

Hamiltonian Alpha Engine

Market-state forecasts and option-book cashflows

Robert Jenkinson Álvarez

Technical summary

We start from the idea that the options market is not a collection of disconnected quotes or assets, but a coupled system of evolving surfaces, underlyings and other market variables. The option surfaces and their relationships are viewed as pricing objects and also high-dimensional observations of market-state. The geometry of the surfaces, constraints and their relationship with the rest of the market are used to construct a canonical phase-space state, which is then evolved using learned Hamiltonian dynamics. However, since these dynamics are deterministic descriptions of the structured evolution of the market state, we incorporate stochastic components to represent uncertainty around the evolutions. We therefore allow distributions of future states, rather than only point forecasts.

Selected forecasting results

Forecast target	Result	Historical evaluation
Relative surface state ^[1]	Innovation $R^2 = \mathbf{0.3654}$	2023; 41-observation horizon
Aggregate state ^[1]	Innovation $R^2 = \mathbf{0.2513}$	Same block and horizon
Physical index–basket covariance ^[2]	IC $\mathbf{0.5051}$; rank IC $\mathbf{0.4867}$	NDX; 20-day validation
Covariance innovation ^[2]	IC $\mathbf{0.4577}$; rank IC $\mathbf{0.3859}$	Same NDX benchmark
Quote-count state ^[1]	Newey–West $t = \mathbf{7.495}$	2021 annual evaluation
Half-spread state ^[1]	Newey–West $t = \mathbf{6.010}$	2023; five-day horizon

IC denotes forecast–outcome correlation; rank IC uses ranks. These are state and liquidity results, not returns. The 2023 figures are previously inspected development evidence. NDX validation has 40 sequences with overlapping horizons and retains index-replication residuals.

Surface history, 2019–2023^[1]

9,346

Surface-state records

1,174

Market dates

11.0 million

Quote-residual records

0

Surface-fit failures

Eight equities; frozen SANOS histories. Record counts are not independent observations or trades.

1 Data and research protocol

Over the 2019–2023 research window, the engine combines daily equity-market data, option surfaces, rates, VIX, credit information and supported alternative data. Option surfaces are obtained from the separate arbitrage-free surface engine which supplies surfaces, derivatives and fit-quality information. Write the information available at date t as

$$\mathcal{D}_t = \left\{ \begin{array}{l} \text{equities, option surfaces, rates,} \\ \text{VIX, credit, alternative data} \end{array} \right\}_{\leq t}. \quad (1)$$

Importantly, every feature used at date t is constructed only from information available by that date,

$$X_t = \Phi(\mathcal{D}_{\leq t}) \in \mathcal{F}_t. \quad (2)$$

Publication times, missing observations, identifier changes and corporate actions are therefore treated as part of the data and we do not let the historical backtest know information available at a later time.

Quantities learned from data are estimated from the training sample:

$$(\hat{\mu}, \hat{\Sigma}, \hat{M}, \hat{T}, \hat{\theta}, \hat{c}_{\text{trade}}) = \mathcal{A}(\mathcal{D}_{\text{train}}). \quad (3)$$

Feature dates and target dates are treated separately. For horizon h , an observation is admissible only when

$$t \leq T_{\text{feature, end}}, \quad L(t, h) \leq T_{\text{target, end}}, \quad (4)$$

with purging across split boundaries where future holding periods overlap. Executable forecasts use the timing

$$\begin{aligned} \text{signal at } t &\longrightarrow \text{entry at } t + 1 \\ &\longrightarrow \text{liquidation at } t + h + 1. \end{aligned} \quad (5)$$

So training data fit the model, validation data select specifications and trading rules, and evaluation data are used for reporting. Statistical inference is performed at the date or trade-cohort level. Candidate signals are compared against simple matched baselines and are subject to negative controls, multiple-testing correction, stability checks and the same execution-cost assumptions before entering a trading rule.

The separate option-surface engine supplies prices, derivatives and fit-quality information. Its call field is

$$C_{i,t}(m, \tau), \quad m = \log \frac{K}{F_{i,t}(\tau)}.$$

2 Hamiltonian market world model

The core model treats the market as a single coupled dynamical system. At date t , the observed market is represented by a canonical phase-space state

$$z_t = (q_t, p_t), \quad (6)$$

where q_t describes the current market configuration (position) and p_t its local direction of motion (momentum). The state contains information across assets and across the different market blocks.

