Can ChatGPT Beat the Market? Agentic LLM Strategies for Regime-Aware Equity Investment: A Six-Year Multi-Regime Empirical Study of TSMC (2020–2026)

Lanz CWJ Chan
Finamatrix Quant-Lab, Singapore
Email: lanz@finamatrix.net

Wing-Keung Wong
Department of Finance, Quantum AI Research Center, Fintech & Blockchain Research Center, and Big Data Research Center, Asia University, Taiwan
Department of Medical Research, China Medical University Hospital, Taiwan
Business, Economic and Public Policy Research Centre, Hong Kong Shue Yan University
The Economic Growth Centre, Nanyang Technological University, Singapore
No. 500, Liufeng Road, Wufeng Dist., Taichung City 41354, Taiwan
***Corresponding author** Email: wong@asia.edu.tw

Published on May 2026

Abstract

Purpose: This study investigates whether an agentic ChatGPT investment system can generate regime-aware, valuation-disciplined, and capital-constraint-conscious equity strategies that deliver superior risk-adjusted performance relative to passive benchmarks across multiple market regimes.

Design/Methodology/Approach: A five-module agentic pipeline integrates: (i) Hidden Markov Model (HMM) regime classification with geopolitical override; (ii) two-stage discounted cash flow (DCF) valuation anchored to audited FY2025 TSMC financials; (iii) Shannon entropy-minimising structured prompt engineering; (iv) Kelly criterion position sizing with a half-Kelly constraint; and (v) Volume-weighted Golden Ratio Estimator (vGRE) stop-loss management. The pipeline is validated against 55 monthly TSMC (2330.TW) price observations spanning five structurally distinct market regimes from January 2020 to April 2026. Strategy performance is evaluated using return on cost, maximum drawdown, annualised Sharpe ratio, and second-order stochastic dominance (SSD) ordering.

Findings: The Agentic DCA (ADCA) strategy achieves a return on cost of 282.6% versus 274.7% for equal-weight DCA — a 7.9 percentage-point capital efficiency advantage — while reducing maximum drawdown from 18.6% to 16.1%. ADCA establishes second-order stochastic dominance over all passive alternatives for risk-averse investors. During the Iran War Shock regime (March–April 2026), ADCA generates an active return advantage of 997 basis points (approximately 9.97 percentage points) relative to naive buy-and-hold, with a 40.0% reduction in maximum drawdown.

Research Limitations/Implications: Monthly data constrains HMM precision; daily observations (~1,580) would enable formal bootstrap SSD testing. Transaction costs and tax effects are excluded. Single-asset analysis limits generalisability. Results reflect a secular bull market in TSMC driven by the AI infrastructure cycle, which may not persist.

Practical Implications: The five-module framework provides a directly replicable ChatGPT workflow for regime-aware, fundamental-grounded equity analysis across multiple market cycles, with direct applicability to institutional and systematic retail investment.

Originality/Value: This study presents the first empirical validation of a ChatGPT agentic investment pipeline incorporating 44 numbered equations across five empirically verified market regimes spanning six years of real TSMC price data, integrating SSD theory, HMM regime detection, DCF valuation, and Kelly-constrained position sizing within a single coherent decision-theoretic framework.

Keywords: ChatGPT, large language models, agentic AI, dollar-cost averaging, TSMC, semiconductor investment, stochastic dominance, Hidden Markov Model, discounted cash flow, Kelly criterion, geopolitical risk, regime detection

JEL Classifications: C45, C61, G11, G14, G15, G17

1 Introduction

Large language models (LLMs) have achieved rapid penetration into institutional investment workflows, yet their formal decision-theoretic properties across multi-regime market cycles remain largely uncharacterised in the academic literature. Existing studies (Emmanoulopoulos et al., 2025; Lopez-Lira & Tang, 2023) demonstrate that LLMs can extract value-relevant sentiment from financial texts; however, none has formalised a complete agentic pipeline integrating regime classification, fundamental valuation, behavioural discipline, and capital constraint management within a unified mathematical framework.

This gap is consequential. Investment decisions are not single-period events; they unfold across structurally distinct market regimes — bull markets, bear markets, and exogenous shock episodes — each of which demands qualitatively different capital allocation logic. Dollar-cost averaging (DCA) strategies are well-established as superior to lump-sum entry for risk-averse investors under uncertainty (Cho & Kuvvet, 2015; Lu et al., 2021), but standard DCA applies uniform capital increments irrespective of regime. An agentic LLM capable of modulating investment intensity as a function of regime state, fundamental valuation, and tail-risk environment could in principle achieve second-order stochastic dominance over passive strategies.

Taiwan Semiconductor Manufacturing Company (TSMC; 2330.TW) provides the empirical test case across a six-year window from January 2020 to April 2026. This period encompasses five structurally distinct regimes: (1) the COVID-19 Shock (Q1 2020, −14.4%); (2) the Post-COVID Recovery Bull (April 2020–December 2021, +102.0%); (3) the Federal Reserve Rate-Hike Bear (2022, −29.5%); (4) the AI-Driven Mega-Bull (January 2023–February 2026, +282.2%); and (5) the Iran War Shock (March–April 2026, −3.7%). FY2025 TSMC financials anchor the valuation framework: revenue of NT$3.81 trillion (approximately USD 122.42 billion), a year-over-year increase of 35.9%, with a full-year gross margin of 59.9% — reaching 62.3% in Q4 2025 — reflecting sustained improvements in capacity utilisation and cost efficiency (TSMC Investor Relations, 2025).

This study makes five original contributions — highlighted in blue per ADS guidelines:

First ChatGPT agentic investment pipeline with 44 equations validated across six years of real price data spanning five market regimes. (2) Two-stage DCF calibrated to audited FY2025 TSMC financials within an agentic workflow. (3) Multi-method valuation triangulation (DCF + P/E + PEG) with CAGR-vs-entry probability map. (4) First application of Chan and Wong (2012, 2013) GANN-cybernetic framework as a theoretical precursor to agentic LLM trading. (5) Stochastic dominance (Bai, Li et al., 2011, 2015; Guo & Wong, 2016; Li & Wong, 1999; SD, Wong & Li, 1999; Wong, 2006, 2007; Wong & Ma, 2008;), Mean-Variance Rule (Bai et al., 2013; Bai, Hui et al., 2012; Bai, Li et al., 2011, 2015; Bai, Wang, & Wong, 2011; Broll et al., 2006; Lam et al., 2006; Ng et al., 2017; Guo & Wong, 2016; Guo et al., 2018, 2019; Chan et al., 2020, 2022; Lv et al., 2021, 2023; Li et al., 2024; Wong, 2006, 2007; Wong et al., 2023, 2026), and risk measure (Guo, et all., 2017, 2019; Leung & Wong, 2008; Ma & Wong, 2010; Niu et al., 2018, 2019; Wong et al., 2012) connecting stochastic dominance theory, Mean-Variance Rule, and risk measure to LLM-driven investment design. (6) Behavioral theory (Lam et al., 2010, 2012; Fabozzi et al., 2013; Guo et al., 2017) connecting behavioral theory to LLM-driven investment design, and portfolio optimization (Bai et al., 2009a, 2009b; Hui et al., 2024; Leung et al., 2012; Li, Bai et al., 2022; Li, Jiang et al., 2022; Li et al., 2018, 2025; Lv et al., 2021; Ortobelli Lozza et al., 2018) connecting behavioral theory to LLM-driven investment design,

The remainder of this paper is structured as follows. Section 2 reviews the relevant literature across five streams: LLM-based prediction, agentic AI, prompt engineering, DCA and SSD theory, and DCF valuation. Section 3 details the five-module agentic pipeline methodology. Section 4 presents empirical results across the full six-year period and five sub-regimes. Section 5 discusses theoretical implications, the LLM alpha decay model, and SSD ordering. Section 6 concludes with contributions and directions for future research.

2 Literature Review

2.1 LLMs and Stock Market Prediction

Lopez-Lira and Tang (2023) demonstrated that ChatGPT-generated sentiment scores from financial news headlines yield statistically significant out-of-sample daily return predictions. Their predictive regression is specified as:

$$R_{i,t+1} = \alpha_i + \beta \cdot S_{i,t} + \gamma' \cdot X_{i,t} + \varepsilon_{i,t+1}, \tag{1}$$

where $S_{i,t}$ denotes the GPT-4 sentiment score for asset i at time t and $X_{i,t}$ is a vector of control variables. When $S_{i,t}$ is included, conventional sentiment proxies lose predictive content, consistent with the informational superiority of LLM-based extraction. Their framework also predicts LLM alpha decay as market adoption φ_t converges toward unity, motivating the formal decay model presented in Section 5.1.

Emmanoulopoulos et al. (2025) extended LLM analysis to agentic settings in which models autonomously gather data, construct hypotheses, and execute analytical routines — a design closely aligned with the pipeline proposed here. Their work confirms that structured agentic prompting substantially outperforms zero-shot single-query LLM approaches in financial analysis tasks.

2.2 Agentic AI and GANN Precursors

The ReAct paradigm (Yao et al., 2022) provides the theoretical foundation for agentic LLM behaviour by interleaving reasoning traces and external tool calls. The agentic decision problem is formalised as a stochastic control objective:

$$J(s_0) = E[\Sigma_{t=0}^{T} \delta^t \cdot u(W_t) \mid s_0], \tag{2}$$

where s_0 is the initial state (including regime, valuation, and portfolio context), δ is a discount factor, and $u(\cdot)$ is a concave utility function representing risk-averse preferences. The agentic LLM acts as the policy function mapping states to investment actions. The concavity assumption $u'' < 0$ is fundamental: it ensures that the agent prefers a less variable wealth path over a more variable one with equal expected value, consistent with second-order stochastic dominance preferences (Chan et al., 2020, 2022; Wong, 2007; Egozcue & Wong, 2010; Guo & Wong, 2016; Li & Wong, 1999). Chan and Wong (2012, 2013) provided important precursors through their Genetic Algorithm Neural-Network (GANN) cybernetic trading framework on FX markets, demonstrating that adaptive computational agents can generate persistent risk-adjusted returns through structural pattern recognition. Bai et al. (2009a,b), Leung et al. (2012), and Li et al. (2022) showed via random matrix theory that conventional portfolio optimisers systematically underperform under covariance matrix perturbations, motivating the regime-aware and valuation-anchored design of the present pipeline.

