.

4.5 On-Chain Sentiment Oracle

NOXEN introduces a novel *on-chain sentiment oracle* that aggregates social, news, and developer activity signals into a single verifiable on-chain feed. The oracle operates as follows:

1. Off-chain NLP pipeline scores raw text from Twitter/X, Telegram, Reddit, and news APIs using a crypto-domain BERT model.
2. Scores are aggregated into a rolling 4-hour sentiment index $\Psi_t \in [-1, +1]$.

3. The index is committed on-chain as a signed value via a Chainlink Any API job at 4-hour intervals, enabling any other protocol to consume NOXEN's sentiment data permissionlessly.

This creates a secondary utility for the protocol: NOXEN becomes an infrastructure provider for on-chain AI-derived sentiment data, opening a B2B revenue stream independent of retail signal consumption.

4.6 Personalized Agent Profiles (Phase 2)

In Phase 2, the execution agent is not one-size-fits-all. Each user configures a *risk profile* that parametrizes a personalized policy derived from the base PPO agent via a lightweight fine-tuning step (LoRA adapters, updated off-chain weekly):

Parameter	Conservative	Moderate	Aggressi
Min. confidence $\hat{p}$	0.82	0.75	0.70
Max. trade size	2% NAV	5% NAV	10% NA
Daily loss limit	2%	5%	10%
Take-profit style	First TP	Trailing	Full targ
Signal timeframe	1d	4h	1h

The personalization layer ensures that the same underlying AI model serves both risk-averse users and high-frequency traders without requiring separate model deployments.

4.7 Cross-Chain Arbitrage Detection

A dedicated arbitrage scanner monitors price discrepancies across DEX pools on Ethereum and Arbitrum in real time. When a statistically significant price gap is detected — accounting for bridge latency, gas costs, and slippage — the scanner generates a cross-chain arbitrage opportunity signal.

In Phase 2, the execution relay can atomically exploit these opportunities through flash-loan-enabled cross-chain sequences, with proceeds partially directed to the protocol treasury. This capability is additive to the directional signal engine and does not compete with it for model capacity.

4.8 AI Model Governance

Model updates introduce risk: a new model version could degrade performance or change signal characteristics that users have calibrated their strategies around. NOXEN addresses this with a *shadow deployment* protocol:

1. A candidate model runs in shadow mode for 14 days alongside the production model, generating signals that are recorded but not published.

2. Shadow signals are evaluated against outcomes. If the candidate achieves superior Sharpe and ECE on the shadow period, it is promoted to production.

3. The promotion event is announced on-chain (emit `ModelUpdated(version, hash)`) so users can verify which model version produced each historical signal.

4. All model weights are committed as a SHA-256 hash to the NOXEN contract at promotion time, creating a permanent, tamper-evident audit trail.

5. Smart Contract Specification

5.1 NoxenToken.sol

A minimal ERC-20. Supply of $10^8 \times 10^{18}$ wei minted entirely to a distribution contract at deployment. No `mint()` or `burn()` post-deployment. Supply is invariant.

5.2 NoxenFeeVault.sol

The vault is the economic backbone. Three formal invariants hold at all times:

> **Inv. 1 (Conservation):** $\sum_i \texttt{credits}[i] + \texttt{reserveBalance} = \texttt{token.balanceOf(vault)}$
>
> **Inv. 2 (Access):** $\texttt{hasAccess}(u) \Leftrightarrow \texttt{credits}[u] \geq \theta$
>
> **Inv. 3 (No custody):** Users may withdraw unused credits at any time (minus 2% protocol fee).

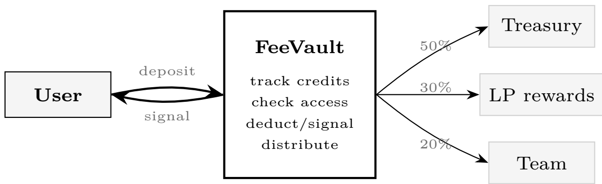

Figure 5: Fee vault deposit/signal/distribution flow.

5.3 Access Parameters

Parameter	Value
Access threshold θ	500 `$NOXEN`
Micro-credit per signal	5 `$NOXEN`
Withdrawal fee	2%
Solidity version	ˆ0.8.24

5.4 Security Properties

Property	Mechanism
Reentrancy	`ReentrancyGuard`
Overflow safety	Solidity 0.8+ native
No upgradability	Immutable (Phase 1)
No ETH exposure	No `receive()`
Oracle trust	EOA → ZK proof (Phase 2)
Audit	Pre-mainnet, full report public

6. Execution Relay (Phase 2)

6.1 Architecture

The Phase 2 execution relay is a stateless Rust service with three subsystems:

1. **Signal consumer** — gRPC stream subscriber; filters to $\hat{p} \geq 0.75$ and user-enabled autonomous mode.