The deterministic (non-stochastic) evolution is generated by a learned Hamiltonian,

$$\begin{aligned} \dot{q}_t &= \nabla_p H_\theta(z_t; c_t), \\ \dot{p}_t &= -\nabla_q H_\theta(z_t; c_t), \end{aligned} \quad (7)$$

where c_t denotes information used to condition the dynamics. We write

$$H_\theta(z; c) = H_0(z; c) + V_\theta(z; c), \quad (8)$$

with H_0 providing the structured reference dynamics and V_θ learning nonlinear residual structure. The Hamiltonian is defined on the joint state, so cross-asset and cross-market relationships are part of the dynamics.

The conservative evolution is computed through a structure-preserving map

$$\begin{aligned} z_{t+1} &= \Phi_{\theta, \Delta t}(z_t; c_t), \\ \Phi_{\theta, \Delta t}^* \omega &= \omega, \end{aligned} \quad (9)$$

so the learned transition retains the canonical phase-space geometry. Longer-horizon forecasts are obtained recursively from the same one-step law,

$$\widehat{z}_{t+h} = \Phi_{\theta, \Delta t}^h(z_t; c_t). \quad (10)$$

HamJEPA (NeurIPS2026 paper on the website) is used to learn this transition in representation space. It predicts the future market state generated by the Hamiltonian flow. In simplified form,

$$\begin{aligned} \widehat{z}_{t+h} &= P_\phi(z_t, c_t, h), \\ \mathcal{L}_{\text{JEPA}} &= \|\widehat{z}_{t+h} - \text{sg}(z_{t+h})\|^2 + \lambda \mathcal{R}(z), \end{aligned} \quad (11)$$

where sg denotes the stop-gradient target and $\mathcal{R}$ prevents representation collapse while respecting the phase-space structure. The output of this stage is therefore a forecast of the evolution of the joint market state. Uncertainty around that evolution is introduced separately in the stochastic layer.

3 Stochastic dynamics, memory and shocks

The Hamiltonian model gives a structured conditional evolution of the market state, but this single deterministic path is not sufficient for forecasting or trading. We therefore embed the Hamiltonian dynamics in a stochastic open-system model. In general form,

$$\begin{aligned} dz_t &= [X_{H_\theta}(z_t; c_t) + r_\psi(z_t; c_t)] dt \\ &+ \sum_{k=1}^M X_{G_k}(z_t; c_t) \circ dW_t^k + dJ_t, \end{aligned} \quad (12)$$

where X_{H_θ} is the learned Hamiltonian evolution, r_ψ represents dissipative or unresolved effects, the G_k generate continuous stochastic forcing and J_t contains discontinuous shocks.

The corresponding stochastic flow retains the canonical geometry pathwise,

$$\Phi_t^* \omega = \omega. \quad (13)$$

Dissipative and open-system components are kept separate and are not assumed to preserve Hamiltonian volume.

The observed state is only a partial description of the market. Unresolved variables can generate persistence and delayed response, which are represented through the finite-memory correction,

$$r_t^{\text{mem}} = - \int_0^t K_\psi(t-s) v(z_s) ds + \eta_t, \quad (14)$$

where K_ψ determines how previous state motion affects the present dynamics. This allows persistence and dissipation to arise from the dynamics rather than from an arbitrary collection of lagged features.

Large market moves are treated separately from continuous fluctuations. For a sequence of shock times

$$T_n = \sum_{j=1}^n \tau_j, \quad (15)$$

$$N_t = \max\{n : T_n \leq t\},$$

the shock time may depend on the elapsed time since the previous event, $a_t = t - T_{N_t}$, $\lambda_t = \lambda_\eta(a_t, c_t)$. This allows non-memoryless arrival dynamics. At a shock time the state changes according to

$$z_{T_n} = z_{T_n^-} + \mathcal{J}_\psi(z_{T_n^-}, c_{T_n^-}, \xi_n), \quad (16)$$

with the arrival process and the size and direction of the state displacement modelled separately.

The output of this layer is a conditional distribution over future market states,

$$\mathcal{P}_{t,h} = \mathcal{L}(z_{t+h} | \mathcal{F}_t), \quad (17)$$

from which quantities such as $\mu_{t,h} = \mathbb{E}[z_{t+h} | \mathcal{F}_t]$, $\Sigma_{t,h} = \text{Cov}(z_{t+h} | \mathcal{F}_t)$ can be obtained.

