

AI-Driven Retail Trading and the Institutional Complementarity Trap: Why China's Market Reform Requires Simultaneous, Not Delayed, Change

Minxing Wang

wangminxing581@gmail.com

HSE University

Pavel Braslavski

Ural Federal University

Dmitry I. Ignatov

HSE University

Article

Keywords: evolutionary game theory, replicator dynamics, AI-driven trading, Chinese stock market, institutional complementarity, multi-steady-state dynamics

Posted Date: September 7th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10448366/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

AI-Driven Retail Trading and the Institutional Complementarity Trap: Why China’s Market Reform Requires Simultaneous, Not Delayed, Change

Anonymous Author(s)

Abstract

We develop a five-population asymmetric evolutionary game to analyze why AI-driven retail trading and the rise of quantitative institutions do not, by themselves, push China’s stock market toward institutional modernization. Unlike prior models that bundle institutional constraints, we decompose the environment into three distinct institutional gates—T+1 settlement, price limits, and short-selling constraints—and examine how their interaction with AI-driven retail strategies creates a *two-timescale dynamic lock-in*: on short timescales ($\sim 5,000$ steps) the system converges to a locally attractive HTT (high-AI, T+1, price-limits) trap that appears stable to market participants, while on long timescales ($\sim 50,000$ steps) the trap proves transient and decays toward the traditional-market attractor. AI adoption does not prepare the market for modernization: it extends the trap’s transient lifetime, and beyond a critical penetration band ($\lambda^* \in [2.2, 2.4]$) the trap hardens into permanent stability, while excessive penetration ($\lambda \geq 2.9$) can collapse even a mature market. Stochastic analysis further shows that the mature-market attractor, once reached, is genuinely persistent—in 60% of low-noise trajectories it is never abandoned—whereas the traditional attractor is not. The policy implications are threefold. *Complementarity*: single-gate reforms fail except within narrow parameter windows, the trading-rule gate always fails alone, and two-gate reform succeeds in only 9.4% of the policy space. *Timeliness*: simultaneous three-gate reform is fastest—about 14 times faster than the slowest viable sequential scheme—and delay beyond a critical interval fails irreversibly. *Evaluation*: reform success cannot be judged at short horizons, since shocks that fail at 5,000 steps succeed in 81.5% of cases at 50,000 steps. The two-attractor architecture is structurally robust: it survives $\pm 50\%$ perturbation of 91.5% of the payoff parameters and is carried by the base payoffs and the joint sign pattern of cross-population incentives.

Keywords: evolutionary game theory, replicator dynamics, AI-driven trading, Chinese stock market, institutional complementarity, multi-steady-state dynamics

1 Introduction

1.1 The Problem: AI-Driven Retail Trading and the Limits of Gradual Institutional Change

A striking feature of China’s stock market is the rapid proliferation of AI-driven retail trading tools—from intelligent stock-picking algorithms to strategy-backtesting software—among a population of over 200 million retail investors (China Securities Depository and Clearing Corporation statistics). Yet despite this technological transformation, the market’s core institutional architecture remains largely unchanged: T+1 settlement (introduced 1995), $\pm 10\%$ price limits ($\pm 20\%$ for the STAR Market and, since August 2020, for ChiNext; $\pm 30\%$ for the Beijing Stock Exchange), and functionally constrained short-selling. This creates a puzzle: why does AI adoption not push the market toward institutional modernization? Instead, the system appears locked in a hybrid state where retail investors use sophisticated tools within a rigid institutional framework.

Prior evolutionary game models of the stock market have overwhelmingly bundled these constraints into a single dimension—typically a “policy state” or “institutional regime” variable. We argue this bundling is methodologically convenient but empirically misleading. T+1, price limits, and short-selling constraints are distinct institutional levers with distinct economic mechanisms, and their interaction with AI-driven retail trading creates what we term an *institutional complementarity trap*: the technology is adopted, but the institutional environment that would make it welfare-enhancing remains unaddressed.