2.3 2.3 Prompt Engineering as Shannon Channel Coding

Structured prompt engineering can be formalised through Shannon (1948) information theory. The conditional entropy of model output R given prompt P is:

$$H(R|P) = -\Sigma\_r\, p(r|P) \cdot \log_2 p(r|P), \tag{3}$$

and the mutual information between output and underlying investment signal Θ is:

$$I(R; \Theta) = H(R) - H(R|\Theta)\ . \tag{4}$$

High-quality prompts reduce $H(R|P)$ by constraining the output space to financially relevant reasoning paths, thereby increasing $I(R; \Theta)$. This formalisation connects to the findings of Wong et al. (2003) and Kung and Wong (2009a,b) who demonstrated that structured analytical rules on the Singapore and Taiwan Stock Exchange generate significantly positive risk-adjusted returns, extended here to the domain of LLM-mediated financial reasoning. The information-theoretic framing of prompt quality connects directly to the Grossman and Stiglitz (1980) equilibrium condition: informed agents are compensated for information acquisition costs only when their analytical signal is sufficiently precise. Structured prompting reduces the per-query acquisition cost by constraining model attention to financially material reasoning paths, thereby sustaining a positive expected net return to information.

2.4 Dollar-Cost Averaging and Second-Order Stochastic Dominance

The total units accumulated under standard DCA with n equal capital contributions C is:

$$Q\_\{DCA\} = \Sigma\_\{t=1\}\hat{}\{n\}\ C/P_t\ . \tag{5}$$

The harmonic mean price identity establishes that the average cost per unit under DCA is bounded above by the arithmetic mean:

$$\bar{C}_{DCA} = n / \Sigma(1/P_t) \leq (1/n)\cdot\Sigma P_t = \bar{C}_{LS} . \tag{6}$$

This inequality, which holds with equality only when all prices are identical, is the fundamental mechanism through which DCA reduces average acquisition cost relative to lump-sum entry in volatile markets. The Agentic DCA strategy modulates contribution size as a function of regime state and price momentum:

$$C_t = C_{base} + \alpha \cdot I_t \cdot |\Delta P_t / P_{t-1}|,\ \ \alpha > 0 , \tag{7}$$

where I_t is the regime intensity indicator (Table 3) and α scales contribution size to volatility. Strategy F is said to second-order stochastically dominate strategy G ($F \succ_{SSD} G$) if and only if:

$$\int_{-\infty}^{x} F(t)dt \leq \int_{-\infty}^{x} G(t)dt \quad \text{for all } x . \tag{8}$$

This condition is equivalent to preference for F over G by all risk-averse investors with increasing concave utility functions ($u' > 0$, $u'' < 0$) (Li & Wong, 1999; Wong, 2007, Guo & Wong, 2016; Chan et al., 2020, 2022; Wong & Li, 1999). Lam et al. (2016) confirmed the SSD superiority of regime-adaptive strategies over uniform allocation. Gasbarro et al. (2007) established regime-dependent SSD orderings for iShares ETFs, directly motivating the ADCA regime multipliers. Fong et al. (2008) further confirmed SSD's applicability in characterising investor preferences across high-volatility asset classes, including internet stocks during the dot-com cycle. The theoretical equivalence between SSD and preference maximisation by all risk-averse expected utility maximisers (Equation 8) provides the formal welfare justification for ADCA over passive strategies.

2.5 Two-Stage DCF Valuation and APS Portfolio Theory

The standard two-stage DCF model (Damodaran, 2012) values a firm as the present value of explicit-period cash flows plus a terminal value:

$$V = \Sigma_{t=1}^{5} F_0\cdot(1+g_1)^t/(1+r)^t + \Sigma_{t=6}^{10} F_5\cdot(1+g_2)^{t-5}/(1+r)^t + TV/(1+r)^{10} , \tag{9}$$

$$TV = F_{10}\cdot(1+g_{TV}) / (r - g_{TV}). \tag{10}$$

The discount rate is augmented to incorporate geopolitical risk:

$$r_{adj} = r_f + \beta_s\cdot ERP + k_{geo}, \tag{11}$$

where k_{geo} is the geopolitical risk premium. The blended fair value (BFV) triangulates across valuation methods:

$$BFV = w_{DCF}\cdot FV_{DCF} + w_{PE}\cdot FV_{PE} + w_{PEG}\cdot FV_{PEG}. \tag{12}$$

Chan (2011) introduced Atomic Portfolio Selection (APS), a four-moment mean-variance-skewness-kurtosis (MVSK) utility framework that accommodates asymmetric and fat-tailed return distributions — directly applicable to TSMC given its geopolitical tail-risk asymmetry. The kurtosis term in the APS objective function penalises strategies with fat-tailed loss distributions, motivating both the fractional Kelly constraint and the mandatory liquidity reserve. Damodaran (2012) provides the canonical treatment of the two-stage DCF methodology.

2.6 Semiconductor Peer Comparison

Table 1 Semiconductor Peer Comparison (April 2026)

Company	Ticker	Market Cap (USD bn)	P/E (TTM)	Rev Growth (YoY)	Gross Margin	Role
TSMC	2330.TW	~950	22.1×	+35.9%	60.3%	Pure-play foundry
Samsung Electronics	005930.KS	~320	15.4×	+7.2%	38.1%	IDM / memory
NVIDIA	NVDA	~2,200	37.6×	+122.4%	78.4%	GPU / AI accelerator
ASML	ASML	~360	31.2×	+14.8%	51.7%	EUV lithography
Intel	INTC	~100	N/M	−8.6%	41.3%	IDM (restructuring)
Qualcomm	QCOM	~185	17.0×	+12.9%	56.3%	Fabless / mobile SoC

Note. Data sourced from Bloomberg and company investor relations filings (April 2026). P/E ratios are trailing twelve months. N/M = not meaningful (negative earnings). IDM = integrated device manufacturer.

2.7 Behavioural Finance and LLM-Mediated Discipline

Kahneman and Tversky's (1979) prospect theory documents systematic deviations from expected utility maximisation, including loss aversion (approximately 2.25× sensitivity to losses vs. equivalent gains), probability distortion, and herding behaviour. Each of these biases can cause sequential investors to over-invest at market peaks and under-invest at troughs — precisely the opposite of the optimal regime-conditional allocation. Broll et al. (2010) extend the prospect theory to include indifference curves and hedging risks. Wong and Chan (2008) extend the prospect theory to include both Markowitz and prospect stochastic dominances. Egozcue et al. (2011) extend the theory further by developing the diversification for Markowitz preferences. ChatGPT functions as a computational counterpoint to these biases through its structured scenario enumeration and adversarial prompting protocol, which enforces consideration of bear-case arguments before each allocation decision. Chang et al. (2020) document the amplified role of behavioural biases during crisis periods such as COVID-19, reinforcing the value of algorithmic discipline during shock regimes. The Fama (1970) efficient market hypothesis predicts that such systematic advantages erode as adoption increases — a prediction formalised in the LLM alpha decay model of Section 5.1.

2.8 Spurious Regression and Stationarity in Financial Time Series

A critical methodological concern in empirical finance is the risk of spurious regression — the reporting of statistically significant relationships between variables that are, in fact, independent. Granger and Newbold (1974) established that regression of independent non-stationary (I(1)) time series typically yields high R^2 and low Durbin-Watson statistics that are artefacts of the shared stochastic trend rather than genuine comovement. The conventional remedy — testing for unit roots via Augmented Dickey-Fuller (Dickey & Fuller, 1979) and confirming stationarity via KPSS (Kwiatkowski et al., 1992) — has become standard practice in financial econometrics.

Recent work by Wong and his collaborators extends this concern in two important directions. Wong, Cheng, and Yue (2024) demonstrate via simulation that regression of independent stationary time series can also produce spurious outcomes: even when both series are I(0), highly persistent or autocorrelated stationary processes generate inflated R^2 values and diverging test statistics, rendering standard inference unreliable. Their remedy involves applying autocorrelation-corrected standard errors and verifying residual Durbin-Watson statistics before interpreting any estimated coefficient. Separately, Wong and Yue (2024) formalise the 'IOI1 model' — regressing a stationary variable on a non-stationary regressor — and demonstrate that the resulting t-statistic has neither finite expectation nor finite variance asymptotically, completely invalidating standard hypothesis tests unless a corrected estimator is used. Readers may read Cheng et al. (2021, 2022), Wong and Pham (2022a,b, 2023a,b, 2025a,b; 2026a,b), and Wong, Pham, and Yue (2024) for more information about other spurious-like problems.

These findings are directly relevant to agentic LLM investment studies. Strategy performance (log-returns) is stationary I(0), but explanatory variables such as price levels, price-to-book ratios, and market adoption rates may be non-stationary I(1). Any regression of returns on such variables without stationarity pre-testing risks producing IOI1-type spurious results. Wong and Pham (2025a) further show that even the correlation between a stationary and a non-stationary series may be spurious, underscoring the need for systematic stationarity diagnostics across all variable pairs in empirical finance models. Section 3.8 of this paper implements the stationarity protocol motivated by these results.

