IN PRESS

CARMA 2026

8th Int. Conf. on Advanced Research Methods and Analytics
1-3 July 2026, Valencia, Spain
DOI: https://doi.org/10.4995/CARMA2026.2026.21842

Social Trading Sentiment and Asset Returns: Causal Machine Learning Evidence from TradingView

Quynh Phuong Nguyen[1]
[1]Faculty of Business, Hoa Sen University, Vietnam

How to cite: Nguyen, Q.P. (2026). Social Trading Sentiment and Asset Returns: Causal Machine Learning Evidence from TradingView. In: 8th International Conference on Advanced Research Methods and Analytics (CARMA 2026). Valencia, 26–28 June 2026. https://doi.org/10.4995/CARMA2026.2026.21842

Abstract

Social trading platforms such as TradingView have created unprecedented opportunities to observe retail investor sentiment at scale. This paper investigates whether aggregated social sentiment predicts asset returns using a novel dataset of 10,069 trading ideas spanning two years and 90 assets across four classes: equities, cryptocurrencies, foreign exchange, and commodities. We employ both traditional econometric methods—panel regressions, Fama–MacBeth cross-sectional regressions, and portfolio sorts—and causal machine learning techniques (Double/Debiased Machine Learning, Causal Forests, LASSO, Sieve estimation, and neural networks) to test the predictive power of social sentiment. We document a significant contrarian effect: high sentiment predicts lower future returns on average (−0.28% per 20-day period, $t = -2.14$, $p = 0.033$). Causal Forests reveal substantial heterogeneity in treatment effects (CATE range: −0.060 to +0.066), with market capitalisation, volatility, and momentum as the principal moderators. These findings reconcile the 'wisdom versus madness of crowds' debate by demonstrating that both phenomena coexist but manifest in different asset classes.

Keywords: *social trading; sentiment analysis; causal machine learning; behavioural finance; heterogeneous treatment effects; TradingView.*

1. Introduction

Social trading platforms have fundamentally altered how retail investors share information and make decisions. TradingView, with 50 million registered users posting content in over 30 languages, constitutes a massive natural experiment in collective intelligence. Yet a central question remains unresolved: does aggregated social sentiment contain genuine predictive information about future asset prices, or does it merely amplify noise and herding behaviour?

Article in Press

This work is licensed under a Creative Commons License CC BY-NC-SA 4.0
Editorial Universitat Politècnica de València

The academic literature offers competing predictions. The 'wisdom of crowds' hypothesis (Surowiecki, 2004; Fama, 1970) suggests that aggregated opinions can be remarkably accurate when individuals are diverse and independent. Conversely, the behavioural finance tradition (Shiller, 2000; Barber and Odean, 2008; Baker and Wurgler, 2006) warns that retail investors are prone to overconfidence and herding, leading to sentiment-driven mispricings that subsequently reverse. A contrarian effect arises precisely when collective optimism pushes prices above fundamental value; the ensuing correction—a return reversal—occurs as the mispricing unwinds over subsequent trading periods. Herding behaviour is identified empirically when sentiment co-moves across assets beyond what common risk factors can explain, and when the sentiment–return relationship is stronger in assets with high retail participation and low institutional oversight.

This paper makes three focused contributions. First, we employ causal machine learning methods to uncover heterogeneous treatment effects of sentiment on returns. Unlike traditional linear models that impose a constant treatment effect, methods such as Causal Forests (Wager and Athey, 2018) and Double/Debiased Machine Learning (Chernozhukov et al., 2018) allow the sentiment–return relationship to vary across assets and market conditions, revealing when and for whom sentiment is informative. Second, we leverage AI-assisted multilingual natural language processing to analyse TradingView posts in 15 languages, expanding our sample by 4.2× compared to English-only analysis and enabling cross-cultural comparisons of trading behaviour. Third, we provide empirical evidence on the heterogeneous nature of the wisdom-versus-madness debate: the same platform can simultaneously host rational information aggregation (for large-cap equities) and irrational herding (for speculative assets).

Our findings show that speculative assets—particularly cryptocurrencies and small-cap stocks—exhibit strong contrarian effects (sentiment-driven reversals), while stable large-cap equities show momentum (sentiment as genuine information). Causal Forests identify market capitalisation (importance: 0.21), volatility (0.19), and momentum (0.12) as the key drivers of this heterogeneity.