- T+1 restricts intraday liquidity, amplifying momentum and preventing immediate price correction. In the presence of AI-driven trading tools, this creates a “liquidity trap”: retail investors use algorithmic signals to enter positions but cannot exit when signals reverse, amplifying losses and reinforcing strategy homogenization.
- Price limits create “magnet effects” and trading halts, distorting price discovery. Combined with T+1, they form a “double lock” that prevents both intraday exit and full price adjustment.
- Short-selling constraints eliminate a natural stabilizing force: the ability of informed traders to express negative views through the market. In the presence of AI-driven herding and T+1, this creates a “one-sided market” where synchronized momentum cannot be arbitrated away.

These three mechanisms are not substitutes; they are complements that jointly determine whether AI-driven retail trading leads to market improvement or institutional entrapment. A model that bundles them cannot distinguish whether a reform is removing one constraint while leaving others intact, or whether the reform is only effective when all three constraints are lifted simultaneously.

Related literature. Our work sits at the intersection of three literatures. First, the evolutionary game theory (EGT) literature on multi-population asymmetric games [1–3] has established that interior mixed equilibria are structurally unstable when three or more populations interact. EGT has been increasingly applied to Chinese institutional settings, including carbon market development [4], e-commerce

quality regulation [5], government regulatory policy for digital transformation [6], carbon-market trading equilibria [7], consumer-finance regulation [8], and greenwashing governance [9], but the complete basin structure of decomposed institutional games remains unexplored. Second, the institutional complementarity and path-dependence literature [10–15] has developed static frameworks for understanding why institutions reinforce each other and why historical contingencies can lock an economy into a sub-optimal configuration. Our contribution is to extend this literature by introducing *dynamic lock-in*: the trap is not a permanent equilibrium but a transient trapping state with a finite lifetime, and the effective reform window is determined by the time the system spends in the transient trapping state before decaying to the traditional-market attractor. Third, the literature on algorithmic and AI-driven trading [16–19] examines how algorithmic tools affect market behavior, but rarely connects this to institutional reform dynamics. Our contribution is to bridge these three literatures by showing that AI adoption within an unreformed institutional framework deepens—rather than alleviates—the complementarity trap.

1.2 The Three-Gates Hypothesis: From AI Adoption to Institutional Complementarity

We hypothesize that AI-driven retail trading has created a new form of institutional lock-in: the technology is adopted, but the institutional environment that would make it welfare-enhancing remains unaddressed. The stock market faces three institutional gates that must be crossed simultaneously to escape this trap and reach the mature-market stable state:

Gate 1: T+1 Settlement. The T+1 rule prevents retail investors from closing positions intraday, forcing them to hold overnight exposure. In the presence of AI-driven trading tools (AL), this creates a “liquidity trap”: retail investors use algorithmic signals to enter positions but cannot exit when signals reverse, amplifying losses and reinforcing homogenization.

Gate 2: Price Limits. The $\pm 10\%$ price limit rule truncates the price distribution, preventing full price discovery. In the presence of T+1, price limits create a “double lock”: investors cannot exit (T+1) and the price cannot fully adjust (price limits), leading to predictable next-day gaps that quant institutions exploit.

Gate 3: Short-Selling Constraints. The limited short-selling mechanism prevents informed traders from arbitraging overvaluation. In the presence of T+1 and price limits, short-selling constraints create a “one-sided market”: only upward momentum can be expressed, while downward corrections are suppressed until they explode.

We ask: Can a standard multi-population evolutionary game—with T+1, price limits, and short-selling as distinct strategy dimensions—predict whether the stock market is locked in a trap, and what complementary reforms are needed to escape?

1.3 Model Architecture: Five Populations, 96 Pure-Strategy Corners

The model consists of five populations with distinct strategy sets:

The P3 strategy space decomposes the institutional environment into four regime configurations that capture the interaction between technology adoption (high vs. low AI) and two institutional policy levers (T+1 settlement and price limits):

- **LTT**: Low AI + T+1 + Price Limits (traditional stock market, pre-2010s)
- **LNT**: Low AI + T+1 + No Price Limits (hypothetical intermediate)
- **HTT**: High AI + T+1 + Price Limits (current stock market, 2020s)
- **HTN**: High AI + No T+1 + No Price Limits (mature market, US-style)

Two features of this decomposition deserve explicit comment. **First, the P3 labels refer to the regime, not to any individual’s strategy choice.** The “H/L” prefix indicates the level of *regime-level AI infrastructure supply*—the data feeds, algorithmic order-routing capacity, and AI-based trading services that the institutional environment makes available to all participants—and is decoupled from P1’s individual adoption margin (AL/AG/E). A market can therefore sit in P3 = HTT while a majority of retail investors remain passive (P1 = E): the infrastructure is in place even though households have not (yet) adopted AI tools. The current-market trap identified in Section 3.3 is exactly such a configuration, so there is no contradiction between a “High-AI” regime label and a passive retail population.