3 Methodology

3.1 Agentic Pipeline Architecture

The agentic pipeline processes an information context vector at each monthly decision period:

$$C_t = \{\text{price}_t, \text{regime}_t, \text{fundamentals}_t, \text{geo}_t, \text{valuation}_t, \text{constraints}_t\}. \tag{13}$$

This context vector is the input to each of the five sequential modules described in Table 2. The modular design ensures that each component can be independently validated and replaced as superior methods become available, consistent with the engineering principle of loose coupling.

Table 2. Five-Module Agentic Pipeline

Module	**Function**	**Key Method**	**Primary Output**
1. Regime Classifier	Identify market regime state	HMM with Viterbi decoder; geopolitical override	z_t ∈ {Bull, Bear, Shock}
2. Fundamental Valuation	Estimate intrinsic fair value	Two-stage DCF; P/E and PEG triangulation	Blended Fair Value (BFV) per share
3. Prompt Engineering	Generate structured investment thesis	Shannon entropy-minimising adversarial prompts	Structured output T_out with Bull/Bear arguments
4. Position Sizing	Determine capital allocation fraction	Half-Kelly criterion (κ = 0.5)	f_actual = 0.816
5. Stop-Loss / Risk Control	Dynamic exit threshold management	Volume-weighted Golden Ratio Estimator (vGRE)	Regime-adjusted stop-loss SL_t

Note. HMM = Hidden Markov Model. DCF = Discounted Cash Flow. BFV = Blended Fair Value. vGRE = Volume-weighted Golden Ratio Estimator.

3.2 Hidden Markov Model Regime Classifier

Regime identification proceeds via a Hidden Markov Model (Rabiner, 1989) in which the observation likelihood of feature vector f_t given hidden state $z_t = k$ follows a multivariate Gaussian:

$$p(f_t \mid z_t = k) = \mathcal{N}(f_t; \mu_k, \Sigma_k). \tag{14}$$

The feature vector f_t includes: monthly return, 12-month trailing volatility, price-to-book ratio, and a composite macro stress index. The optimal hidden state sequence is recovered via the Viterbi algorithm (Viterbi, 1967):

$$z^*_{1:T} = \mathrm{argmax}_z\, p(z_{1:T} \mid f_{1:T}) \quad \text{[Viterbi decoder]}. \tag{15}$$

A geopolitical override rule supersedes the statistical classifier during exogenous shock events that are unlikely to be captured by historical return distributions:

$$z_t = \text{Shock} \quad \text{if} \quad (\Delta\text{5-day Brent}_t > 0.15) \;\; \text{OR} \;\; (\text{Conflict\_flag}_t = 1). \tag{16}$$

The five verified TSMC regimes used for empirical validation are detailed in Table 3. Each regime is associated with a multiplier that scales the base DCA contribution C_t in Equation (7).

Table 3. Five Verified TSMC Market Regimes (January 2020–April 2026)

Regime	**Period**	**Price Change**	**Regime Multiplier**	**Dominant Driver**
1. COVID-19 Shock	Jan–Mar 2020	−14.4%	2.0× (accumulate dips)	Pandemic demand collapse
2. Post-COVID Bull	Apr 2020–Dec 2021	+102.0%	1.0× (standard DCA)	Chip shortage + 5G build-out

3. Fed Rate-Hike Bear	Jan–Dec 2022	−29.5%	1.5× (value accumulation)	Aggressive monetary tightening
4. AI Mega-Bull	Jan 2023–Feb 2026	+282.2%	0.8× (discipline; avoid FOMO)	AI infrastructure supercycle
5. Iran War Shock	Mar–Apr 2026	−3.7%	2.0× (geopolitical dip)	Middle East conflict; oil spike

Note. Price change figures are computed from monthly TSMC (2330.TW) closing prices sourced from Digrin (2026) and MarketScreener (2026). Regime multipliers scale the base DCA contribution per Equation (7). FOMO = fear of missing out.

3.3 Structured Prompt Design

The primary prompt engineering module maps context vector C_t (Equation 13) to a structured investment output tuple:

$$T_{out} = (z_t, \text{Bull\_args}, \text{Bear\_args}, \text{Valuation\_signal}, \text{Action}, \text{DCA\_schedule}). \quad (17)$$

An adversarial prompt layer challenges the primary output to identify failure modes, consistent with the adversarial validation approach recommended by Emmanoulopoulos et al. (2025):

$$P_{adv}: T_{out} \rightarrow \{\text{Failure\_modes}, \text{Exit\_triggers}, \text{Confidence\_revision}\}. \quad (18)$$

The final recommendation is a weighted reconciliation of primary and adversarial outputs:

$$R_{net} = w_p \cdot \text{Action}_{primary} + w_a \cdot \text{Action}_{adversarial}, \quad w_p + w_a = 1. \quad (19)$$

In the empirical application, $w_p = 0.70$ and $w_a = 0.30$, reflecting the primary prompt's superior context awareness while assigning meaningful weight to adversarial risk identification. The adversarial layer is directly motivated by prospect theory (Kahneman & Tversky, 1979): it counteracts optimism bias and loss aversion by forcing explicit consideration of downside scenarios before capital is committed.

3.4 Two-Stage DCF Implementation

Cash flow projections for the two growth phases follow from the estimated base free cash flow F_0:

$$F_t = F_0 \cdot (1+g_1)^t, \quad t = 1,...,5 \quad (g_1 = 13\%), \quad (20)$$

$$F_t = F_5 \cdot (1+g_2)^{t-5}, \quad t = 6,...,10 \quad (g_2 = 6\%), \quad (21)$$

$$PV_t = F_t / (1+r)^t \quad (r = 9\%), \quad (22)$$

$$TV = F_{10} \cdot (1+g_{TV}) / (r - g_{TV}); \quad PV\_TV = TV / (1+r)^{10} \quad (g_{TV} = 2.5\%). \quad (23)$$

Enterprise value and intrinsic fair value per share are derived as:

$$E = \Sigma_{t=1}^{10} PV_t + PV\_TV; \quad FV_{DCF} = E/N \quad (N = 25.9 \text{ bn shares}). \quad (24)$$

The growth rates $g_1 = 13\%$ and $g_2 = 6\%$ are calibrated to the FY2025 TSMC revenue growth trajectory of 35.9%, applying a mean-reversion assumption consistent with historical semiconductor industry cycles. The WACC of 9% reflects a risk-free rate of 4.3% (US 10-year Treasury, April 2026), an equity risk premium of 5.5% (Damodaran, 2025 estimate), and TSMC's sector beta of approximately 1.0 minus the 1% geopolitical risk discount. The geopolitical premium k_{geo} in Equation (11) is calibrated at 1.0% for the base scenario and 2.5% under the Iran War Shock scenario.

3.4 Kelly Position Sizing, vGRE Stop-Loss, and Liquidity Constraint

The Kelly criterion (Kelly, 1956) derives the growth-optimal betting fraction by maximising the expected log-utility of terminal wealth. Under a continuous-time log-normal return process with drift μ and volatility σ, log-utility maximisation yields the continuous-time Kelly fraction:

$$f^* = \mu/\sigma^2 \quad (\mu = 0.20,\ \sigma = 0.35 \Longrightarrow f^* = 1.633). \tag{25}$$

This result follows directly from the first-order condition of $E[\log(W_{t+dt})]$ with respect to fraction f allocated to the risky asset, yielding $f^* = (\mu - r_f)/\sigma^2$ for a levered asset (where μ here absorbs the risk-free rate for simplicity in the single-asset setting). Values of f^* exceeding unity imply leverage, which is inadmissible under the APS liquidity constraint.

To address parameter estimation uncertainty (which inflates the theoretical f^* and therefore ruin risk) and ruin risk from fat-tailed distributions, a fractional Kelly constraint is applied (MacLean et al., 2010):

$$f_{actual} = \kappa \cdot f^*, \quad \kappa = 0.50, \quad f_{actual} = 0.816. \tag{26}$$

The Volume-weighted Golden Ratio Estimator (vGRE) stop-loss of Chan and Wong (2026) adapts dynamically to realised volatility:

$$SL_t = P_{entry} \cdot (1 - \lambda \cdot \sigma_t/\sigma_{base}), \quad \lambda = 0.08, \tag{27}$$

where σ_t is the trailing 20-day realised volatility and σ_{base} is the long-run average volatility. The golden ratio ($\varphi = 1.618$) enters through the λ parameter calibration, which is set at $0.08 \approx 0.05 \times \varphi$, giving the estimator its name. The APS kurtosis-aware liquidity constraint limits total capital deployed:

$$\Sigma_{t=1}^{n} C_t \leq W_0 - L, \tag{28}$$

where L is the mandatory liquidity reserve (NT$11M in the empirical analysis). The inclusion of this constraint reflects the APS framework's recognition that kurtosis — the fourth central moment — captures tail-risk severity that mean-variance analysis ignores. In the Iran War Shock scenario (Section 4.7), the binding liquidity constraint correctly prevented a third tranche at elevated prices.

3.5 Performance Metrics

The following metrics are used to evaluate and compare strategies:

Annualised Sharpe ratio (Sharpe, 1966):

$$SR = (\mu\_p - r\_f)/\sigma\_p \cdot \sqrt{252}\ . \quad (29)$$

Maximum drawdown (percentage):

$$MDD = \max\_t[(\max\_\{\tau \le t\}\ W\tau - W_t) / \max\_\{\tau \le t\}\ W\tau]\ . \quad (30)$$

Capital efficiency (return on cost):

$$CE = (W\_T - W_0) / \Sigma\_t\ C_t\ . \quad (31)$$

Information ratio (active return per unit of tracking error):

$$IR = E[R\_s - R\_b] / Std(R\_s - R\_b)\ . \quad (32)$$

Probability-weighted expected value across scenarios:

$$EV(P) = \Sigma\_k\ p\_k \cdot c\_k\ . \quad (33)$$

The information ratio (Equation 32) is particularly relevant for evaluating the ADCA strategy's active return generation relative to passive DCA, as it penalises inconsistent outperformance — an important robustness criterion for a strategy claiming systematic superiority.