4 Regime and foliation geometry

A large displacement in the market state does not necessarily distinguish a temporary dislocation from a change in the underlying market regime. We introduce a low-dimensional regime map

$$\begin{aligned} \rho : \mathcal{P} &\rightarrow \mathbb{R}^k, \\ \rho(z) &= (\rho_1(z), \dots, \rho_k(z)), \end{aligned} \quad (18)$$

whose level sets $\mathcal{L}_c = \{z \in \mathcal{P} : \rho(z) = c\}$ define a foliation (theory explained in foliation paper) of the relevant region of phase space.

The market dynamics are separated locally into motion along the current regime and motion across regimes,

$$X_t = X_t^\parallel + X_t^\perp, \quad (19)$$

where $X_t^\parallel$ is tangent to the current leaf and $X_t^\perp$ changes the regime coordinates. Equivalently,

$$\dot{\rho}_t = D\rho(z_t) X_t. \quad (20)$$

For the slowly varying regime coordinates, the structured Hamiltonian dynamics are approximately tangential,

$$\{\rho_a, H_\theta\} \approx 0, \quad a = 1, \dots, k, \quad (21)$$

so ordinary state evolution occurs primarily within the current regime, while stronger transverse motions represent a change in the relationships defining that regime.

This provides a geometric distinction between two superficially similar market states:

$$\begin{aligned} X_t^\parallel &\rightarrow \text{within-regime dislocation,} \\ X_t^\perp &\rightarrow \text{regime transition.} \end{aligned} \quad (22)$$

The foliation therefore separates uncertainty about motion within the current market structure from uncertainty about whether that structure itself is changing. The trading layer uses this distinction when converting state forecasts into positions.

5 Contract-level trading and execution

The Hamiltonian model does not forecast trading signals. For a candidate book b consisting of fixed option contracts, let V_b denote the market value. Its expected net return over holding period h is evaluated

$$\hat{\mu}_{b,t,h}^{\text{net}} = \mathbb{E} \left[\Delta V_b^{\text{mid}} + G_b^{\text{hedge}} \mid \mathcal{F}_t \right] - \hat{C}_{b,t,h}, \quad (23)$$

where G_b^{hedge} denotes gross hedge cashflows and $\hat{C}_{b,t,h}$ includes estimated execution, hedging, impact and funding costs.

The direction and size of a position are determined using information available at the *signal date*. Candidate books must satisfy liquidity, capacity and risk constraints, and their expected net gains must exceed an allowance for forecast uncertainty:

$$\hat{\mu}_{b,t,h}^{\text{net}} > \mathcal{U}_{b,t,h}, \quad b \in \mathcal{B}_t^{\text{feasible}}. \quad (24)$$

When no candidate satisfies these conditions, the decision is not to trade. Trading follows a fixed chronological convention,

$$\begin{aligned} \text{signal at } t &\rightarrow \text{entry at } t+1 \\ &\rightarrow \text{exit at } t+h+1. \end{aligned} \quad (25)$$

Holding horizon h is specified prior to entry. Each position retains its original contracts, strikes and expiries throughout the holding period, with hedging cashflows accounted for separately. Realized performance is measured using the complete contract-level ledger,

$$\begin{aligned} \text{PnL}_b^{\text{net}} &= \Delta V_b^{\text{mid}} + G_b^{\text{hedge}} \\ &\quad - C_b^{\text{spread}} - C_b^{\text{impact}} \\ &\quad - C_b^{\text{hedge}} - C_b^{\text{funding}}. \end{aligned} \quad (26)$$

A predictable movement in the market state is only a trading opportunity if it can be converted into a feasible position with sufficient expected net return after costs and uncertainty. The term ΔV_b^{mid} measures the change in the value of the midpoint and does not mean that we trade at the midpoint. Trades buy at the *ask* and sell at the *bid*, with the resulting spread costs accounted in C_b^{spread} .

Recorded numerical and accounting checks^[3]

Conservative symplectic residual

2.98×10^{-14}

Maximum net cashflow discrepancy

5.23×10^{-12} dollars

Software regression tests

741 / 741 passed

The geometry check applies to the tested conservative step. The cash audit covers 2,185 complete admissions in overlapping alternative books.

Evidence. Selected internal historical results from different model versions, not a pooled track record or live performance. State forecasts and numerical checks do not establish validated net alpha.

Sources. [1] *Prequential Hamiltonian 2019–2023 Review*, pass five. [2] *Exact Canonical Alpha Program: 2022–2023 Review*, second pass. [3] *Response-Law Campaign Review*, 23 September 2026. [4] *Payoff-Jet Hamiltonian Program Review*, final stable pass. [5] *Forecast Engine Repair Review*, 18 September 2026.