Second, the bundling of the three gates is asymmetric by design, and the asymmetry mirrors the regulatory structure. T+1 settlement and price limits are bundled into P3 because they are joint attributes of a single trading-rule regime: both are set by the same exchange-level rulebook and their interaction (the “double lock”) is precisely the complementarity we wish to represent as a within-regime property. Short-selling, by contrast, is governed by a distinct policy lever—the eligible-securities list and the refinancing mechanism—with its own history of expansion and tightening (Section 2.3.1), and it can be liberalized or restricted without touching the settlement or price-limit rules; it is therefore assigned its own population (P5) so that “reforming the short-selling gate” is a well-defined experiment separate from changing the trading-rule regime. The counterfactual regime LNT (Low AI + T+1 + no price limits) is retained for the same reason of non-prejudgment: no major market has combined T+1 with unrestricted prices, but excluding LNT would force the model to jump directly between LTT and HTN and would assume, rather than demonstrate, that such hybrid configurations are dynamically non-viable. As the stability analysis shows (Section 3.2), no LNT configuration is among the stable attractors—the dynamics, not the modeler, discard the hybrid.

The total pure-strategy space is $3 \times 2 \times 4 \times 2 \times 2 = 96$ corners. Each corner represents a candidate steady state. The question is not whether the system has a single equilibrium but which corners are stable attractors and which are unstable traps.

1.4 Terminology: “Structural” vs. “Path-Dependent”

Before stating the contributions, we clarify two terminological distinctions that structure the paper’s claims.

Structural (attractor existence). A conclusion is *structural* if it concerns the existence, number, or identity of stable attractors in the replicator dynamics, and if this conclusion survives large parameter perturbations (e.g., $\pm 50\%$). Structural conclusions

describe the qualitative architecture of the basin landscape, not quantitative parameter values.

Path-dependent (attractor accessibility). A conclusion is *path-dependent* if it concerns which attractor is reached from a specific initial condition, or the width of a particular basin. Path-dependent conclusions are sensitive to parameter values and initial conditions, even though the underlying attractor structure is structural.

Calibration-dependent vs. calibration-independent. A conclusion is calibration-dependent if it requires precise parameter values. A conclusion is calibration-independent if it holds across a wide range of parameter values.

Transient trapping, saddle, and lock-in. Three dynamical terms recur throughout the paper, and we use them in the following strict senses. A state is a *transient trapping state* if trajectories enter its neighborhood and reside there for a long but finite time before escaping: it is not locally asymptotically stable (its Jacobian has a positive dominant eigenvalue), but the escape is slow enough that the state appears stable to observers at the short-run horizon. The HTT trap under the baseline parameterization is the central example. A *saddle* is a fixed point whose Jacobian possesses both stable and unstable directions, so that trajectories may approach along the stable manifold yet eventually depart along the unstable one; the unstable mixed states of Section 3 are saddles in this sense. *Lock-in* is used in the institutional, not the purely dynamical, sense: a self-reinforcing configuration whose complementarities make departure costly for policymakers, regardless of whether the underlying dynamical state is a stable attractor or a transient trap. Our “dynamic lock-in” thus refers to trapping by the transient HTT state, while the two stable attractors represent two alternative institutional lock-ins of markedly different stochastic persistence (Section 3.7).

Throughout the paper, we use “structural” as a synonym for “calibration-independent” (attractor existence). The apparent tension between “structural conclusions” and “9.4% policy space” is resolved by recognizing that the former refers to attractor existence, while the latter refers to attractor accessibility from a specific initial condition.

