

A Cointegrated Ising Spin Model for Asynchronously Traded Futures Contracts: Spread Trading with Crude Oil futures

Kostas Giannopoulos^{*1}

¹Dep. of Accounting and Finance, Neapolis University, Pafos, Cyprus 8042
`k.giannopoulos@nup.ac.cy`

13 October 2025

Abstract

Pairs trading via futures calendar spreads offers a robust market-neutral approach to exploiting transient mispricings, yet real-time implementation is hindered by asynchronous trading. This paper introduces a Cointegrated Ising Spin Model for real-time signal generation in high-frequency spread trading. The model links the macro-level equilibrium of cointegration with micro-level agent interactions, representing prices as magnetizations in an agent-based system. A novel Δ -weighted arbitrage force dynamically adjusts agents' corrective behavior to account for information staleness. Calibrated on tick-by-tick Brent crude oil futures, the model produces a time-varying probability of spread reversion, enabling probabilistic trading decisions. Backtesting demonstrates a 74.65% success rate, confirming the model's ability to generate stable, data-driven arbitrage signals in asynchronous environments.

The model bridges macro-level cointegration with micro-level agent interactions, representing prices as magnetizations within an agent-based Ising system. A novel feature is a **Δ -weighted arbitrage force**, where the corrective pressure applied by agents in response to the standard Error Correction Term is

^{*}The author would like to thank David H Bolton for programming support.

dynamically amplified based on information staleness. The model is calibrated on historical tick data and designed to operate in real time, continuously updating its probability-based trading signals as new quotes arrive.

1 Introduction

1.1 Statistical Arbitrage and Futures Spread Trading

Trading futures calendar spreads—the simultaneous purchase and sale of futures contracts on the same asset but with different delivery months—is a cornerstone of high-frequency statistical arbitrage (SA) [10]. Such strategies aim to profit from temporary deviations in the price relationship between contracts, relying on the statistical likelihood that these mispricings will revert to a historical or theoretical equilibrium [15]. Unlike pure arbitrage, SA strategies involve risk but seek to generate market-neutral returns by exploiting these predictable patterns, which are often identified through quantitative models [15, 18].

Futures contracts, due to their standardized nature and liquidity, are frequently employed in SA strategies. Inter-calendar spreads, in particular, are a popular strategy involving the simultaneous purchase of a futures contract for one delivery month and the sale of another futures contract for a different delivery month, on the same underlying asset [10]. These spreads can reflect market expectations about future supply and demand, storage costs (cost of carry), and convenience yields. Deviations in these spreads from their historical norms or fundamentally expected levels can present SA opportunities, assuming mean reversion or predictable patterns in the term structure. Spread trading strategies, particularly those based on pairs of liquid instruments, are widely employed across financial markets. These strategies rely on identifying temporary deviations in price relationships—often modeled via cointegration—that are expected to mean-revert over time. Pairs may consist of contracts on the same underlying asset with different maturities (calendar spreads), or different but closely related assets (inter-market spreads), such as futures, ETFs, or equities with fundamental or statistical linkages. The effectiveness of such strategies depends critically on the availability of high-frequency trading data, adequate market depth, and minimal execution delays.

Although the framework developed in this paper is general and applicable to any

such pair exhibiting liquid and asynchronous trading, we illustrate its implementation using the Brent crude oil futures of the first and second-month. These contracts are among the most actively traded globally, offering a natural testbed for evaluating high-frequency arbitrage strategies under real-world conditions.

2 Literature Review

The model presented in this paper draws upon several distinct but interconnected streams of financial and econophysics literature: statistical arbitrage focusing on pairs trading and cointegration, agent-based modeling using Ising spin systems, and the analysis of high-frequency market dynamics.

2.1 Statistical Arbitrage, Pairs Trading, and Cointegrated Models

A foundational concept in quantitative finance is the exploitation of temporary deviations from statistically identified equilibrium relationships between assets, a field formally known as Statistical Arbitrage (SA) [15]. The primary goal of SA is to construct market-neutral portfolios that profit from the expected correction of these mispricings, thereby isolating alpha from broader market movements. The most prominent implementation of this concept is pairs trading, a strategy involving a long position in an underpriced asset paired with a short position in a related, overpriced asset [9, 18, 10]. The success of this approach hinges on identifying assets whose prices exhibit a stable long-run relationship.

The theoretical underpinnings for identifying such relationships often stem from concepts like Arbitrage Pricing Theory (APT). APT suggests that if two securities share identical risk factor exposures, their expected returns over a given period should be the same [18]. Deviations from this parity, manifested in the spread between their (log) prices, are expected to be temporary and mean-reverting, forming the basis for pairs trading.

The econometric cornerstone for identifying mean-reverting spreads between non-stationary price series is the concept of cointegration, famously developed by Engle and Granger [7]. If two non-stationary time series, say the log-prices p_t^A and p_t^B , are

cointegrated, there exists a linear combination

$$S_t = p_t^A - \gamma p_t^B - \mu_z \quad (1)$$

(the spread, or cointegrating residual) that is stationary. The parameter γ is the cointegration coefficient, and μ_z (or an intercept in the cointegrating relation) represents the long-run equilibrium level of the spread.

The dynamics of cointegrated systems are often captured by a Vector Error Correction Model (VECM). A VECM describes how the individual series adjust when the spread deviates from its long-run equilibrium. A common representation of the VECM, including first-order lagged difference terms for short-run dynamics, is:

$$\Delta p_t^A = \alpha_A(p_{t-1}^A - \gamma p_{t-1}^B - \mu_z) + \phi_A \Delta p_{t-1}^A + \psi_A \Delta p_{t-1}^B + \epsilon_t^A \quad (2)$$

$$\Delta p_t^B = \alpha_B(p_{t-1}^A - \gamma p_{t-1}^B - \mu_z) + \phi_B \Delta p_{t-1}^A + \psi_B \Delta p_{t-1}^B + \epsilon_t^B \quad (3)$$

where α_A and α_B are the speeds of adjustment, $(p_{t-1}^A - \gamma p_{t-1}^B - \mu_z)$ is the error correction term (ECT) from the previous period, and $\phi_A, \psi_A, \phi_B, \psi_B$ are coefficients for the first-order lagged differences of Δp^A and Δp^B [18, 8]. More general VECM representations would include higher-order lags of these differences.