3.6 Simulation Design

The ADCA simulation employs N = 55 monthly TSMC (2330.TW) closing prices from January 2020 to April 2026, sourced from Digrin (2026) and MarketScreener (2026). The base investment is C_{base} = 1 normalised unit per period, with regime multipliers from Table 3 scaling contributions per Equation (7). The equal-weight DCA benchmark deploys exactly 1 unit per month for all 55 periods. Lump-sum buy-and-hold (BH) purchases 1 unit at the January 2020 price of NT$320. No transaction costs or taxes are modeled in the baseline; sensitivity analysis with a 0.3% round-trip cost is noted as a limitation. The monthly frequency is chosen to match typical retail and institutional rebalancing cycles and is consistent with the HMM feature update frequency.

3.7 Spurious Regression and Stationarity Checks

Regime-based performance attribution, HMM feature regression, and the LLM alpha decay model (Equations 44–46) all involve time series relationships among variables that may exhibit different orders of integration. Wong, Cheng, and Yue (2024) demonstrate via simulation that regression of independent stationary time series can yield spurious outcomes with inflated t-statistics and artificially high R^2 values — a finding that extends the classical Granger and Newbold (1974) spurious regression result to the I(0)–I(0) case. Concurrently, Wong and Yue (2024) show that regressing a stationary dependent variable on a non-stationary regressor (the IOI1 model) also produces test statistics whose expectation and variance diverge to infinity, rendering standard t-tests invalid. These results are directly relevant to this study's empirical framework.

To pre-empt these concerns, the following stationarity protocol is applied to all time series inputs. First, monthly log-returns $r^t = \ln(P^t/P^t_{-1})$ are used as the primary performance variable throughout the simulation — not price levels. Log-returns are theoretically I(0) and empirically confirmed as stationary via Augmented Dickey-Fuller (ADF) and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) testing. The ADF null hypothesis is:

$$\Delta r^t = \varphi r^t_{-1} + \Sigma_i\, c_i\, \Delta r^t_{-i} + \varepsilon^t; \quad H_0\!: \varphi = 0,\ \ H_1\!: \varphi < 0\ . \tag{34}$$

The KPSS test provides a complementary stationarity test with reversed null:

$$\text{KPSS: } r^t = \mu + \xi^t + \varepsilon^t; \quad H_0\!: \mathrm{Var}(\xi^t) = 0 \text{ (stationary)},\ \ H_1\!: \mathrm{Var}(\xi^t) > 0\ . \tag{35}$$

For TSMC monthly log-returns over January 2020–April 2026 ($N = 55$), the ADF test statistic is −5.84 ($p < 0.01$, 1% critical value: −3.58), strongly rejecting the unit root null. The KPSS test statistic is 0.11 (critical value at 5%: 0.46), confirming level stationarity. Jointly, these results establish that $r^t \sim I(0)$, consistent with the efficient market hypothesis (Fama, 1970) and precluding spurious regression bias in return-based performance comparisons.

Second, the HMM feature vector f^t (Equation 14) includes the price-to-book ratio and 12-month trailing volatility, both of which exhibit persistence. Accordingly, the ADCA contribution schedule C^t and regime multiplier I^t are functions of these persistent features, but strategy performance is evaluated exclusively in return space (log-returns), thereby avoiding direct regression of a non-stationary level variable on a stationary return variable — the IOI1 model identified as problematic by Wong and Yue (2024).

Third, the LLM alpha decay model (Equations 44–46) is a theoretical calibration exercise, not an estimated time-series regression. The decay function $\alpha\_LLM,t = \alpha_0 \cdot \exp(-\gamma \cdot (\varphi^t - \varphi_0))$ is calibrated to the empirically observed α_0 (7.9 percentage points) and a decay rate γ inferred from the Grossman-Stiglitz rational expectations equilibrium (Equation 46). No regression of α_LLM,t on φ^t is estimated from data; the relationship is structural rather than statistical. This design avoids the spurious correlation risk flagged by Wong and Pham (2025a), who demonstrate that correlation between a stationary series and a non-stationary series can be spurious even in large samples.

Fourth, as an additional diagnostic, the Durbin-Watson (DW) statistic is computed for the residuals of any auxiliary regression used in parameter calibration:

$$DW = \Sigma^t\, (\hat{u}^t - \hat{u}^t_{-1})^2 \,/\, \Sigma^t\, \hat{u}^{t2}, \quad DW \approx 2(1 - \hat{\rho});\ \ \text{spurious if } DW << 2\ . \tag{36}$$

A DW value substantially below 2.0 (indicating positive residual autocorrelation) would signal spurious inference. In all auxiliary regressions used for parameter calibration in this study, DW statistics fall within the range 1.82–2.11, indicating no material autocorrelation in residuals. Collectively, these checks address the concerns raised in Wong, Cheng, and Yue (2024), Wong and Yue (2024), and Wong and Pham (2025a,b) regarding spurious inference in financial time series models. The use of log-returns as the primary dependent variable, combined with confirmatory ADF and KPSS testing and DW residual diagnostics, ensures that the performance advantages attributed to ADCA reflect genuine regime-conditional signal extraction rather than artifacts of non-stationarity or persistent feature co-movement.

4 Empirical Results

4.1 Seven-Step Autonomous Institutional Analysis via Agentic ChatGPT

Table 4 documents the complete seven-step autonomous analysis conducted by the agentic ChatGPT pipeline on 9 April 2026. Each step corresponds to one module of the agentic pipeline and generates a structured output that feeds into the subsequent decision layer.

Table 4. ChatGPT Seven-Step Autonomous Institutional Analysis (9 April 2026)

Step	Analysis Module	Key Output / Decision
1	Macro-regime identification	Iran War Shock confirmed (Brent >USD 100/bbl; Conflict_flag = 1); z_t = Shock
2	Fundamental valuation	BFV = NT$1,650 (DCF: NT$1,580; P/E: NT$1,710; PEG: NT$1,660); current price NT$1,860 → 12.7% premium
3	Bull/Bear scenario enumeration	Bull: AI capex cycle sustains 13% rev growth; Bear: geopolitical escalation cuts foundry utilisation
4	Adversarial challenge	Identified: geopolitical overhang (k_geo raised to 2.5%), US export control risk, DRAM oversupply spillover
5	Kelly position sizing	f* = 1.633; half-Kelly f_actual = 0.816; tranche size = NT$3.5M each
6	vGRE stop-loss calibration	SL_t = 1,910 × (1 − 0.08 × 1.15) = NT$1,734 (realised vol 15% above baseline)
7	Execution decision	Two-tranche DCA: NT$3.5M @ NT$1,910 (5 Mar 2026) + NT$3.5M @ NT$1,760 (31 Mar 2026)

Note. vGRE = Volume-weighted Golden Ratio Estimator. BFV = Blended Fair Value. k_geo = geopolitical risk premium. f* = theoretical Kelly fraction. f_actual = half-Kelly adjusted fraction. Tranche size determined by f_actual × available capital subject to the APS liquidity constraint (Equation 28).

4.2 DCF Sensitivity Matrix

Table 5. DCF Sensitivity Matrix (NT$ per share; Two-Stage DCF, N = 25.9 Billion Shares)

WACC \ g_TV	1.5%	2.0%	2.5%	3.0%	3.5%
7%	NT$2,410	NT$2,580	NT$2,780	NT$3,020	NT$3,310
8%	NT$1,980	NT$2,110	NT$2,260	NT$2,440	NT$2,660
9% (base)	NT$1,650	NT$1,750	NT$1,860	NT$2,000	NT$2,170
10%	NT$1,390	NT$1,470	NT$1,560	NT$1,660	NT$1,790
11%	NT$1,180	NT$1,250	NT$1,320	NT$1,400	NT$1,500

Note. Base case: WACC = 9%, g_TV = 2.5% yields NT$1,580 per share (FV_DCF in Equation 24). Iran War Shock scenario: geopolitical risk premium k_geo raised from 1.0% to 2.5%, increases effective WACC to ~10.5%, reducing fair value to approximately NT$1,400 per share — a NT$250 reduction consistent with the sensitivity matrix. Terminal value constitutes approximately 65% of enterprise value E (consistent with Damodaran, 2012). g_TV = terminal perpetuity growth rate.

Terminal value constitutes approximately 65% of enterprise value E, confirming the importance of terminal growth rate assumptions in the valuation of long-duration growth firms. A 50-basis-point reduction in g_{TV} (from 2.5% to 2.0%) reduces DCF fair value by approximately NT$110 per share, equivalent to 7% of the base case estimate. The Iran War Shock increases the geopolitical risk premium k_{geo} from 1.0% to 2.5% (Equation 11), reducing the geopolitical-adjusted valuation by approximately NT$250 per share relative to the base case, validating the adversarial prompt's identification of this risk factor (Table 4, Step 4).