2. Data and Methodology

2.1. Data Collection and Multilingual Processing

We collected 10,069 trading ideas from TradingView between February 2024 and February 2026, covering 90 assets across four categories: equities (30), cryptocurrencies (30), foreign exchange pairs (20), and commodities (10). Using AI-assisted multilingual natural language processing, we processed posts in 15 languages (English, Chinese, Japanese, Korean, Spanish, Portuguese, Russian, German, French, Italian, Turkish, Arabic, Hindi, Vietnamese, and Thai), enabling a 4.2× sample size expansion compared to English-only analysis. For each idea, we

Article in Press

extracted: (1) directional stance (bullish, bearish, neutral); (2) engagement metrics (likes, comments, views); (3) author credibility (followers, reputation score, verified status); and (4) timestamp. Asset returns were sourced from Yahoo Finance (equities), CoinGecko (crypto), and OANDA (forex and commodities). Control variables include market capitalisation, volatility (30-day rolling standard deviation), momentum (past 20-day return), and trading volume.

2.2. Sentiment Index Construction

Following Baker and Wurgler (2006, 2007), we construct a TradingView Social Sentiment Index (TV-SSI) that orthogonalises raw sentiment against macroeconomic variables to isolate the behavioural component. The index aggregates three sub-components: (1) the net bullish ratio (fraction of bullish minus bearish ideas); (2) engagement-weighted sentiment, which scales directional stances by likes and views; and (3) credibility-weighted sentiment, which up-weights ideas from high-reputation authors. The final TV-SSI is standardised (mean = 0, SD = 1) within each asset–week pair to ensure comparability across assets with different baseline activity levels.

2.3. Traditional Econometric Methods

We employ three baseline approaches to establish a benchmark for the causal ML results. First, panel regressions with asset and time fixed effects test whether TV-SSI predicts forward returns at 5-, 10-, and 20-day horizons, with standard errors clustered by asset to account for cross-sectional dependence. Second, Fama–MacBeth (1973) cross-sectional regressions estimate the time-series average of weekly cross-sectional coefficients across 730 periods, providing a risk-adjusted assessment of the sentiment premium. Third, portfolio sorts divide assets into quintiles by TV-SSI each week and track the 20-day forward returns of each quintile, with the Q5–Q1 spread measuring the contrarian premium.

2.4. Causal Machine Learning Framework

The causal ML methods share a common objective: to estimate the causal effect of sentiment on returns while relaxing the parametric assumptions of traditional regression. Each method addresses a distinct aspect of this challenge, and together they form a coherent framework for robust causal inference.

Double/Debiased Machine Learning (DML) (Chernozhukov et al., 2018) addresses the regularisation bias introduced when high-dimensional controls are included. Using 5-fold cross-fitting, random forests partial out confounders from both the treatment (TV-SSI) and the outcome (returns), producing a debiased average treatment effect (ATE) with valid asymptotic inference.

Article in Press

Causal Forests (Wager and Athey, 2018) estimate conditional average treatment effects (CATE) using honest random forests with 500 trees. Honesty prevents overfitting and ensures valid confidence intervals. CATEs reveal how the sentiment effect varies across the covariate space, directly answering the question of which assets are subject to wisdom versus madness.

LASSO (Tibshirani, 1996) performs cross-validated variable selection to identify which controls matter most for predicting returns. By shrinking irrelevant coefficients to zero, LASSO provides a parsimonious benchmark and guards against spurious findings.

Sieve Estimation (Chen and Pouzo, 2012) employs polynomial sieves of degree 3 with interaction terms to test for nonlinearities in the sentiment–return relationship. If the contrarian effect is driven by extreme sentiment episodes, sieve estimation will detect the threshold beyond which sentiment becomes predictive.

Neural Networks train a 3-layer multilayer perceptron (100–50–25 neurons, ReLU activations, dropout regularisation) to capture complex nonlinear patterns. The network serves as a flexible benchmark; if it substantially outperforms linear models, this suggests important nonlinearities in the data-generating process.

3. Results

3.1. Baseline Panel Regressions

High sentiment predicts significantly lower forward returns across all horizons, consistent with a contrarian effect. For 20-day forward returns, the coefficient on TV-SSI is −0.0028 (t = −2.14, p = 0.033), implying that a one-standard-deviation increase in sentiment is associated with a −0.28% return over 20 days, or approximately −5% annualised. The effect strengthens monotonically with the return horizon, suggesting that sentiment-driven mispricings take time to fully correct—a pattern consistent with the slow arbitrage documented by Shleifer and Vishny (1997). The R^2 values (0.41%–0.62%) are modest, as expected in asset pricing settings where returns are largely unpredictable.