For the VECM system to be stable and mean-reverting, the adjustment speeds α_A and α_B must possess specific signs. With the ECT defined as $p_{t-1}^A - \gamma p_{t-1}^B - \mu_z$, stability requires $\alpha_A < 0$ and (for $\gamma > 0$) $\alpha_B > 0$. These signs ensure that a positive deviation (ECT > 0) is corrected by a decrease in p^A and/or an increase in p^B , and vice versa for a negative deviation. The model presented here adapts this error-correction principle, translating the statistical adjustment speeds (α_A, α_B) into a 'PullTerm' that directly influences agent behavior at the micro-level, as detailed in Section 3.

Trading strategies are then built by taking positions when the spread deviates significantly (e.g., by a threshold Δ) from its mean μ_z , expecting a reversion [18, 9]. The Stock-Watson common trends model provides another perspective, suggesting that cointegrated series share common underlying non-stationary components that are nullified in the cointegrating relationship [17, 18]. Advanced pairs trading frameworks may also incorporate tools like Kalman filters for dynamic hedge ratio estimation and other indicators like the Hurst exponent to assess mean-reversion persistence [10].

However, a critical limitation of the classic VECM framework (Eqs. 2 and 3) is the assumption that error correction is driven solely by the disequilibrium of the

immediately preceding period. In high-frequency markets, this assumption is often violated, as information lags and gradual price discovery can lead to delayed adjustments [10]. This issue of "delayed cointegration" has recently been explored from a continuous-time, mathematical finance perspective. For instance, Yan et al. [19] model the cointegrated price process as a path-dependent stochastic delay differential equation (SDDE), representing the limit of a high-order VECM. Their work focuses on deriving a mathematically optimal portfolio allocation by solving a system of Riccati partial differential equations.

While such analytical approaches provide a rigorous theoretical foundation for delayed adjustment, they are often difficult to implement and calibrate directly with discrete, asynchronous tick data. The present study addresses the same fundamental problem of delayed error correction but from a complementary, micro-founded perspective. Instead of modeling the continuous price path, this study uses an agent-based framework to model the discrete event-time dynamics. The model's key innovation—a Δ -weighted arbitrage force—explicitly incorporates the empirically observed time staleness of trades, offering a practical, data-driven mechanism to account for delayed information processing in a high-frequency environment.

2.2 Ising Models in Finance

The Ising model, originating from statistical mechanics to describe ferromagnetism, has found extensive application in finance and econophysics as a framework for modeling agent-based interactions and collective phenomena [6, 14]. The application of ferromagnet theory to social imitation, where individual agent decisions (spins) are influenced by neighbors and global fields, provides a theoretical basis for these models [4]. In this context, agents (traders) are represented by spins, typically binary ($S_i = \pm 1$), indicating buy/sell decisions or bullish/bearish sentiment. The aggregate state of the system is often measured by the total magnetization $M(t) = \frac{1}{N} \sum_i S_i(t)$, representing the overall market mood or trend [6, 2]. Callen and Shapero [4] discuss how concepts like order parameters (e.g., alignment of fish, phase angle of fireflies) and even a "social temperature" can be analogous to physical systems, further justifying the use of such physics-inspired models for collective human behavior.

Agent decisions (spin flips) in financial Ising models are governed by a local field $h_i(t)$. Bornholdt's influential work [2], for example, models this local field to include:

- **Neighbor Interactions (J_{ij}):** A ferromagnetic coupling ($J > 0$) encourages agents to align with their neighbors, representing herding behavior or imitation [6, 4].
- **Global Field/External Influence (α_{global}):** An external field, often related to the general state of the market (e.g. magnetization $|M(t)|$), can influence individual spins. In many financial Ising models, such as Bornholdt's, this term often induces an anti-ferromagnetic tendency (minority game characteristic), where agents are incentivized to take positions contrary to the majority if they believe profits lie in being contrarian [6, 2]. This component is crucial for generating complex dynamics such as "expectation bubbles" and intermittency. Note: this α_{global} is distinct from the ECM adjustment speeds α_A, α_B .
- **Idiosyncratic Preferences/Strategy Spin (C_i):** Some models introduce heterogeneous agent types, such as fundamentalists (who might believe in a "true" value) versus chartists or noise traders. This can be incorporated through strategy spins that modify how an agent reacts to the global field [6, 2].
- **Stochasticity/Temperature (β_{response}):** The probability of a spin flip is often a logistic function of the local field, $P(\text{flip}) \propto (1 + \exp(-2\beta_{\text{response}}h_i))^{-1}$, where β_{response} (inverse temperature) controls the randomness of agent decisions. High β_{response} (low temperature) implies more deterministic behavior based on the local field [6]. This can be seen as analogous to the concept of "social temperature", where higher randomness equals higher social temperature [4].

Ising models, particularly configurations like those proposed by Bornholdt [2] and further explored by others (e.g., Kaizoji et al. [12], Dvořák [6]), have been successful in replicating several stylized facts of financial markets. These include volatility clustering, heavy tails in return distributions, and the slow decay of autocorrelation in absolute returns [6, 5, 14]. Kukacka and Kristoufek [14] found that the Bornholdt (2001) model exhibited a strong tendency towards multifractal behavior, a complex dynamic feature observed in real financial data, arising from its inherent correlation structure. The "frustration" mechanism in such models, stemming from competing influences (e.g., herding vs. contrarian global field), is key to generating rich market dynamics.