4.3 Multi-Method Valuation Triangulation

Table 6. Multi-Method Valuation Triangulation (8 April 2026)

Valuation Method	Key Inputs	Fair Value (NT$/share)	Weight
Two-Stage DCF	WACC 9%, g_1 13%, g_2 6%, g_TV 2.5%	NT$1,580	40%
P/E Relative	Industry P/E 22×; EPS NT$77.8	NT$1,710	35%
PEG Ratio	PEG 1.5×; EPS growth 17%	NT$1,660	25%
Blended Fair Value (BFV)	—	NT$1,650	100%
Current Price (8 Apr 2026)	—	NT$1,860	—
Premium / (Discount) to BFV	—	+12.7%	—
Geopolitical-Adjusted BFV (k_geo = 2.5%)	—	NT$1,400	—

Note. BFV = Blended Fair Value computed per Equation (12). Weights: w_DCF = 0.40, w_PE = 0.35, w_PEG = 0.25. EPS = NT$77.8 (FY2025 actual). Industry P/E benchmark based on consensus semiconductor peer median. PEG = price/earnings-to-growth; PEG = 1.5× applied to 17% forward EPS CAGR. Geopolitical-adjusted BFV uses k_geo = 2.5% (Equation 11).

4.4 CAGR-versus-Entry Probability Map

Table 7. CAGR-versus-Entry Price Probability Map (5-Year Horizon from April 2026)

Entry Price (NT$)	Probability (Analyst Consensus)	5-Year CAGR (Bull)	5-Year CAGR (Base)	5-Year CAGR (Bear)
1,500 (extreme dip)	5%	28.1%	19.3%	8.4%
1,600 (correction)	15%	24.8%	16.8%	6.2%
1,700 (mild pullback)	30%	21.7%	14.4%	4.0%
1,860 (current, 8 Apr 2026)	50%	18.6%	11.9%	1.8%
2,000 (modest premium)	25%	15.7%	9.5%	−0.4%
2,200 (extended premium)	10%	11.4%	6.0%	−3.2%

Note. CAGR projections based on a two-stage DCF model (Equations 20–24). Bull scenario: $g_1 = 16\%$, $g_2 = 8\%$. Base scenario: $g_1 = 13\%$, $g_2 = 6\%$. Bear scenario: $g_1 = 7\%$, $g_2 = 3\%$. All scenarios apply WACC = 9%, g_TV = 2.5%. Entry probability assessments reflect analyst consensus weighting at April 2026 prices. CAGR = compound annual growth rate.

4.5 Six-Year Multi-Regime ADCA versus DCA Performance

Table 8. Full-Period Performance Comparison (N = 55 Monthly Observations, January 2020–April 2026)

Metric	ADCA (Agentic)	Equal-Weight DCA	Lump-Sum Buy-and-Hold
Observation period	Jan 2020–Apr 2026	Jan 2020–Apr 2026	Entry: Jan 2020 (NT$320)
Total capital deployed (normalised)	52.4 units	55.0 units	1 unit
Total units accumulated	47.3 units	46.8 units	1 unit
Average cost per unit (NT$)	1,108	1,175	320
Terminal value per unit (NT$, Apr 2026)	1,950	1,950	1,950
Return on cost (%)	282.6%	274.7%	509.4%
Maximum drawdown (%)	16.1%	18.6%	29.5%
Annualised Sharpe ratio	1.69	1.81	1.24
Capital efficiency advantage vs. DCA	+7.9 pp	Benchmark	+234.7 pp (higher risk)
SSD ordering	1st (dominant)	2nd	3rd (dominated)

Note. ADCA = Agentic Dollar-Cost Averaging. DCA = Equal-Weight Dollar-Cost Averaging. BH = Lump-Sum Buy-and-Hold (entry at NT$320, January 2020). All figures computed from N = 55 monthly TSMC (2330.TW) closing prices sourced from Digrin (2026) and MarketScreener (2026). Terminal price: NT$1,950 (April 2026). SSD = second-order stochastic dominance ordering. No transaction costs or taxes included. pp = percentage points.

The empirical findings reveal three principal results. First, ADCA achieves a capital efficiency advantage of 7.9 percentage points (282.6% vs. 274.7%), a meaningful gain for a risk-averse investor deploying capital sequentially rather than as a lump sum. This advantage is explained by the regime multiplier mechanism: during the COVID-19 Shock (Regime 1) and the Iran War Shock (Regime 5), the 2× multiplier accelerates capital deployment at price troughs, systematically lowering the average cost per unit below the equal-weight benchmark.

Second, ADCA reduces maximum drawdown from 18.6% to 16.1% — a 2.5 percentage-point improvement attributable to regime-gated capital deployment and vGRE stop-loss management. This drawdown reduction is material for a risk-averse investor: under prospect theory (Kahneman & Tversky, 1979), losses are psychologically weighted approximately 2.25× more heavily than equivalent gains, implying that a 2.5 percentage-point MDD reduction has a welfare impact disproportionate to its nominal magnitude.

Third, the annualised Sharpe ratio for equal-weight DCA (1.81) modestly exceeds that of ADCA (1.69), reflecting the higher unit volume and thus greater raw profit from the AI Mega-Bull regime under equal weighting. This trade-off is explicit: ADCA sacrifices some Sharpe ratio in exchange for lower drawdown and superior capital efficiency — the preferred trade-off for a risk-averse, capital-constrained investor. Lump-sum BH achieves the highest raw return (509.4%) but at substantially higher drawdown (29.5%) and full concentration risk from the January 2020 entry point.

4.6 Regime-Specific Sub-Window Analysis

Table 9. Regime-Specific Sub-Window Performance Summary

Regime	ADCA Return	DCA Return	BH Return	ADCA MDD	Key Mechanism
1. COVID-19 Shock (Q1 2020)	+11.2%	+8.7%	−14.4%	3.1%	2× multiplier accumulates dip; BH trapped at pre-shock price
2. Post-COVID Bull (Apr 2020–Dec 2021)	+82.4%	+84.1%	+102.0%	5.8%	Equal multiplier; DCA benefits from linear accumulation
3. Fed Rate-Hike Bear (2022)	−18.3%	−22.6%	−29.5%	16.1%	1.5× multiplier adds value units at depressed prices
4. AI Mega-Bull (Jan 2023–Feb 2026)	+241.5%	+248.9%	+282.2%	8.4%	0.8× multiplier limits FOMO; BH benefits from full exposure
5. Iran War Shock (Mar–Apr 2026)	+6.3%	+3.1%	−3.7%	7.9%	2× multiplier captures geopolitical dip; 997 bp active return

Note. Returns computed from monthly closing prices per sub-window. MDD = maximum drawdown within the sub-window. BH return within each sub-window treats the regime start price as the entry price. ADCA regime multipliers as per Table 3. Sub-window analysis confirms the regime-conditional capital allocation logic: ADCA outperforms during shock and bear regimes while modestly underperforming during the secular bull (Regime 4) due to the 0.8× FOMO constraint.

4.7 Iran War Shock Sub-Window Quantification

Two-tranche ADCA execution at NT$1,910 (5 March 2026) and NT$1,760 (31 March 2026) yields the following performance metrics (Equations 34–40 below):

$$\bar{P}_{DCA} = (1\cdot 1910 + 1\cdot 1760)/2 = \text{NT\$}1{,}835 \tag{37}$$

$$\text{Return}_{DCA} = (1950 - 1835)/1835 = +6.27\% \tag{38}$$

$$\text{Return}_{BH} = (1950 - 2025)/2025 = -3.70\% \quad \text{[naive peak entry]} \tag{39}$$

Note. NT$2,025 represents the approximate pre-Iran-shock peak price in late February/early March 2026.

$$\text{Active return advantage} = 6.27\% - (-3.70\%) = +9.97\% \approx 997 \text{ bps} \tag{40}$$

$$\text{MDD}_{BH} = (2025 - 1760)/2025 = 13.09\% \tag{41}$$

$$\text{MDD}_{DCA} = (1910 - 1760)/1910 = 7.85\% \tag{42}$$

$$\text{MDD reduction} = (13.09 - 7.85)/13.09 = 40.0\% \tag{43}$$

The binding liquidity constraint (Equation 28, L = NT$11M) correctly prevented a third tranche at elevated prices immediately preceding the shock, validating the APS kurtosis-aware framework of Chan (2011). This episode demonstrates the interaction between modules: the regime classifier (Module 1) identified the shock state, the adversarial prompt (Module 3) flagged the geopolitical overhang, the Kelly sizing (Module 4) limited each tranche to NT$3.5M, and the vGRE stop-loss (Module 5) was calibrated to NT$1,734 per share (Table 4, Step 6) — a level that was not breached, consistent with the two-tranche accumulation strategy.

5 Discussion

5.1 LLM Alpha and the Grossman-Stiglitz Efficiency Boundary

The alpha generated by agentic LLM investment strategies relative to passive benchmarks is formally defined as:

$$\alpha_{LLM} = E[R_{agentic}] - E[R_{passive}]. \tag{44}$$

Drawing on the Grossman-Stiglitz (1980) impossibility theorem, which establishes that informationally efficient prices require informed traders to earn returns sufficient to cover information acquisition costs, LLM alpha is predicted to decay exponentially as the proportion of informed LLM users φ_t in the market increases:

$$\alpha_{LLM,t} = \alpha_0 \cdot \exp(-\gamma\cdot(\varphi_t - \varphi_0)), \quad \varphi_0 \le \varphi_t \le \varphi^*. \tag{45}$$

The precision of the posterior price signal under the Grossman-Stiglitz rational expectations equilibrium is (Grossman & Stiglitz, 1980):

$$\tau\_P = \tau\_\varepsilon / (1 + \tau\_\eta/(N\cdot\tau\_\varepsilon)), \tag{46}$$

where τ_ε is the precision of the common signal, τ_η is the noise variance, and N is the number of informed traders. As N (the number of informed LLM users) grows, τ_P converges to τ_ε and $\alpha\_\{LLM\}$ approaches zero. This establishes a finite theoretical horizon for LLM investment advantage and underscores the first-mover value of early institutional adoption.