1.5 Contribution

The paper offers four contributions, positioned relative to the existing literature on institutional complementarity and multi-population game theory:

1. **Theoretical:** We state as a formal conjecture, and support with systematic numerical evidence, that a five-population asymmetric EGT with $3+2+4+2+2$ strategies admits no stable interior mixed equilibrium (Conjecture 2); instead, all 96 pure-strategy corners are fixed points; under the baseline parameterization, only 2 are stable attractors, though bifurcation analysis shows the number can increase to 3 when AI and HTT institutions become complementary. The remaining 94 are unstable saddles. This finding extends the theoretical frontier by mapping the complete basin structure of a decomposed institutional game. While [3] established that interior equilibria are structurally unstable in asymmetric games with three

or more populations, our contribution is to characterize the full basin structure—which corners are attractors, which are traps, and which basins lead where—in a game with institutional interpretation.

2. **Methodological:** We decompose three institutional constraints (T+1, price limits, short-selling) into distinct strategy dimensions and demonstrate that they form structural institutional complementarity. The two stable stacks—{NS, LTT, FO, SC} and {TI, HTN, FI, SF}—are self-reinforcing configurations that cannot be mixed without creating unstable traps, operationalizing the static complementarity concept of [10, 13] in a dynamic, multi-actor setting. A systematic parameter importance scan further decomposes the model’s 146 parameters (141 payoff parameters plus 5 learning rates) into load-bearing (structural), modulating, and neutral categories: at $\lambda = 1.0$, 13 parameters are load-bearing; at $\lambda = 2.0$, the list expands to 15, confirming that the institutional structure becomes more rigid as AI penetration increases. Two learning rates (η_1, η_2) shift from load-bearing to modulating as λ rises: learning speed modulates the transient lifetime of the trap without changing its attractor status.
3. **Empirical (simulation):** Through seven counterfactual simulations (50,000 steps), 2D policy space scanning, parameter sensitivity analysis, and transition barrier analysis, we identify that most single-point reforms converge to the traditional stable attractor (Stable A) or remain trapped in unstable mixed states. The mature market (Stable B) is accessible only under complementary multi-parameter reform that simultaneously addresses all three institutional gates, or within narrow Goldilocks calibration bands (e.g., $0.06 < \text{base}_5^{SF} < 0.08$; Table 7); a 2D scan confirms accessibility from only 9.4% of the two-parameter policy space. The transition barrier reveals that reform success cannot be evaluated at the short-run horizon—shocks that fail at 5,000 steps succeed in 81.5% of cases at 50,000 steps—and a sequencing experiment shows that simultaneous three-gate reform reaches the mature market fastest, that sequential reform succeeds only within a bounded inter-gate interval, and that delay beyond a critical interval fails irreversibly (Section 3.10). We further identify an AI penetration paradox: AI adoption within an unreformed institutional framework (HTT) deepens the complementarity trap for most initial conditions, rather than preparing the market for modernization, and the paradox is driven by cross-population reinforcing interactions, not by base payoffs alone.
4. **Normative:** We shift the policy framework from “which instrument to adjust?” to “which complementary stack to commit to?” and further to “when to act and at what intensity?” The two stable attractors represent two irreconcilable institutional architectures, and the policy challenge is to commit to one stack and engineer the complementary transition across all three gates simultaneously. We identify a “reform window”: the system is most effectively reformed from mixed initial conditions with all three gates initiated within a bounded interval—simultaneous initiation is fastest, and delay beyond a critical inter-gate interval forfeits the mature market altogether (Section 3.10)—and the reform intensity must be precisely calibrated to navigate the narrow Goldilocks basin of the mature market.

2 Model

2.1 Five-Population Asymmetric EGT

The model consists of five populations, each with distinct strategy sets as shown in Table 1. The total strategy space is 96 pure-strategy profiles. Each profile (a corner of the full strategy space) is a candidate steady state. The expected payoff for each population depends on the strategy distributions of all other populations.

Table 1 Strategy sets for the five populations.