3 A Probabilistic Agent-Based Model of Spread Reversion

The analytical framework presented herein bridges the gap between purely statistical time series models (like VECM for cointegration) and the micro-foundations of simulation-based Agent-Based Models. It proposes a hybrid framework where the perceived market disequilibrium—derived from a macro-level cointegrating relationship—directly influences micro-level agent decisions, creating a novel feedback loop for modeling asynchronously traded, cointegrated assets. Conventional econometric models, such as VECM, are inherently retrospective and assume synchronous price discovery. In contrast, the model proposed here is designed for live implementation, continuously ingesting market quotes to generate real-time spread-trading signals

3.1 Theoretical Foundation and Empirical Implementation

The formal specification of this framework requires distinguishing between the true underlying market sentiment and what is observable at any given moment, explicitly accounting for the asynchronicity of trades. In the theoretical model, for each contract $X \in \{A, B\}$, a **latent magnetization** M_t^X (latent) represents the true, unobserved aggregate sentiment of all agents. The **observed magnetization**, $M_t^{X,obs}$, only updates when a trade occurs, with the time elapsed since the last trade denoted by Δ_t^X .

A core challenge in high-frequency markets is that asset prices do not update simultaneously. To conceptualize how agents might internalize this, we can define a theoretical, **staleness-adjusted Error Correction Term (ECT')**, where the perceived spread is adjusted based on the age (staleness) of the last trade in each contract. This adjusted ECT' reflects a hypothetical perception of disequilibrium that discounts stale price observations. Conceptually, this could be calculated using the last observed market sentiment (magnetization) for each contract, weighted by its staleness:

$$ECT'_{t-1} = \left(M_{\tau_A(t-1)}^{A,obs} \cdot w(\Delta_{t-1}^A) \right) - \gamma \left(M_{\tau_B(t-1)}^{B,obs} \cdot w(\Delta_{t-1}^B) \right) - \mu_z \quad (4)$$

where $w(\Delta) = e^{-k\Delta \cdot \Delta}$ is an exponential weighting function modeling the decay of information relevance. While this ECT' provides the theoretical motivation, this study's practical implementation models this effect more directly by applying a Δ -weighting to the *arbitrage force* that agents exert in response to the standard, empirically observed

ECT, as detailed in Section 3.3.

To operationalize this framework for empirical application, we ground the model in the results of the preliminary VECM analysis. The observable spread at any time t , denoted S_t , is defined as the cointegrating residual, or the empirical **Error Correction Term (ECT)**, from the estimated long-run relationship:

$$S_t \equiv \text{ECT}_t = P_t^A - \gamma P_t^B - \mu_z \quad (5)$$

where γ and μ_z are the parameters for the cointegrating vector and the long-run mean, respectively, estimated from the VECM. **This empirically observed disequilibrium, S_t , serves as a direct and tractable proxy for the theoretical disequilibrium otherwise captured by the difference in latent magnetizations.** The model's core function is to compute the conditional probability that this spread will revert towards its mean. This is achieved by simulating the collective behavior of a heterogeneous agent population partitioned into two distinct behavioral archetypes: trend-followers and contrarians.

3.2 The Agent-Based Ising Spin Model

The evolution of the latent magnetization is driven by the collective decisions of individual agents, modeled as spin flips. The probability of an agent i adopting a positive spin (e.g., a bullish stance) for contract X at time t is governed by a logistic function of a "local field," $h_{i,t}^X$:

$$\text{Prob}(S_{i,t}^X = +1 | \mathcal{F}_{t-1}) = \frac{1}{1 + \exp(-2\beta_{\text{response}}^X h_{i,t-1}^X)} \quad (6)$$

The local field $h_{i,t-1}^X$ synthesizes the forces influencing an agent's decision at time t , based on information available at $t - 1$:

$$h_{i,t-1}^X = \underbrace{\sum_{j \in \text{Nb}^X(i)} J_{ij}^X S_{j,t-1}^X}_{\text{Herding Term}} - \underbrace{\alpha_{\text{global}}^X S_{i,t-1}^X |M_{\tau_X(t-1)}^X|}_{\text{Global Contrarian Term}} + \underbrace{\lambda^X \cdot \text{PullTerm}_{t-1}^X}_{\text{Cointegration Pull Term}} \quad (7)$$

The first two terms represent standard forces in financial Ising models: herding/imitation and a contrarian reaction to the market trend. The third term, the '**PullTerm**', is the novel component that directly links the agent's decision to the cointegrating relationship. It quantifies the arbitrage pressure derived from the per-

ceived disequilibrium, as defined in the empirical implementation by Equations (10) and (11).

3.3 Microfoundations of Agent Behavior and Local Field Specification

The model's dynamics are driven by the aggregate decisions of this heterogeneous agent population. Let Δ_t^A and Δ_t^B represent the time elapsed (staleness) since the last trade for contracts A and B, respectively. The decision calculus for each agent type is informed by a synthesis of momentum and mean-reversion signals.

Modeling Agent Influences. The forces acting upon the agents are designed to capture well-documented financial phenomena: herding, momentum-chasing, and arbitrage-driven error correction.

First, a **Herding/Momentum Influence** (I_t^{herd}) captures the tendency for social imitation, where agents' decisions are influenced by the perceived aggregate sentiment of the market. This is formalized by the simulated market sentiment from the previous period:

$$I_t^{\text{herd}} = M_{t-1}^{\text{sim}} = 2 \cdot P_{t-1}^{\text{market buy}} - 1 \quad (8)$$

This term creates a feedback loop where a tendency to buy (or sell) in one period increases the probability of similar behavior in the next.

Second, a pure **Follower Influence** (I_t^{follow}) represents the behavior of chartists or momentum traders who act on recent price changes. This force is responsible for amplifying spread deviations. Its strength is governed by the parameter J_{mom} :

$$I_t^{\text{follow}} = J_{\text{mom}} \cdot (S_{t-1} - S_{t-2}) \quad (9)$$

A positive and significant J_{mom} ensures the model can generate the positive feedback dynamics that temporarily drive prices away from equilibrium.

Third, the **Contrarian Influence (The 'PullTerm')** is the critical error-correction force that underpins the mean-reverting tendency of the spread. It represents the behavior of arbitrageurs who identify deviations from the long-run equilibrium and trade to correct them. This term is the practical implementation of the staleness-weighted arbitrage response motivated by the theoretical ECT' concept. It is a function of the

spread's deviation from its mean, scaled by the adjustment speed parameter, α . Critically, it incorporates information staleness, where the parameter k_Δ models the degree to which arbitrageurs amplify their reaction based on the age of price data:

$$I_t^{\text{contr, A}} = -\alpha \cdot (S_{t-1} - \mu_z) \cdot (1 + k_\Delta \cdot \Delta_t^A) \quad (10)$$

$$I_t^{\text{contr, B}} = -\alpha \cdot (S_{t-1} - \mu_z) \cdot (1 + k_\Delta \cdot \Delta_t^B) \quad (11)$$

The negative sign ensures that a positive spread deviation ($S_{t-1} > \mu_z$) generates a negative (sell) pressure. The term $(1 + k_\Delta \cdot \Delta_t^X)$ means that the corrective force is amplified by staleness.