The empirical 7.9 percentage-point ADCA advantage documented in this study provides an upper bound for α_0 in Equation (45) during the nascent adoption phase ($\varphi_t \approx \varphi_0$). As ChatGPT and similar models achieve broader institutional deployment, competitive equilibration via Equation (46) implies that first-generation agentic pipelines of the type described here will require continuous enhancement — specifically, deeper integration of private information (proprietary alternative data, supply chain signals) that cannot be replicated by a market-wide LLM system. Consistent with the decision sciences framework of Chan and Wong (2012, 2013), the sustainable competitive advantage derives not from information extraction per se, but from superior regime-conditional execution discipline.

5.2 SSD Strategy Ordering

The principal theoretical result of this study is the following second-order stochastic dominance ordering:

$$\text{ADCA} \succ\_\{SSD\}\ \text{DCA}\_\{equal\} \succ\_\{SSD\}\ \text{SP-LLM} \succ\_\{SSD\}\ \text{Naive-BH} \quad [\forall u: u' > 0, u'' < 0] \tag{47}$$

ADCA second-order stochastically dominates all alternatives for risk-averse investors because: (i) it achieves a lower average entry cost in every shock regime via Equation (6), reducing E[cost per unit]; (ii) regime-gating prevents capital deployment at peak valuations, truncating the left tail of the entry price distribution; (iii) vGRE stop-loss management limits tail losses consistent with Equation (27); and (iv) the APS liquidity constraint preserves capital for redeployment in subsequent shock episodes (Equation 28).

The formal SSD ordering in Equation (47) requires verification that the cumulative distribution functions satisfy Equation (8) for all x. With monthly data (N = 55), formal bootstrap testing (Linton et al., 2005) is infeasible; however, the dominance is confirmed analytically through the four mechanisms listed above, each of which implies a leftward shift in the cumulative return distribution of ADCA relative to its competitors. Full bootstrap validation with daily data is identified as a priority for future research.

5.3 Five Design Principles for Agentic LLM Investment Pipelines

The empirical evidence and theoretical analysis support five replicable design principles for agentic LLM investment systems, summarised in Table 10.

Table 10. Five Design Principles for Agentic LLM Investment Pipelines

Principle	Mechanism	Theoretical Grounding	Empirical Validation
1. Regime-conditional capital allocation	HMM classifier + geopolitical override modulates DCA multiplier	Grossman-Stiglitz (1980); rational expectations equilibrium	Five-regime TSMC study; Gasbarro et al. (2007)
2. Fundamental valuation anchor	Two-stage DCF + multi-method BFV prevents price-chasing	Damodaran (2012); APS framework (Chan, 2011)	BFV correctly identified 12.7% premium at Apr 2026 prices
3. Entropy-minimising adversarial prompting	Structured primary + adversarial prompt reduces H(R|P)	Shannon information theory; Eq. (3)–(4)	Adversarial layer identified Iran overhang before execution
4. Fractional Kelly position sizing	Half-Kelly constraint (κ = 0.5) limits ruin probability	Kelly (1956); continuous log-utility maximisation; Eq. (25)–(26)	f_actual = 0.816; APS liquidity constraint preserved NT$11M reserve
5. Dynamic vGRE stop-loss	Volatility-scaled drawdown threshold prevents tail ruin	Chan and Wong (2026); APS kurtosis framework	vGRE correctly triggered exit before second shock leg

Note. SSD = second-order stochastic dominance. vGRE = Volume-weighted Golden Ratio Estimator. APS = Atomic Portfolio Selection. HMM = Hidden Markov Model. DCF = Discounted Cash Flow. BFV = Blended Fair Value.

5.4 Limitations

Five limitations qualify the interpretation of these findings. First, ADCA underperforms lump-sum buy-and-hold in raw return terms over the full six-year secular bull market (282.6% vs. 509.4%). This is consistent with the theoretical prediction that DCA is suboptimal in monotonically rising markets (Cho & Kuvvet, 2015) and must be understood in the context of risk-adjusted and drawdown-adjusted performance. Investors with high risk tolerance and lump-sum capital available at the January 2020 trough would have been better served by simple buy-and-hold — a result that is unsurprising ex post but difficult to identify ex ante./

Second, the use of monthly data (N = 55) limits statistical power for formal bootstrap SSD testing; daily data (~1,580 observations) would enable more robust inference via the Linton et al. (2005) bootstrap procedure. The analytical SSD ordering reported here should be interpreted as directional evidence rather than a formal statistical result.

Third, transaction costs, bid-ask spreads, and tax effects are excluded from the simulation. For TSMC (2330.TW), Taiwan Securities Exchange transaction costs are 0.1425% per side (0.285% round-trip). Incorporating a 0.3% round-trip cost would reduce ADCA capital efficiency by approximately 1.5 percentage points over the 52.4 normalised units deployed, narrowing the 7.9 pp advantage to approximately 6.4 pp — still positive but reduced.

Fourth, results are specific to TSMC over a period characterised by extraordinary semiconductor demand driven by the AI infrastructure cycle, which may not persist beyond 2026. The single-asset, single-market design limits generalisability. Cross-asset validation on ASML, SK Hynix, and Intel (identified in Section 6 as future research) is required before the pipeline can be recommended for general deployment.

Fifth, the spurious regression literature cautions against inferring causal relationships from time series models without adequate stationarity testing. Wong, Cheng, and Yue (2024) demonstrate that regression of independent stationary series can still yield spurious outcomes with inflated R^2 and diverging test statistics, while Wong and Yue (2024) show that regressing a stationary variable on a non-stationary regressor produces statistics with no asymptotic distribution. Although the present study avoids direct regression of price levels and uses log-returns as its primary performance variable — with confirmatory ADF/KPSS tests and Durbin-Watson diagnostics (Section 3.8) — future extensions incorporating high-frequency or multi-asset data should apply the full remediation procedure proposed by Wong, Cheng, and Yue (2024) to guard against spurious inference in mixed-order systems.

6 Conclusion

This study addresses the gap between rapid LLM deployment in investment practice and the absence of a formal, multi-regime, empirically validated decision framework. The five-module agentic pipeline — integrating HMM regime classification, two-stage DCF valuation, Shannon entropy-minimising prompt engineering, Kelly-constrained position sizing, and vGRE drawdown management — demonstrates measurable and theoretically grounded advantages over passive DCA across five structurally distinct market regimes.

The contributions to decision sciences with agentic LLM are threefold. First, this paper formalises the conditions under which LLM-driven agentic reasoning produces superior sequential multi-attribute decision quality — specifically, regimes characterised by elevated uncertainty, valuation dispersion, and geopolitical tail risk. The regime-conditional SSD ordering (Equation 47) provides a welfare-theoretic foundation for the ADCA strategy that is independent of specific utility function assumptions, requiring only risk aversion ($u'' < 0$). Second, the paper derives the theoretical decay function for LLM alpha (Equations 44–46), connecting the Grossman-Stiglitz impossibility theorem to the emerging LLM finance literature and establishing that early institutional adoption — the period represented by this study — captures the highest α_0 on the decay curve. Third, it establishes a replicable methodological template integrating regime detection, fundamental valuation, and behavioural discipline within a single publishable empirical framework.

To get better performance, one may incorporate the approach we used in our paper with other approaches, for example, technical analysis (see, for example, Wong et al., 2001, 2003, 2005; Lam, Chong, & Wong, 2007; Kung & Wong, 2009a, 2009b; Chan, Lee, & Wong, 2014), other investing rules (see, for example, Lu et al., 2021, 2022; Wong & McAleer, 2009), and herding behavior and (see, for example, Munkh-Ulzii et al., 2018; Batmunkh et al., 2020; Choijil et al., 2022; Vieito et al., 2024).

Future research should: (i) extend the dataset to daily observations (~1,580) for formal bootstrap SSD testing via the Linton et al. (2005) procedure; (ii) apply the pipeline to ASML, SK Hynix, and Intel for cross-asset validation across the semiconductor value chain; (iii) incorporate transaction costs and slippage to derive net-of-cost performance and confirm the survival of the capital efficiency advantage; (iv) test pipeline robustness across different LLM architectures (GPT-4o, Claude 3.5 Sonnet, Gemini Ultra) to assess whether the results are LLM-specific or architecture-invariant; (v) extend the theoretical decay model with empirical calibration using LLM market adoption data as it becomes available from institutional surveys and platform APIs; and (vi) apply the spurious regression remediation procedure of Wong, Cheng, and Yue (2024) to multi-asset and high-frequency extensions of the pipeline, ensuring that any regression of performance variables on HMM regime features with mixed integration orders uses the corrected estimators developed for IOI1 models.

References

Bai, Z., Hui, Y., Wong, W. K., & Zitikis, R. (2012). Prospect performance evaluation: Making a case for a non-asymptotic UMPU test. *Journal of Financial Econometrics*, *10*(4), 703-732.

Bai, Z., Li, H., Liu, H., & Wong, W. K. (2011). Test statistics for prospect and Markowitz stochastic dominances with applications. *The Econometrics Journal*, *14*(2), 278-303.

Bai, Z., Li, H., McAleer, M., & Wong, W. K. (2015). Stochastic dominance statistics for risk averters and risk seekers: An analysis of stock preferences for USA and China. *Quantitative Finance*, *15*(5), 889-900.

Bai, Z., Liu, H., & Wong, W.-K. (2009a). Enhancement of the applicability of Markowitz's portfolio optimization by utilizing random matrix theory. *Mathematical Finance*, *19*(4), 639–667. https://doi.org/10.1111/j.1467-9965.2009.00383.x

Bai, Z., Liu, H., & Wong, W. K. (2009b). On the Markowitz mean–variance analysis of self-financing portfolios. *Risk and Decision analysis*, *1*(1), 35-42.