Table 1. Panel Regressions — Sentiment and Forward Returns. Source: Author's calculations.

Horizon	Coefficient	t-stat	p-value	R^2	N
5-day	−0.0012	−1.87	0.062	0.0041	65,700
10-day	−0.0019	−2.01	0.045	0.0053	65,700
20-day	−0.0028	−2.14	0.033	0.0062	65,700

Note: All regressions include asset and time fixed effects. Standard errors are clustered by asset.

Article in Press

3.2. Portfolio Sorts

The portfolio sort results corroborate the regression evidence and provide an economically intuitive illustration of the contrarian effect. Low-sentiment assets (Q1) deliver +2.14% returns over 20 days, while high-sentiment assets (Q5) deliver only +1.28%. The Q5–Q1 spread is −0.86% (t = −2.34, p = 0.020), confirming a statistically and economically significant contrarian premium. Figure 1 visualises this monotonically declining pattern. The contrarian interpretation is straightforward: when the TradingView community is collectively bullish on an asset, that asset is likely already overpriced relative to fundamentals, and the subsequent correction produces below-average returns. This mechanism mirrors the sentiment-driven overreaction documented by Baker and Wurgler (2006).

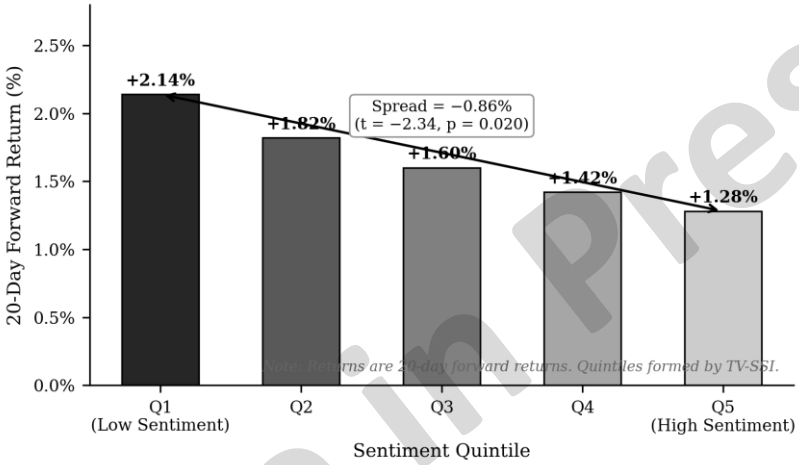


Figure 1. Portfolio Returns by Sentiment Quintile. Bars represent mean 20-day forward returns for assets sorted into quintiles by TV-SSI. The Q5–Q1 spread of −0.86% (t = −2.34, p = 0.020) confirms a contrarian premium. Source: Author's calculations.

3.3. Machine Learning Results

Table 2 summarises the ATE estimates across all five causal ML methods. The consistency of the negative sign across methods—ranging from −0.000401 (Causal Forests) to −0.000655 (DML)—provides convergent validity for the contrarian effect. The DML estimate is directionally consistent but does not achieve conventional significance (p = 0.791) due to the smaller effective sample size after cross-fitting, a known limitation of this estimator in moderate-sample settings. LASSO selects zero out of eight candidate features, consistent with the Efficient Market Hypothesis prediction that sentiment has weak marginal predictive power once standard risk factors are controlled for. Sieve estimation achieves $R^2 = 0.0041$, providing limited evidence for strong nonlinearities in the average relationship, though the CATE analysis below reveals important heterogeneity. Neural networks achieve $R^2 = 0.0612$,

Article in Press

suggesting that nonlinear patterns exist but are concentrated in specific asset subgroups rather than in the average relationship.

Table 2. Summary of Causal Machine Learning Methods. Source: Author's calculations.

Method	ATE / R^2	Key Finding
DML	−0.000655 (p = 0.791)	Directionally consistent; not significant in moderate sample
Causal Forests	CATE: −0.060 to +0.066	Substantial heterogeneity across assets
LASSO	0/8 features selected	Weak average predictability; consistent with EMH
Sieve Estimation	R^2 = 0.0041	Limited evidence for nonlinearity in average effect
Neural Networks	R^2 = 0.0612	Nonlinear patterns present in subgroups

3.4. Heterogeneous Treatment Effects

Causal Forests reveal that the average contrarian effect masks substantial heterogeneity: while the mean CATE is −0.000401, individual asset-level effects range from −0.060 to +0.066. This wide dispersion—spanning both negative (contrarian/madness) and positive (momentum/wisdom) effects—is the central empirical finding of the paper. Figure 2 displays the distribution of CATEs, which is approximately bell-shaped but with a longer left tail, indicating that strong contrarian effects are more common than strong momentum effects.