Agent Decision Probabilities. These distinct informational influences are synthesized into a "local field", h_t , for each agent type, representing the net decision-making force.

$$h_t^{\text{follow}} = I_t^{\text{herd}} + I_t^{\text{follow}} \quad (12)$$

$$h_t^{\text{contr, A}} = I_t^{\text{herd}} + I_t^{\text{contr, A}} \quad (13)$$

$$h_t^{\text{contr, B}} = I_t^{\text{herd}} + I_t^{\text{contr, B}} \quad (14)$$

The translation of this deterministic field into a stochastic decision is governed by a logistic function. The parameter β_{response} , analogous to the inverse temperature in statistical mechanics, controls the rationality or noise level of the agents.

$$P_t^{\text{follow buy}} = \frac{1}{1 + \exp(-2\beta_{\text{response}}h_t^{\text{follow}})} \quad (15)$$

$$P_t^{\text{contr, A buy}} = \frac{1}{1 + \exp(-2\beta_{\text{response}}h_t^{\text{contr, A}})} \quad (16)$$

$$P_t^{\text{contr, B buy}} = \frac{1}{1 + \exp(-2\beta_{\text{response}}h_t^{\text{contr, B}})} \quad (17)$$

The average buy probability for the contrarian agent group is the arithmetic mean: $P_t^{\text{contr avg buy}} = (P_t^{\text{contr, A buy}} + P_t^{\text{contr, B buy}})/2$.

Market-Level Aggregation and Reversal Probability. The micro-level decisions are aggregated to yield a market-level signal, reflecting the market ecology. The parameter N_C represents the proportion of contrarian agents in the total population,

N_{total} .

$$P_t^{\text{market buy}} = \frac{N_C}{N_{total}} P_t^{\text{contr avg buy}} + \frac{N_{total} - N_C}{N_{total}} P_t^{\text{follow buy}} \quad (18)$$

Finally, this aggregate probability is translated into the model's ultimate prescriptive output: the conditional probability of mean reversion, $P_t(\text{reversal})$. If the spread is too high ($S_{t-1} > \mu_z$), reversion implies selling.

$$P_t(\text{reversal}) = \begin{cases} 1 - P_t^{\text{market buy}} & \text{if } S_{t-1} > \mu_z \\ P_t^{\text{market buy}} & \text{if } S_{t-1} < \mu_z \end{cases} \quad (19)$$

3.4 Parameter Estimation via Trading Simulation

The estimation of parameters for the proposed agent-based model presents a formidable challenge. The high-dimensional latent state space, representing the configuration of all agent spins, renders conventional methods such as Maximum Likelihood Estimation (MLE) computationally infeasible due to an intractable likelihood function. Similarly, while the Simulated Method of Moments (SMM) is a viable alternative for many Agent-Based Models, it focuses on matching a set of pre-selected statistical properties (moments) of the empirical data. Given that the primary goal of this research is to develop a framework for a profitable trading strategy, we adopt a more direct and performance-oriented estimation approach: "simulation-based optimization". Instead of matching statistical moments, this methodology defines an objective function based on the success rate of a trading strategy derived from the model itself. The model parameters are then optimized to maximize this performance metric over the historical dataset, directly aligning the parameter estimation process with the model's ultimate application. The model's free parameters, encapsulated in the vector $\Theta_{ABM} = \{J_{\text{mom}}, \alpha, k_{\Delta}, \beta_{\text{response}}, N_C, \theta_{\text{buy}}, \theta_{\text{sell}}\}$, are optimized by maximizing the historical performance of a trading strategy.

Trading Signal Generation. A trade signal is generated when the model's probabilistic output exceeds a calibrated threshold, indicating a high likelihood of mean reversion.

- **Sell Signal Generation:** A signal to sell the spread is generated at time t if $S_{t-1} > \mu_z$ and the reversal probability exceeds a sell threshold ($P_t(\text{reversal}) \geq \theta_{\text{sell}}$).

- **Buy Signal Generation:** A signal to buy the spread is generated at time t if $S_{t-1} < \mu_z$ and the reversal probability exceeds a buy threshold ($P_t(\text{reversal}) \geq \theta_{buy}$).

For the backtest, the trading strategy’s logic is defined by a sophisticated, dual-filter trigger mechanism that requires a specific alignment of statistical conviction with confirming market dynamics. A trade signal is generated only when a set of stringent conditions are met: for a sell signal, for example, the spread must be overvalued ($S_{t-1} > \mu_z$), the model must indicate a high probability of reversion ($P_t(\text{reversal}) \geq \theta_{sell}$), and recent positive momentum in the spread must validate the entry point. Buy signals are generated symmetrically. This dual requirement acts as a powerful safety mechanism, ensuring the strategy enters a trade precisely as a strong deviation shows signs of exhaustion, rather than attempting to trade against adverse momentum.

Defining a Successful Trade. A trade executed at time t is assessed for success based on the price movement in the subsequent tick, $t+1$. A trade is deemed successful if the spread moves in the predicted direction.

- A **sell trade** is successful if the spread decreases at $t+1$ (i.e., $\Delta S_{t+1} < 0$).
- A **buy trade** is successful if the spread increases at $t+1$ (i.e., $\Delta S_{t+1} > 0$).

Hybrid Objective Function for Optimization. Optimizing for the raw success rate is a direct approach, but it ignores the probabilistic confidence of each signal. A more robust methodology evaluates the model’s probabilistic accuracy, but only for those instances that trigger a trade. This ensures the optimization focuses on improving actionable forecasts. To balance signal quality with overall trading performance, we define a *hybrid objective function*.