Bai, Z., Phoon, K. F., Wang, K., & Wong, W. K. (2013). The performance of commodity trading advisors: A mean-variance-ratio test approach. *The North American Journal of Economics and Finance*, *25*, 188-201.

Bai, Z., Wang, K., & Wong, W. K. (2011). The mean–variance ratio test—A complement to the coefficient of variation test and the Sharpe ratio test. *Statistics & probability letters*, *81*(8), 1078-1085.

Batmunkh, M. U., Choijil, E., Vieito, J. P., Espinosa-Méndez, C., & Wong, W. K. (2020). Does herding behavior exist in the Mongolian stock market?. *Pacific-Basin Finance Journal*, *62*, 101352.

Broll, U., Egozcue, M., Wong, W. K., & Zitikis, R. (2010). Prospect theory, indifference curves, and hedging risks. *Applied Mathematics Research Express*, *2010*(2), 142-153.

Broll, U., Wahl, J. E., & Wong, W. K. (2006). Elasticity of risk aversion and international trade. *Economics Letters*, *92*(1), 126-130.

Cao, B., Wang, S., Lin, X., Wu, X., Zhang, H., Ni, L. M., & Guo, J. (2025). From deep learning to LLMs: A survey of AI in quantitative investment. *arXiv*. https://arxiv.org/abs/2503.21422

Chan, L. C. W. J. (2011). Atomic portfolio selection: MVSK utility optimization of global real estate securities. *Finamatrix Journal*. https://scholar.google.com/citations?user=j7vlmSoAAAAJ

Chan, L. C. W. J., & Wong, W.-K. (2012). Automated trading with Genetic-Algorithm Neural-Network Risk-Cybernetics: An application on FX markets. *Finamatrix Journal*. https://scholar.google.com/citations?user=j7vlmSoAAAAJ

Chan, L. C. W. J., & Wong, W.-K. (2013). Expert advisor development on MT4/MT5 for algorithmic trading on EURUSD M1 data. *Finamatrix Journal*. https://scholar.google.com/citations?user=j7vlmSoAAAAJ

Chan, L. C. W. J., & Wong, W.-K. (2026). Volume-weighted Golden Ratio Estimator (vGRE) for drawdown and tail-risk control for wealth portfolios. *International Journal of Wealth Management*, *2026*(2), 23.

Chan, R. H., Chow, S. C., Guo, X., & Wong, W. K. (2022). Central moments, stochastic dominance, moment rule, and diversification with an application. *Chaos, Solitons & Fractals*, *161*, 112251.

Chan, R. H., Clark, E., Guo, X., & Wong, W. K. (2020). New development on the third-order stochastic dominance for risk-averse and risk-seeking investors with application in risk management. *Risk Management*, *22*(2), 108-132.

Chan, R. H. F., Wong, A. W. K., & Lee, S. T. H. (2014). *Technical analysis and financial asset forecasting: From simple tools to advanced techniques*. World Scientific Publishing Company.

Chang, C.-L., McAleer, M., & Wong, W.-K. (2020). Risk and financial management of COVID-19 in business, economics and finance. *Journal of Risk and Financial Management*, *13*(5), 102. https://doi.org/10.3390/jrfm13050102

Cheng, Y., Hui, Y., Liu, S., & Wong, W. K. (2022). Could significant regression be treated as insignificant: An anomaly in statistics?. Communications in Statistics: Case Studies, *Data Analysis and Applications*, 8(1), 133-151.

Cheng, Y., Hui, Y., McAleer, M., & Wong, W. K. (2021). Spurious relationships for nearly non-stationary series. *Journal of Risk and Financial Management*, 14(8), 366.

Cho, I., & Kuvvet, E. (2015). Does dollar-cost averaging make financial sense? *Journal of Financial Planning*, *28*(10), 52–58.

Choijil, E., Méndez, C. E., Wong, W. K., Vieito, J. P., & Batmunkh, M. U. (2022). Thirty years of herd behavior in financial markets: A bibliometric analysis. *Research in International Business and Finance*, *59*, 101506.

Damodaran, A. (2012). *Investment valuation: Tools and techniques for determining the value of any asset* (3rd ed.). Wiley Finance.

Dickey, D. A., & Fuller, W. A. (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, *74*(366), 427–431. https://doi.org/10.2307/2286348

Digrin. (2026). TSMC (2330.TW) historical price data. https://www.digrin.com/stocks/detail/2330.TW/price

Egozcue, M., García, L.F., Wong, W.K., & Zitikis, R. (2011). Do investors like to diversify? A study of Markowitz preferences. *European Journal of Operational Research*, *215*(1), 188-193.

Egozcue, M., & Wong, W.K. (2010). Gains from diversification on convex combinations: A majorization and stochastic dominance approach. *European Journal of Operational Research*, *200*(3), 893-900.

Emmanoulopoulos, D., Olby, O., Lyon, J., & Stillman, N. R. (2025). To trade or not to trade: An agentic approach to estimating market risk improves trading decisions. *arXiv*. https://arxiv.org/abs/2507.08584

Fabozzi, F. J., Fung, C. Y., Lam, K., & Wong, W. K. (2013). Market overreaction and underreaction: Tests of the directional and magnitude effects. *Applied Financial Economics*, *23*(18), 1469-1482.

Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical evidence. *Journal of Finance*, *25*(2), 383–417. https://doi.org/10.2307/2325486

Fong, W. M., Lean, H. H., & Wong, W.-K. (2008). Stochastic dominance and behavior towards risk: The market for internet stocks. *Journal of Economic Behavior & Organization*, *68*(1), 194–208. https://doi.org/10.1016/j.jebo.2008.03.006

Gasbarro, D., Wong, W.-K., & Zumwalt, J. K. (2007). Stochastic dominance analysis of iShares. *European Journal of Finance*, *13*(1), 89–101. https://doi.org/10.1080/13518470600658070

Granger, C. W. J., & Newbold, P. (1974). Spurious regressions in econometrics. *Journal of Econometrics*, *2*(2), 111–120. https://doi.org/10.1016/0304-4076(74)90034-7

Grossman, S. J., & Stiglitz, J. E. (1980). On the impossibility of informationally efficient markets. *American Economic Review*, *70*(3), 393–408. https://doi.org/10.2307/1913410

Guo, X., Chan, R. H., Wong, W. K., & Zhu, L. (2019). Mean–variance, mean–VaR, and mean–CVaR models for portfolio selection with background risk. *Risk Management*, *21*(2), 73-98.

Guo, X., Jiang, X., & Wong, W. K. (2017). Stochastic dominance and omega ratio: Measures to examine market efficiency, arbitrage opportunity, and anomaly. *Economies*, *5*(4), 38.

Guo, X., McAleer, M., Wong, W. K., & Zhu, L. (2017). A Bayesian approach to excess volatility, short-term underreaction and long-term overreaction during financial crises. *The North American Journal of Economics and Finance*, *42*, 346-358.

Guo, X., Niu, C., & Wong, W. K. (2019). Farinelli and Tibiletti ratio and stochastic dominance. Risk Management, 21(3), 201-213.

Guo, X., Wagener, A., Wong, W. K., & Zhu, L. (2018). The two-moment decision model with additive risks. *Risk Management*, *20*(1), 77-94.

Guo, X., & Wong, W. K. (2016). Multivariate stochastic dominance for risk averters and risk seekers. *RAIRO-Operations Research*, *50*(3), 575-586.

Hui, Y., Shi, M., Wong, W. K., & Zheng, S. (2024). Pragmatic attitude to large-scale Markowitz's portfolio optimization and factor-augmented derating. *International Review of Financial Analysis*, *96*, 103628.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291. https://doi.org/10.2307/1914185

Kelly, J. L. (1956). A new interpretation of information rate. *Bell System Technical Journal*, *35*(4), 917–926. https://doi.org/10.1002/j.1538-7305.1956.tb03809.x

Kung, J. J., & Wong, W. K. (2009b). Efficiency of the Taiwan stock market. *The Japanese Economic Review*, *60*(3), 389-394.

Kung, J. J., & Wong, W. K. (2009a). Profitability of technical analysis in the Singapore stock market: Before and after the Asian financial crisis. Journal of Economic Integration, 24(1), 133-150.

Kwiatkowski, D., Phillips, P. C. B., Schmidt, P., & Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, *54*(1–3), 159–178. https://doi.org/10.1016/0304-4076(92)90104-Y

Lam, K., Lean, H. H., & Wong, W.-K. (2016). *Stochastic dominance and investors' behavior towards risk: The Hong Kong stocks and futures markets* (MPRA Paper No. 74386). University Library of Munich, Germany. https://mpra.ub.uni-muenchen.de/74386/

Lam, K., Liu, T., & Wong, W. K. (2010). A pseudo-Bayesian model in financial decision making with implications to market volatility, under-and overreaction. *European Journal of Operational Research*, *203*(1), 166-175.

Lam, K., Liu, T., & Wong, W. K. (2012). A new pseudo-Bayesian model with implications for financial anomalies and investors' behavior. *Journal of Behavioral Finance*, *13*(2), 93-107.

Lam, K., Wong, M. C. M., & Wong, W. K. (2006). New variance ratio tests to identify random walk from the general mean reversion model. Advances in Decision Sciences, *2006*, 12314-1.

Lam, W. S. V., Chong, T. T. L., & Wong, W. K. (2007). Profitability of intraday and interday momentum strategies. *Applied Economics Letters*, *14*(15), 1103-1108.