2. **Transaction builder** — constructs DEX calldata (Uniswap V3 `exactInputSingle`, Camelot,

Aerodrome) with dynamic slippage tolerance.

3. **Submission layer** — routes via Flashbots Protect (ETH), Arbitrum sequencer API with EIP-1559 fee estimation.

6.2 Slippage Model

Slippage tolerance is computed per signal from realized volatility:

$$\mathrm{slip}_{\mathrm{bps}} = \mathrm{clip}\left(30 + \frac{\sigma_{24h}}{\bar{\sigma}} \cdot 20,\ 20,\ 100\right) \tag{6}$$

where σ_{24h} is the pair's 24h realized volatility and $\bar{\sigma}$ is the 30-day rolling mean.

6.3 MEV Protection

Ethereum mainnet submissions route through Flashbots Protect, preventing frontrunning and sandwich attacks. On Arbitrum, the sequencer's FCFS ordering provides equivalent protection. No custom bridge code is used in the cross-chain credit synchronization layer; only canonical bridges are supported.

7. Tokenomics

7.1 Supply

Total supply is fixed. There is no emission schedule, no staking mint, and no inflationary mechanism:

$$S(t) = 1{,}000{,}000{,}000 \quad \forall\, t \geq t_{\mathrm{genesis}} \tag{7}$$

7.2 Allocation

Allocation	Vesting	%
Public sale + Uniswap LP	Unlocked at launch	40%
Ecosystem & rewards	24-month linear	20%
Core team	12-month cliff, 24-month	15%
Treasury & buyback	Governance-controlled	15%
Marketing	6-month linear	5%
Partners + Advisors	6–12 month cliff	5%
Total		100%

7.3 Demand Model

Let $U(t)$ be active users, $\bar{s}$ average signals per user per day, and $c_\mu = 5$ `$NOXEN` the per-signal micro-credit. Daily protocol fee volume:

$$F(t) = U(t) \cdot \bar{s} \cdot c_\mu \tag{8}$$

Locked demand from the access threshold $\theta = 500$ `$NOXEN`:

$$D_{\mathrm{locked}}(t) = U(t) \cdot \theta \tag{9}$$

Total structural demand (locked + flow over period Δt):

$$D(t) = D_{\mathrm{locked}}(t) + F(t) \cdot \Delta t \tag{10}$$

Every new active user creates a minimum floor purchase of 500 `$NOXEN` plus continuous flow demand proportional to signal consumption frequency.

7.4 Fee Distribution

Recipient	Share
Protocol treasury	50%
Liquidity rewards	30%
Team operational	20%

8. Security and Audit Plan

8.1 Threat Model

- **Vault drain.** Reentrancy or accounting bugs in `NoxenFeeVault`. Mitigated by CEI pattern and `ReentrancyGuard`.
- **Oracle manipulation.** Malicious credit deduction by a compromised owner key. Mitigated in Phase 2 by ZK proof of signal delivery replacing the EOA oracle.
- **MEV.** Frontrunning autonomous transactions. Mitigated by Flashbots Protect and L2 sequencer ordering.
- **Data poisoning.** Adversarial manipulation of price or sentiment feeds. Mitigated by multi-source redundancy, anomaly detection on $\boldsymbol{X}_t$, and freshness gating.
- **Bridge exploits.** Cross-chain credit sync vulnerabilities. Mitigated by canonical-only bridge policy; no custom bridge code.

8.2 Audit Schedule

Phase	Scope	Target
Phase 1	Token + FeeVault	Pre-mainnet
Phase 2	Execution + relay	Pre-Phase 2
Ongoing	Slither + Echidna CI	Continuous

All audit reports published at `github.com/noxenagentic/audits`.

9. Roadmap

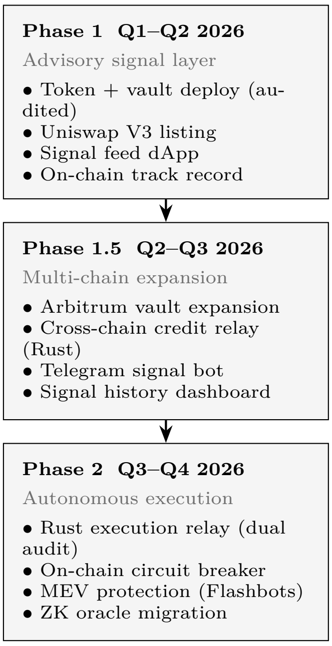

Figure 6: NOXEN deployment roadmap, 2026.