First, we define the realized outcome of a trade signal issued at time t . Let R_t be a binary variable indicating if the spread reverted in the subsequent period, $t+1$:

$$R_t = \begin{cases} 1, & \text{if } (S_{t-1} > \mu_z \text{ and } \Delta S_{t+1} < 0) \text{ or } (S_{t-1} < \mu_z \text{ and } \Delta S_{t+1} > 0), \\ 0, & \text{otherwise.} \end{cases} \quad (20)$$

Next, we define the set of time steps where a trade signal is generated, denoted

$\mathcal{T}_{\text{trade}}$:

$$\mathcal{T}_{\text{trade}} = \{t \mid (S_{t-1} > \mu_z \wedge P_t(\text{reversal}) \geq \theta_{\text{sell}}) \vee (S_{t-1} < \mu_z \wedge P_t(\text{reversal}) \geq \theta_{\text{buy}})\} \quad (21)$$

The objective is to maximize the statistical consistency between the forecast $P_t(\text{reversal})$ and the outcome R_t for all $t \in \mathcal{T}_{\text{trade}}$. This is captured by a conditional Mean Squared Error (MSE) objective:

$$\mathcal{O}_{\text{C-MSE}}(\Theta) = 1 - \frac{1}{|\mathcal{T}_{\text{trade}}|} \sum_{t \in \mathcal{T}_{\text{trade}}} (P_t(\text{reversal}) - R_t)^2 \quad (22)$$

where $|\mathcal{T}_{\text{trade}}|$ is the number of trade signals. If no trades are generated, the objective is 0.

To reconcile signal quality with trading frequency, this calibration metric is combined with the raw trading success rate, $\mathcal{O}_{\text{TradeSuccess}}(\Theta)$, forming the final hybrid objective function:

$$\mathcal{O}_{\text{HYB}}(\Theta) = \omega \mathcal{O}_{\text{C-MSE}}(\Theta) + (1 - \omega) \mathcal{O}_{\text{TradeSuccess}}(\Theta), \quad (23)$$

where $\omega \in [0, 1]$ is a weighting parameter. This objective function is maximized over the training dataset using a global optimization algorithm.

4 Empirical Implementation and Calibration of the Cointegrated Ising Spin Model

The empirical validity and operational viability of the Cointegrated Ising Spin Model (CISM) are established through its implementation and calibration as a novel framework for generating real-time, probabilistic trading signals in high-frequency environments. The Cointegrated Ising Spin Model establishes a sophisticated synthesis between macro-level econometric principles and micro-level agent dynamics by conceptualizing the prices of two cointegrated futures contracts as aggregate magnetizations emerging from an underlying Ising spin system. Within this system, agent decisions are governed by a triad of forces—herding, momentum, and arbitrage—the last of which incorporates a central methodological advance: a novel Δ -weighted adjustment mechanism. This mechanism dynamically corrects for the informational staleness in-

herent in asynchronous trade arrivals, a feature of paramount importance for live operational viability. The empirical exercise presented herein is therefore not a conventional backtest but rather the crucial training and calibration phase for a real-time decision engine. Its dual objectives are, first, to demonstrate the model’s capacity to be robustly trained on high-frequency, event-time tick data, and second, to verify that the estimated parameters yield economically meaningful and stable dynamics requisite for reliable deployment in live trading conditions.

The empirical validation is conducted using high-frequency tick data for the two most liquid Brent crude oil futures contracts—the front-month (LC0c1) and second-month (LC0c2)—sourced from the LSEG Eikon platform and traded on the Intercontinental Exchange (ICE) in London. Each contract represents 1,000 barrels of crude oil, is quoted in U.S. dollars per barrel with a minimum tick size of \$0.01, and adheres to a standard expiration schedule central to calendar-spread trading. The ICE Brent futures market operates electronically on a near-continuous basis (02:00 to 23:00 London time) and exhibits a well-defined intraday liquidity profile that peaks during the London–New York session overlap, rendering it an ideal environment for this study. For the calibration phase, the training dataset comprises a representative trading day, 20 June 2025, from 09:30 to 23:00 London time. A raw tick-by-tick capture at a one-second sampling rate initially yielded approximately 42,000 time-stamped observations. To construct an event-time series reflective of actual market activity, non-trading intervals were systematically removed, retaining only timestamps with at least one trade in either contract. This filtering process resulted in a final dataset of approximately 16,000 observations, corresponding to an average trade arrival in the spread every 2.6 seconds. This resultant asynchronicity is not treated as a data imperfection but is instead embraced as a fundamental market characteristic that the CISM is architected to exploit. The elapsed time between consecutive trades, denoted Δ_t , becomes a direct and essential input that modulates the model’s arbitrage intensity. When operationalized, the same Δ_t measure is computed continuously from the incoming market data feed, allowing the model to dynamically adjust its internal state and generate updated trading probabilities at the native frequency of the market itself. To our knowledge, this work represents the first empirical realization of an Ising-based model specifically designed for, and validated on, asynchronous financial time series. Before proceeding to the calibration of the Cointegrated Ising Spin Model, it is essential to formally establish the statistical properties of the underlying price series and confirm the existence of

a long-run equilibrium relationship, which is a fundamental prerequisite for the model. The initial step in this preliminary analysis is to examine the descriptive statistics of the processed, event-time data series, which are presented in Table 1.

Table 1: Summary Statistics for Futures Prices

Series	Observations	Mean (\$)	Std. Dev.	Minimum (\$)	Maximum (\$)
P_t^A (LC0c1)	16,318	76.78	0.359	75.61	77.50
P_t^B (LC0c2)	16,318	75.24	0.327	74.09	75.86

A Vector Error Correction Model (VECM) is estimated to confirm the presence of a long-run equilibrium. The spread, or cointegrating residual, is defined as:

$$S_t = P_t^A - \gamma P_t^B - \mu_z, \quad (24)$$

where γ is the cointegrating coefficient and μ_z is the long-run equilibrium level. The VECM specification is given by:

$$\Delta P_t^A = \alpha_A(P_{t-1}^A - \gamma P_{t-1}^B - \mu_z) + \phi_{AB}\Delta P_{t-1}^B + \epsilon_t^A, \quad (25)$$

$$\Delta P_t^B = \alpha_B(P_{t-1}^A - \gamma P_{t-1}^B - \mu_z) + \phi_{BA}\Delta P_{t-1}^A + \epsilon_t^B, \quad (26)$$

where α_A and α_B represent the speeds of adjustment towards equilibrium, and ϕ_{AB} , ϕ_{BA} capture short-run dynamic effects.