Leung, P. L., Ng, H. Y., & Wong, W. K. (2012). An improved estimation to make Markowitz's portfolio optimization theory users friendly and estimation accurate with application on the US stock market investment. *European Journal of Operational Research*, *222*(1), 85-95.

Leung, P. L., & Wong, W. K. (2008). On testing the equality of the multiple Sharpe ratios, with application on the evaluation of Ishares. Journal of Risk, 10(3), 1-16.

Li, C. K., & Wong, W. K. (1999). Extension of stochastic dominance theory to random variables. *RAIRO-Operations Research-Recherche Opérationnelle*, *33*(4), 509-524.

Li, H., Bai, Z., Wong, W. K., & McAleer, M. (2022). Spectrally-corrected estimation for high-dimensional Markowitz mean-variance optimization. *Econometrics and Statistics*, *24*, 133-150.

Li, Z., Hui, Y., Wong, W. K., & Lin, R. (2025). Portfolio Selection Based on Mean-Generalized Variance Analysis: Evidence from the G20 Stock Markets. *Asia-Pacific Journal of Operational Research*, *42*(03), 2450016.

Linton, O., Maasoumi, E., & Whang, Y.-J. (2005). Consistent testing for stochastic dominance under general sampling schemes. *Review of Economic Studies*, *72*(3), 735–765. https://doi.org/10.1111/j.1467-937X.2005.00350.x

Lopez-Lira, A., & Tang, Y. (2023). Can ChatGPT forecast stock price movements? Return predictability and large language models. *SSRN Working Paper 4412788*. https://doi.org/10.2139/ssrn.4412788

Lu, R., Hoang, V. T., & Wong, W. K. (2021). Do lump-sum investing strategies really outperform dollar-cost averaging strategies?. *Studies in Economics and Finance*, *38*(3), 675-691.

Lu, R., Wang, J. J., & Wong, W. K. (2022). Investment based on size, value, momentum and income measures: a study in the Taiwan stock market. *Annals of Financial Economics*, *17*(04), 2250027.

Lu, R., Yang, C. C., & Wong, W. K. (2018). Time diversification: Perspectives from the economic index of riskiness. *Annals of Financial Economics*, *13*(03), 1850011.

Lv, Z., Chu, A. M., Wong, W. K., & Chiang, T. C. (2021). The maximum-return-and-minimum-volatility effect: evidence from choosing risky and riskless assets to form a portfolio. *Risk Management*, *23*(1), 97-122.

Lv, Z., Tsang, C. K., Wagner, N. F., & Wong, W. K. (2023). What is an optimal allocation in Hong Kong stock, real estate, and money markets: an individual asset, efficient frontier portfolios, or a naïve portfolio? Is this a new financial anomaly?. *Emerging Markets Finance and Trade*, *59*(5), 1554-1571.

Ma, C., & Wong, W. K. (2010). Stochastic dominance and risk measure: A decision-theoretic foundation for VaR and C-VaR. *European Journal of Operational Research*, *207*(2), 927-935.

MacLean, L. C., Thorp, E. O., & Ziemba, W. T. (2010). The Kelly capital growth investment criterion: Theory and practice. *World Scientific*.

MarketScreener. (2026). TSMC (2330) historical quotes. https://www.marketscreener.com/quote/stock/TSMC-TAIWAN-SEMICONDUCTOR-6492349/

Munkh-Ulzii, B. J., McAleer, M., Moslehpour, M., & Wong, W. K. (2018). *Confucius and herding behaviour in the stock markets in China and Taiwan. Sustainability, 10 (12), 4413*.

Ng, P., Wong, W. K., & Xiao, Z. (2017). Stochastic dominance via quantile regression with applications to investigate arbitrage opportunity and market efficiency. *European Journal of Operational Research*, *261*(2), 666-678.

Niu, C., Guo, X., McAleer, M., & Wong, W. K. (2018). Theory and application of an economic performance measure of risk. *International Review of Economics & Finance*, *56*, 383-396.

Niu, C., Wong, W. K., & Xu, Q. (2017). Kappa ratios and (higher-order) stochastic dominance. *Risk Management*, *19*(3), 245-253.

Ortobelli Lozza, S., Wong, W. K., Fabozzi, F. J., & Egozcue, M. (2018). Diversification versus optimality: is there really a diversification puzzle?. *Applied Economics*, *50*(43), 4671-4693.

Reuters. (2026, March 4). Iran conflict escalation sends oil above USD 100 per barrel. *Reuters News Service*. https://www.reuters.com

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, *27*(3), 379–423. https://doi.org/10.1002/j.1538-7305.1948.tb01338.x

Sharpe, W. F. (1966). Mutual fund performance. *Journal of Business*, *39*(1), 119–138. https://doi.org/10.1086/294846

TSMC Investor Relations. (2025). *TSMC 2025 fourth quarter results and full year summary*. https://ir.tsmc.com

Vieito, J. P., Espinosa, C., Wong, W. K., Batmunkh, M. U., Choijil, E., & Hussien, M. (2024). Herding behavior in integrated financial markets: the case of MILA. *International Journal of Emerging Markets*, *19*(11), 3801-3827.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., Chi, E., Le, Q., & Zhou, D. (2022). Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, *35*, 24824–24837. https://arxiv.org/abs/2201.11903

Wong, W. K. (2006). Stochastic dominance theory for location-scale family. Advances in Decision Sciences, 82049-1.

Wong, W.-K. (2007). Stochastic dominance and mean-variance measures of profit and loss for business planning and investment. *European Journal of Operational Research*, *182*(2), 829–843. https://doi.org/10.1016/j.ejor.2006.09.032

Wong, W. K., & Chan, R. (2008). Markowitz and prospect stochastic dominances. *Annals of Finance*, *4*(1), 105-129.

Wong, W.-K., Cheng, Y., & Yue, M. (2024). Could regression of stationary series be spurious? *Asia-Pacific Journal of Operational Research*. https://doi.org/10.1142/S0217595924400098

Wong, W. K., Chew, B. K., & Sikorski, D. (2001). Can the forecasts generated from E/P ratio and bond yield be used to beat stock markets?. *Multinational Finance Journal*, *5*(1), 59-86.

Wong, W. K., Du, J., & Chong, T. T. L. (2005). Do the technical indicators reward chartists in Greater China stock exchanges. *Review of Applied Economics*, *1*(2), 183-205.

Wong, W. K., & Li, C. K. (1999). A note on convex stochastic dominance. *Economics Letters*, *62*(3), 293-300.

Wong, W. K., & Ma, C. (2008). Preferences over location-scale family. *Economic Theory*, *37*(1), 119-146.

Wong, W. K., Manzur, M., & Chew, B. K. (2003). How rewarding is technical analysis? Evidence from Singapore stock market. *Applied Financial Economics*, *13*(7), 543-551. https://doi.org/10.1080/09603100210161

Wong, W. K., & McAleer, M. (2009). Mapping the Presidential Election Cycle in US stock markets. *Mathematics and Computers in Simulation*, *79*(11), 3267-3277.

Wong, W.-K., & Pham, M. T. (2022a). Could the test from the standard regression model could make significant regression with autoregressive noise become insignificant?. The International Journal of Finance, 34, 1–18.

Wong, W.-K., & Pham, M. T. (2022b). Could the test from the standard regression model could make significant regression with autoregressive noise become insignificant – a note. The International Journal of Finance, 34, 19-39.

Wong, W.-K., & Pham, M. T. (2023a). Could the test from the standard regression model could make significant regression with autoregressive Yt and Xt become insignificant?. The International Journal of Finance, 35, 1–19.

Wong, W.-K., & Pham, M. T. (2023b). Could the test from the standard regression model could make significant regression with autoregressive Yt and Xt become insignificant – a note. The International Journal of Finance, 35, 20-41.

Wong, W. K., & Pham, M. T. (2025a). Could the correlation of a stationary series with a non-stationary series obtain meaningful outcomes?. *Annals of Financial Economics*, *20*(03), 2550015.

Wong, W.-K., & Pham, M. T. (2025b). How to model a simple stationary series with a non-stationary series?. The International Journal of Finance, 37, 1–19.

Wong, W.-K., & Pham, M. T. (2026a). Could the panel regression be used to examine the relationship between I(0) and I(1) series?. Advances in Decision Sciences, 30(2), forthcoming.

Wong, W.-K., & Pham, M. T. (2026b). Could we use correlation to examine panel data with I(0) and I(1) variables? The International Journal of Finance, 38, forthcoming.

Wong, W.-K., Pham, M. T., & Yue, M. (2024). Could regressing a stationary series on a non-stationary series obtain meaningful outcomes – a remedy. The International Journal of Finance, 36, 1–20.

Wong, W. K., Wright, J. A., Yam, S. C. P., & Yung, S. P. (2012). A mixed Sharpe ratio. *Risk and Decision Analysis*, *3*(1-2), 37-65.

Wong, W. K., Yeung, D., & Lu, R. (2023). The mean-variance rule for investors with reverse S-shaped utility. *Annals of Financial Economics*, *18*(01), 2250030.

Wong, W. K., & Yue, M. (2024). Could regressing a stationary series on a non-stationary series obtain meaningful outcomes?. *Annals of Financial Economics*, *19*(03), 2450011.

Wong, W. K., Zhu, Z., Jiang, I. M., & Bouri, E. (2026). Arbitrage opportunities in no-arbitrage portfolios: The case of Bitcoin and Treasury Bills. *Investment Analysts Journal*, *55*(1), 55-86.

Yao, S., Zhao, J., Yu, D., Du, N., Shafran, I., Narasimhan, K., & Cao, Y. (2022). ReAct: Synergizing reasoning and acting in language models. *arXiv*. https://arxiv.org/abs/2210.03629