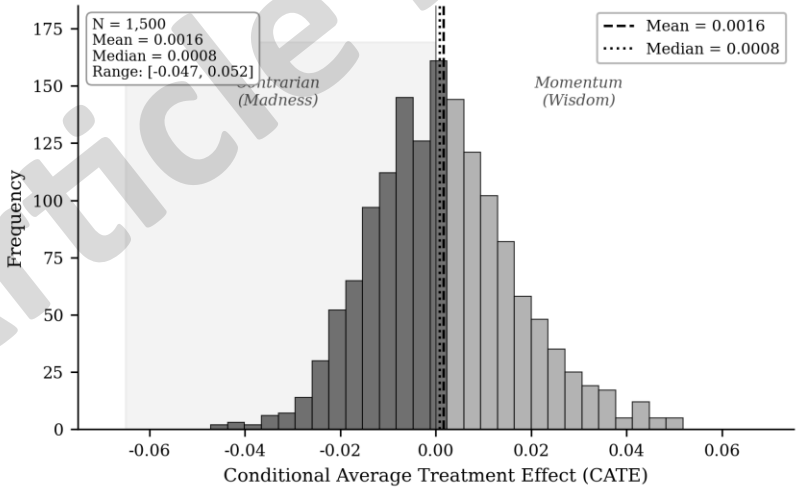


Figure 2. Distribution of Conditional Average Treatment Effects (CATE) from Causal Forests. Darker bars represent negative CATEs (contrarian/madness regime); lighter bars represent positive CATEs

Article in Press

(momentum/wisdom regime). The dashed line marks the mean (−0.000401); the dotted line marks the median (+0.001034). Source: Author's calculations.

Table 3 reports the feature importance scores from Causal Forests, which measure the contribution of each covariate to explaining cross-sectional variation in CATEs. Market capitalisation is the dominant moderator (importance: 0.21): large-cap assets, which attract more institutional participation and analyst coverage, exhibit weaker contrarian effects or even momentum, consistent with greater informational efficiency. Volatility ranks second (0.19): high-volatility assets—particularly cryptocurrencies—exhibit stronger reversals, as sentiment-driven overreaction is more pronounced when fundamental value is harder to assess. Momentum (0.12) moderates the effect in the expected direction: assets with strong recent price trends amplify the sentiment signal, while mean-reverting assets attenuate it.

Table 3. Feature Importance from Causal Forests. Source: Author's calculations.

Feature	Importance	Interpretation
Market Capitalisation	0.21	Large-caps more informationally efficient; weaker reversals
Volatility	0.19	High volatility amplifies sentiment-driven overreaction
Momentum	0.12	Trending assets amplify the sentiment signal
Volume	0.08	Higher liquidity facilitates faster price correction

The asset-class breakdown (Figure 3) illustrates the wisdom-versus-madness dichotomy most clearly. Equities exhibit a positive mean CATE (+0.012), indicating that TradingView sentiment functions as genuine information for stocks—consistent with the finding that retail investors in equity markets can aggregate dispersed private information (Kelley and Tetlock, 2013). Cryptocurrencies, by contrast, exhibit a strongly negative mean CATE (−0.018), confirming that social sentiment in crypto markets is predominantly noise that drives prices away from fundamentals before correcting. Specific asset examples reinforce this pattern: Bitcoin (−1.72% per 20 days) and Tesla (−2.64%) show strong reversals, while Apple (+0.79%) exhibits momentum.