The VECM estimation results, presented in Table 2, provide strong statistical evidence for a stable cointegrating relationship. The estimated cointegrating coefficient $\gamma = 1.02045$ is highly significant and economically intuitive, reflecting a market in weak contango where the forward price trades at a slight premium to the spot price, consistent with storage and financing costs between contract expiries.

The adjustment speed parameters are both statistically significant and possess the correct signs for a stable system. The negative sign of $\alpha_A = -0.02574$ indicates that the front-month contract decreases in price to correct a positive disequilibrium, while the positive sign of $\alpha_B = 0.02011$ shows the second-month contract increases. The relative magnitudes suggest an asymmetric adjustment process where the more liquid front-month contract bears approximately 56% of the total correction. The significant short-run coefficients $\phi_{AB} = 0.414$ and $\phi_{BA} = 0.337$ reveal strong positive feedback and momentum spillover between contracts, characteristic of cointegrated futures trading

Table 2: VECM Estimated Coefficients

Variable	Coefficient	Std. Error	T-Statistic	p-value
<i>Adjustment Speeds</i>				
α_A	-0.02574	0.00214	-12.01	0.0000
α_B	0.02011	0.00225	8.93	0.0000
<i>Cointegrating Parameter</i>				
γ	1.02045	0.00005	20596.37	0.0000
<i>Short-Run Dynamics</i>				
ϕ_{AB}	0.41446	0.00257	161.15	0.0000
ϕ_{BA}	0.33730	0.00210	160.28	0.0000

in high-frequency environments.

Figure 1 visually confirms the stationarity and mean-reverting behavior of the Error Correction Term, providing the statistical foundation upon which the real-time trading strategy is built.




Figure 1: Time Series of the Estimated Error Correction Term (ECT). The series represents the deviation from the long-run equilibrium, calculated as $ECT_t = P_t^A - 1.02045 \cdot P_t^B$. The plot visually confirms the stationary, mean-reverting nature of the spread.

4.1 Agent-Based Model Calibration and Trading Performance

Following the establishment of cointegration, this study calibrates the agent-based model using a performance-driven methodology specifically designed for real-time operation. The model's key innovation is a Δ -weighted arbitrage force that dynamically

amplifies corrective pressure based on information staleness, making it uniquely suited for asynchronous high-frequency environments where trade arrivals are irregular.

Model parameters are optimized using a hybrid objective function that balances probabilistic forecast accuracy with raw trading success:

$$\mathcal{O}_{\text{HYB}}(\Theta) = \omega \mathcal{O}_{\text{C-MSE}}(\Theta) + (1 - \omega) \mathcal{O}_{\text{TradeSuccess}}(\Theta), \quad (27)$$

where $\mathcal{O}_{\text{C-MSE}}$ measures the conditional mean squared error between predicted and realized reversals on trade events, and $\mathcal{O}_{\text{TradeSuccess}}$ is the raw in-sample success rate. Optimization employs a staged approach combining coarse grid search, simulated annealing, and derivative-free local solvers.

The calibrated parameters, presented in Table 3, reveal a compelling market ecology dominated by sophisticated arbitrage activity. The high proportion of contrarian agents ($N_C = 78\%$) suggests an environment rich with latent arbitrage capital poised to enforce mean reversion. The nearly equal magnitudes of momentum strength ($J_{\text{mom}} = 34.225$) and contrarian strength ($\alpha = 34.888$) describe a high-tension equilibrium where momentum-driven deviations meet aggressive corrective responses.

Table 3: Calibrated Parameters for the Agent-Based Model

Parameter	Calibrated Value
Proportion of Contrarians, N_C	78
Agent Determinism, β_{response}	5.387
Momentum Strength, J_{mom}	34.225
Contrarian Strength, α	34.888
Staleness Weight, k_{Δ}	25.282
Sell Threshold, θ_{sell}	0.504
Buy Threshold, θ_{buy}	0.454

The high agent determinism ($\beta_{\text{response}} = 5.387$) indicates rational, purposeful response to market signals. Crucially, the substantial staleness weight ($k_{\Delta} = 25.282$) confirms the importance of the Δ -weighted arbitrage mechanism, revealing that information staleness amplifies rather than degrades perceived arbitrage opportunities in this market.

The model’s practical performance is evaluated through comprehensive backtesting. Table 4 details the trading results, demonstrating the effectiveness of using the model’s probabilistic output for real-time signal generation.

Table 4: Trading Strategy Performance

Trade Execution Summary		Performance Summary	
Sell Trades	80	Overall Success Rate	74.65%
Successful Sells	63	Sell Success Rate	78.75%
Buy Trades	137	Buy Success Rate	72.26%
Successful Buys	99	Total Trades	217

The backtest yielded 217 trades with an overall success rate of 74.65%. The strategy demonstrates nuanced state-dependent behavior: stricter sell conditions generated fewer trades (80) but higher success (78.75%), while slightly more lenient buy thresholds permitted more entries (137 trades) to capture violent snap-back reversals while maintaining strong performance (72.26%). This asymmetry likely reflects day-specific microstructural effects that the model successfully navigates.

4.2 Model Diagnostics and Dynamics

An examination of the key statistical series provides crucial insight into the model’s inner workings and validates the trading strategy’s premises. Table 5 presents distributional properties of both the Error Correction Term and the model’s primary output, the conditional probability of reversion ($P_t(\text{reversal})$).

The ECT’s stationarity is confirmed by a runs test that decisively rejects the random walk null hypothesis (Z-Score = -45.72, $p < 0.01$), establishing that spread deviations are indeed temporary. The distribution of $P_t(\text{reversal})$ is well-behaved and centered (mean 0.3943, median 0.3682), indicating the model is discerning and does not perpetually signal high reversion likelihood. The negative excess kurtosis (-1.326) suggests a platykurtic distribution capable of generating a wide range of probability values.