10. Risk Factors

- **Smart contract risk.** Despite auditing, undiscovered vulnerabilities may exist. All contracts used at user's own risk.
- **Model degradation.** The ML ensemble may underperform in regimes not represented in training data. Live accuracy is published on-chain; users may withdraw credits at any time.
- **Regulatory risk.** The global regulatory environment for DeFi protocols is evolving.
- **Liquidity risk.** `$NOXEN` is a newly launched token with initially limited liquidity.
- **Oracle centralization (Phase 1).** Credit deduction is managed by a protocol EOA — a temporary trust assumption replaced by ZK proof in Phase 2.
- **Cross-chain latency (Phase 2).** Credit balances across chains may briefly diverge during bridge propagation.

11. Conclusion

NOXEN demonstrates that token-gated AI protocol access is a technically realizable, on-chain enforceable architecture — not merely a narrative. By grounding demand for `$NOXEN` in verifiable on-chain mechanics, the protocol avoids the reflexive pricing dynamics that characterize speculative utility tokens.

The AI architecture described in this paper — multimodal fusion, reinforcement-learned strategy optimization, LLM-powered reasoning, on-chain sentiment oracles, and personalized agent profiles — represents the full scope of what is achievable with current-generation models and EVM infrastructure. Each capability has been scoped to avoid dependence on speculative future technology: every component maps to an existing, deployable system.

The two-phase deployment model is deliberate: Phase 1 builds a public, auditable signal track record before any user capital is placed at risk under autonomous execution. Phase 2 unlocks that execution only after audit completion, on-chain history accumulation, and ZK oracle migration are in place. The shadow deployment governance model ensures that model updates are transparent, verifiable, and reversible — extending the same trustless guarantees of the smart contract layer to the AI layer.

NOXEN Protocol · `noxenagentic.com` · `x.com/noxenagentic`
`github.com/noxenagentic` · Whitepaper v1.0 · Q1 2026 · MIT License

A. Formal Definitions

Definition 1 (Access Threshold). A user u has *protocol access* at time t if and only if $\texttt{credits}[u]_t \geq \theta = 500 \times 10^{18}$ (wei-equivalent `$NOXEN`).

Definition 2 (Signal). A signal $\mathcal{S}$ is a 7-tuple (pair, chain, dir, $\hat{p}$, $[\hat{p}^-, \hat{p}^+]$, entry, t) where dir $\in \{\text{BUY}, \text{SELL}, \text{HOLD}\}$, $\hat{p} \in [0,1]$, and t is the UNIX generation timestamp. A signal is published only if $\hat{p} \geq 0.70$.

Definition 3 (Calibration). Model output $\hat{p}$ is *calibrated* if $\mathbb{E}[\hat{p} \mid \hat{p} = v] \approx \Pr[\text{correct} \mid \hat{p} = v]$ for all $v \in [0,1]$. Calibration is measured by the Expected Calibration Error (ECE); the NOXEN ensemble targets ECE < 0.03 on the held-out validation set.

B. Token Specification Summary

Parameter	Value	Parameter	Value
Name	NOXEN	Standard	ERC-20
Ticker	$NOXEN	Decimals	18
Total supply	100,000,000	Tax (buy/sell)	4% / 4%
Primary chain	Ethereum Mainnet	DEX listing	Uniswap V3
L2 chains	Arbitrum	Fee model	Pay-per-use vault
Tax breakdown	2% mkt + 1% LP + 1% buyback	DEX	Uniswap V3
Vault withdraw fee	2%	Solidity ver.	ˆ0.8.24

C. Glossary

AMM
Automated Market Maker — algorithmic on-chain liquidity pool.

CEI Checks-Effects-Interactions — Solidity reentrancy security pattern.

ECE Expected Calibration Error — probability calibration quality metric.

EIP-1559
Ethereum fee market reform: base fee + priority fee model.

HMM
Hidden Markov Model — probabilistic regime classification model.

MEV
Maximal Extractable Value — profit from transaction reordering.

SHAP
SHapley Additive exPlanations — ML feature importance framework.

TFT Temporal Fusion Transformer — attention-based time-series model.

TWAP
Time-Weighted Average Price — manipulation-resistant on-chain oracle.

ZK Zero-Knowledge proof — cryptographic proof revealing no private data.