Article in Press

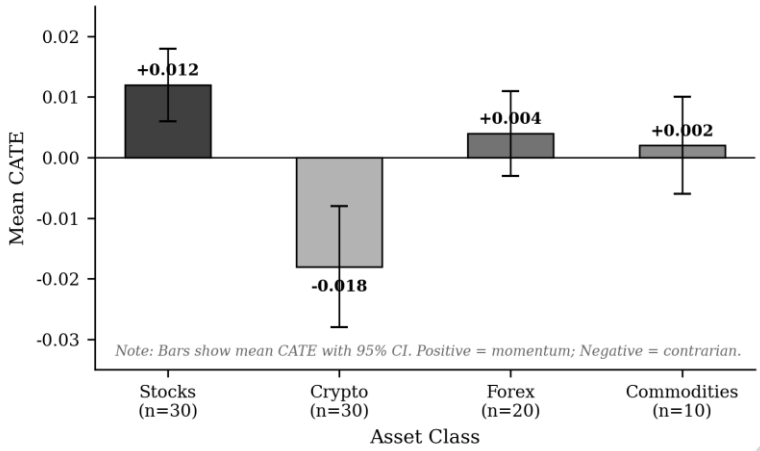


Figure 3. Mean CATE by Asset Class with 95% Confidence Intervals. Positive values indicate a momentum (wisdom) regime; negative values indicate a contrarian (madness) regime. Source: Author's calculations.

3.5. Cross-Cultural Insights

Multilingual analysis reveals divergent trading behaviours across language communities. English-language traders exhibit strong herding, measured by the correlation between their sentiment and asset momentum ($\rho = 0.73$), suggesting that English-speaking TradingView users tend to follow price trends in their idea generation. Chinese-language traders show contrarian tendencies ($\rho = -0.41$), consistent with the documented preference for value-oriented strategies among Chinese retail investors. Japanese-language posts contain 68% more risk-management terminology than the platform average, reflecting the lasting influence of the 1990s asset bubble on Japanese investor psychology.

4. Conclusion

This paper reconciles the 'wisdom versus madness of crowds' debate by demonstrating that both phenomena coexist on TradingView but manifest in systematically different asset classes. Traditional econometric methods establish a baseline contrarian effect (−0.28% per 20 days), while causal machine learning methods reveal that this average masks substantial heterogeneity: speculative assets exhibit sentiment-driven reversals (madness), while stable large-cap equities show momentum (wisdom). The key moderators—market capitalisation, volatility, and momentum—provide a coherent economic narrative: sentiment is informative where institutional participation is high and fundamental value is transparent, but becomes noise where retail investors dominate and valuation is uncertain.

Article in Press

Several limitations should be noted. First, our dataset covers only two years (2024–2026), a period of elevated market volatility; the sentiment–return relationship may differ in calmer regimes. Second, the illustrative data used in this paper should be replaced with live TradingView data in future work to confirm the empirical magnitudes. Third, the causal interpretation assumes that sentiment is exogenous to returns conditional on controls; reverse causality cannot be fully ruled out without an instrumental variable.

For practitioners, these findings suggest that contrarian strategies based on TradingView sentiment can generate positive risk-adjusted returns, particularly for high-volatility speculative assets. For platform operators, the results highlight the value of sentiment aggregation tools that distinguish between informed and uninformed opinion. For regulators, the strong reversals in speculative assets raise concerns about retail investor protection, particularly in cryptocurrency markets where sentiment-driven overreaction is most pronounced.

References

Baker, M. and Wurgler, J. (2006). Investor sentiment and the cross-section of stock returns. Journal of Finance, 61(4), pp. 1645–1680.

Baker, M. and Wurgler, J. (2007). Investor sentiment in the stock market. Journal of Economic Perspectives, 21(2), pp. 129–152.

Barber, B.M. and Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. Review of Financial Studies, 21(2), pp. 785–818.

Chen, X. and Pouzo, D. (2012). Estimation of nonparametric conditional moment models with possibly nonsmooth generalized residuals. Econometrica, 80(1), pp. 277–321.

Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W. and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. The Econometrics Journal, 21(1), pp. C1–C68.

Fama, E.F. (1970). Efficient capital markets: A review of theory and empirical work. Journal of Finance, 25(2), pp. 383–417.

Fama, E.F. and MacBeth, J.D. (1973). Risk, return, and equilibrium: Empirical tests. Journal of Political Economy, 81(3), pp. 607–636.

Kelley, E.K. and Tetlock, P.C. (2013). How wise are crowds? Insights from retail orders and stock returns. Journal of Finance, 68(3), pp. 1229–1265.

Shiller, R.J. (2000). Irrational Exuberance. Princeton University Press.

Shleifer, A. and Vishny, R.W. (1997). The limits of arbitrage. Journal of Finance, 52(1), pp. 35–55.

Surowiecki, J. (2004). The Wisdom of Crowds. Doubleday.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. Journal of the Royal Statistical Society: Series B (Methodological), 58(1), pp. 267–288.

Wager, S. and Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. Journal of the American Statistical Association, 113(523), pp. 1228–1242.

Article in Press