The model’s effectiveness emerges most clearly in the distribution tails. A powerful relationship exists where large, infrequent spread deviations trigger sharp increases in reversion probability. When the ECT reaches its 95th percentile (deviation of +0.0874), the reversal probability jumps to 0.8645—far exceeding trading thresholds. This relationship strengthens further at the 99th percentile. This dynamic is precisely what the trading strategy exploits: lying in wait for high-conviction forecasts that occur only during significant market dislocations (top 5-10% of spread deviations).

Table 5: Descriptive Statistics: ECT and Reversal Probability

Metric	ECT (Spread)	Prob(Reversal)
Summary Statistics		
Mean	0.0002	0.3943
Std. Dev.	0.0525	0.2905
Skewness	0.5127	0.2411
Kurtosis (excess)	0.3251	-1.3260
Minimum	-0.2815	0.0000
Maximum	0.2238	0.9489
Percentiles		
1st	-0.0946	0.0003
5th	-0.0767	0.0109
25th	-0.0377	0.1170
50th (Median)	-0.0036	0.3682
75th	0.0293	0.6823
95th	0.0874	0.8645
99th	0.1510	0.8934
Stationarity Test		
Runs Test (Z-Score)	-45.72***	—

4.3 Robustness and Sensitivity Analysis

This study conducts targeted sensitivity analyses to validate the model’s architecture against overfitting and establish parameter necessity.

1. **Staleness Weight (k_{Δ}):** Reducing the calibrated value ($k_{\Delta} = 25.282$) by an order of magnitude materially decreased both trade signals and success rate, providing strong evidence that the Δ -weighted arbitrage force is a primary source of predictive edge.
2. **Agent Determinism (β_{response}):** Systematically lowering β_{response} compressed the predictive distribution of $P_t(\text{reversal})$ toward 0.5, eroding signal clarity and reducing performance. This validates that strategy success depends on identifying high-conviction moments requiring agent rationality.
3. **Relative Balance of Forces (J_{mom}/α):** Altering the near-equal ratio of momentum to contrarian strength induced regime shifts. Higher ratios created trend-dominated markets with lower success rates; lower ratios led to passive markets

with insufficient trades. This demonstrates critical sensitivity to the precise competitive balance between opposing market forces.

These analyses collectively demonstrate that the model’s performance arises not from any single parameter but from the synergistic interplay of its core features: staleness amplification, agent determinism, and the delicate balance between trend and counter-trend forces. The architecture proves robust to perturbations, with performance degradation following predictable patterns when key mechanisms are compromised.

5 Discussion

The empirical results demonstrate the potential of the proposed Cointegrated Ising Spin framework. Before concluding, it is useful to contextualize the model’s contribution by comparing its core features to those of traditional financial and econophysical models, and to discuss the study’s limitations and avenues for future research.

5.1 Model Novelty and Strategic Value

While the proposed framework builds upon concepts from econophysics, it introduces key innovations that tailor it specifically to high-frequency statistical arbitrage, offering distinct advantages over traditional approaches. This model is designed as a **real-time signal generator for non-directional arbitrage strategies**. Its core logic is not limited to futures calendar spreads but can be applied to any pair of cointegrated assets—such as cross-listed equities, related ETFs, or inter-commodity spreads—that exhibit strong mean-reverting forces and are traded asynchronously at high frequencies.

The model’s novelty can be summarized in three key departures from the standard approaches in both econometrics and agent-based modeling:

1. **Endogenous, Econometrically-Grounded Arbitrage Force:** Traditional agent-based and Ising models (e.g., [2, 12]) often rely on abstract or exogenous fundamental values, deriving agent decisions from forces like local herding and a global contrarian field. In contrast, the model’s primary driver is the novel **‘Pull-Term’**. This term is an endogenous force derived directly from an empirically estimated, macro-level econometric relationship—the cointegration vector. This

elegantly anchors micro-level agent behavior to an observable, mean-reverting macro-level equilibrium, bridging the gap between econometrics and econophysics and providing a theoretically sound foundation for the trading signals.

2. **Explicit Modeling and Exploitation of Asynchronous Time:** Standard time-series models and Ising models struggle with asynchronous data, often assuming synchronous time steps $(t, t + 1, \dots)$. The model’s key innovation is how it operationalizes a Δ -weighted arbitrage force. By explicitly incorporating the real-world time elapsed since the last trade (Δ_t) for each asset, this framework turns asynchronicity from a data problem into a source of alpha. This dynamic weighting of the corrective force yields a more accurate, real-time response to market disequilibrium that is uniquely suited for generating signals from tick-level data.
3. **A Prescriptive, Probabilistic Output:** The output of many quantitative models is a latent state (e.g., aggregate magnetization) or a binary signal. The model’s primary output is a **prescriptive, conditional probability of mean reversion**. By aggregating agent decisions, the model directly calculates an instantaneous, actionable probability that the spread will converge. This transforms the model from a purely descriptive tool into a sophisticated, probabilistic signal generator, allowing for the construction of more nuanced, risk-aware strategies that move beyond static threshold rules.

5.2 Limitations and Future Research

Despite the promising results, it is crucial to acknowledge the limitations of this study, which provide clear avenues for future work.

First, the model parameters were calibrated and backtested on a historical dataset from a single trading day. This approach, while useful for demonstrating the framework’s internal consistency and potential, carries a significant risk of overfitting. A crucial next step is to perform rigorous out-of-sample validation by training the model on data from one period and testing its performance on subsequent, unseen data to assess its true predictive power and robustness across different market regimes.

Second, the backtest was conducted without explicit **transaction costs**. In a live market, the bid-ask spread, commissions, and potential slippage represent direct and unavoidable costs that would reduce the reported profitability. A comprehensive

assessment requires incorporating these frictions, though the highly liquid nature of the Brent futures market provides a robust foundation for the model’s initial validation.

Third, the **small trade sample** (217 trades) generated in the one-day backtest is a limitation for conclusively validating a high-frequency strategy. A longer testing period generating a larger number of trades is required to confirm the statistical significance and robustness of the performance.

Finally, the framework itself has scope for improvement. Parameter calibration could be enhanced with more dynamic optimization methods to better adapt to changing market conditions. The objective function could also be modified to maximize risk-adjusted metrics like the Sharpe ratio instead of raw success rates. Further extensions could involve endogenizing the trade arrival process or incorporating a wider variety of heterogeneous agent types with different memory lengths and risk tolerances.

6 Conclusions

This study demonstrates that combining econometric cointegration with agent-based modeling can yield a powerful, probabilistic trading framework for asynchronous markets. The Cointegrated Ising Spin Model integrates the statistical structure of the VECM with the behavioral microdynamics of agents, translating macro-level error correction into real-time, probability-based trading signals. Applied to Brent crude oil futures, the model achieved a 74.65% predictive success rate, validating the economic relevance of the Δ -weighted arbitrage mechanism.

This paper has addressed the significant challenge of modeling and trading cointegrated, asynchronously traded assets, with a direct application to Brent crude oil futures spreads. We proposed a novel hybrid framework that integrates the macro-level dynamics of a Vector Error Correction Model (VECM) with the micro-foundations of an agent-based Ising spin model. The key innovation is the introduction of a ‘PullTerm’ into the agent decision-making process, ...which is driven by a Δ -weighted arbitrage force that explicitly accounts for information staleness.

The parameters of the agent-based model were estimated through a direct, performance-driven optimization of a probability-based trading strategy, aligning the model’s calibration with its practical goal. The in-sample backtest of the optimized strategy yielded a strong overall success rate of 74.65%, highlighting the potential of using such hybrid agent-based models to move beyond static-threshold arbitrage strategies toward

more dynamic, probability-driven approaches.

By calibrating the Cointegrated Ising Spin Model on historical tick data and subsequently feeding it with live market quotes, the proposed framework establishes a deployable architecture for automated spread trading. This dual design—training offline and executing online—extends the conventional boundaries of statistical arbitrage by transforming an econometric model into an adaptive, continuously operating decision system. In doing so, it enables real-time monitoring and signal generation in markets where asynchronicity and information staleness are inherent features of price formation.

Overall, this study demonstrates that bridging statistical time-series modeling with agent-based simulation yields a robust and versatile framework for capturing the microdynamics of spread reversion. The methodology provides a practical blueprint for next-generation, non-directional arbitrage strategies capable of functioning in live trading environments. Its flexibility allows seamless adaptation across asset classes characterized by liquidity, asynchronous price discovery, and mean-reverting behavior, paving the way for a new class of real-time, probabilistic trading systems grounded in sound econometric principles.

References

- [1] Alfarano, S., Lux, T., Wagner, F. (2005). Estimation of agent-based models: the case of an asymmetric herding models. *Computational Economics*, 26, 19–49.
- [2] Bornholdt, S. (2001). Expectation Bubbles in a Spin Model of Markets: Intermittency from Frustration across Scales. *International Journal of Modern Physics C*, 12(05), 667–674.
- [3] Brock, W. A. & Hommes, C. H. (1998). Heterogeneous Beliefs and Routes to Chaos in a Simple Asset Pricing Model. *Journal of Economic Dynamics and Control*, 22(8-9), 1235–1274.
- [4] Callen, E., & Shapero, D. (1974). A theory of social imitation. *Physics Today*, 27(7), 23–28.
- [5] Cont, R. (2001). Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance*, 1(2), 223–236.

- [6] Dvořák, P. (2012). *From Microscopic Rules to Macroscopic Phenomena: Ising Model in Finance*. (Unpublished bachelor thesis). Charles University in Prague.
- [7] Engle, Robert F. and C. W. Granger. (1987). Co-integration and Error Correction: Representation, Estimation and Testing. *Econometrica*, 55(2), 251–276.
- [8] Enders, Walter. (1995). *Applied Econometric Time Series*. New York: John Wiley & Sons, Inc.
- [9] Gatev, Evan, G., William, N. Goetzmann, and K. Greet Rouwenhorst. (1999). *Pairs Trading: Performance of a Relative Value Arbitrage Rule*. NBER Working Papers 7032, National Bureau of Economic Research Inc.
- [10] He, C., Wang, T., Liu, X., & Huang, K. (2023). An innovative high-frequency statistical arbitrage in Chinese futures market. *Journal of Innovation Knowledge*, 8(2023), 100429.
- [11] Hommes, C. H. (2006). Heterogeneous Agent Models in Economics and Finance. In L. Tesfatsion & K. L. Judd (editors), *Handbook of Computational Economics*, volume 2, chapter 23, pp. 1109–1186. Elsevier.
- [12] Kaizoji, T., S. Bornholdt, & Y. Fujiwara (2002). Dynamics of Price and Trading Volume in a Spin Model of Stock Markets with Heterogeneous Agents. *Physica A: Statistical Mechanics and its Applications*, 316(1-4), 441–452.
- [13] Kirman, A. (1993). Ants, rationality, and recruitment. *The Quarterly Journal of Economics*, 108(1), 137–156.
- [14] Kukacka, J., & Kristoufek, L. (2020). Do ‘complex’ financial models really lead to complex dynamics? Agent-based models and multifractality. *Journal of Economic Dynamics and Control*, 113, 103855.
- [15] Lazzarino, M., Berrill, J. and Šević, A. (2018). What Is Statistical Arbitrage? *Theoretical Economics Letters*, 8, 888-908.
- [16] Sieczka, P. & J. A. Holyst (2007). A Threshold Model of Financial Markets. arXiv:0711.3106 [q-fin.ST].
- [17] Stock, James H. and Mark W. Watson. (1988). Testing for Common Trends. *Journal of the American Statistical Association*, 83(404), 1097–1107.

- [18] Vidyamurthy, G. (2004). *Pairs Trading: Quantitative Methods and Analysis*. John Wiley & Sons.
- [19] Yan, T., Chiu, M. C., & Wong, H. Y. (2022). Pairs trading under delayed cointegration. *Quantitative Finance*, 22(9), 1627–1648.