Population	Strategies	Description
P1 (Retail)	E, AL, AG	Passive holding (E), AI-driven learning (AL), AI-driven gambling (AG)
P2 (Institutions)	NS, TI	State-controlled (NS), Quant-friendly (TI)
P3 (Institutional Regime)	LTT, LNT, HTT, HTN	Low/High AI $\times$ T+1 $\times$ Price Limits
P4 (Foreign Capital)	FI, FO	Inflow (FI), Outflow (FO)
P5 (Short-Selling)	SC, SF	Constrained (SC), Free (SF)

Methodological clarification on P3. Unlike P1 (retail investors) and P2 (institutions), which represent populations of strategic agents who choose their type, P3 is a characterization of the institutional regime. T+1 and price limits are exogenous policy rules set by regulators; they do not “evolve” through individual choice. Modeling the regime as a strategy dimension in the replicator dynamics is a standard technique in multi-population EGT: it treats the institutional environment as a meta-game in which the “frequency” of each regime represents the effective weight of that institutional configuration in the cross-population payoff structure [13, 20]. When P3 = LTT with frequency 0.7, it means that 70% of the institutional interactions in the system occur under the traditional rules. This is analogous to modeling the “state of technology” as a strategy in evolutionary innovation games (e.g., [21]). The benefit of this representation is that it allows us to embed AI penetration as a continuous variable that interacts with discrete policy rules within a unified payoff matrix, without requiring a separate law-of-motion equation for institutional change.

A fully endogenous institutional model would require a separate regulatory-reaction function (e.g., as in [22]), but calibrating such a function would require estimating policy-response elasticities that are not observable. The P3-as-strategy representation is therefore a tractable approximation that keeps the focus on cross-population interactions while remaining silent on the precise mechanism of policy reform. We emphasize three interpretive principles: (1) the “frequency” of P3 = HTT is not a literal population share but the effective weight of that regime in the cross-population payoff structure; (2) we do not model the regulator as a strategic agent—the dynamics reveal which institutional configurations are self-reinforcing, not how regulators should change rules; (3) the representation is silent on the speed of institutional change, focusing instead on structural attractors and basins.

Economic interpretation of P3’s “payoff”: institutional persistence fitness. Because institutions are not strategic agents with utility functions, the quantity π_3 in Eq. (A4) (Section A) should not be read as an institutional “utility.” Instead, we interpret π_3 as a *persistence fitness*: a reduced-form measure of how self-reinforcing each institutional configuration is, given the current strategy composition of the other four populations. This interpretation rests on three theoretical foundations. First, [23] argues that institutions exhibit increasing returns: once a set of rules is in place, the political and organizational adaptation to those rules makes them progressively harder to reverse. The base term $\text{base}_3^{s_3}$ captures this intrinsic stickiness—the political cost of dismantling an established regime. Second, [13] formalizes institutional complementarity: a configuration is more persistent when it is compatible with the strategies adopted by other economic actors. The cross-terms $\text{cross}_{3j}^{s_3, s_j}$ encode exactly this compatibility: a positive cross-effect means that the strategy mix in population j reinforces the persistence of regime s_3 , while a negative cross-effect means that the strategy mix erodes it. Third, [24] model institutional change as constrained by the political equilibrium: “viable” institutions are those that serve the interests of powerful coalitions, and their viability depends endogenously on the economic structure. In our framework, the replicator dynamics for P3 select toward configurations with higher persistence fitness—i.e., configurations that, once established, are politically and economically harder to dislodge. The learning rate $\eta_3 = 0.5$ (slower than $\eta_1 = \eta_2 = 1.0$) reflects the well-documented regulatory and political inertia that slows institutional adjustment relative to market-participant learning [25].

2.2 Payoff Structure: Institutional Decomposition

The payoff structure embeds three institutional intuitions through cross-population interactions:

Intuition 1: AI hallucination erodes retail algorithmic learning. When P3 = HTT (high AI + T+1 + price limits), retail investors who adopt algorithmic learning (P1 = AL) suffer from overconfidence and pattern hallucination, compounded by the inability to exit positions intraday. The cross-effect is negative: $\text{cross}_{13}^{\text{AL}, \text{HTT}} = -0.035$. This is much more negative than $\text{cross}_{13}^{\text{AL}, \text{LNT}} = +0.015$ (T+1 only, no price limits) or $\text{cross}_{13}^{\text{AL}, \text{HTN}} = +0.030$ (neither), capturing the “double lock” of T+1 and price limits.

Intuition 2: T+1 and price limits create a double lock. When P3 = HTT, the combination of T+1 and price limits prevents both intraday exit and full price adjustment. This is captured by the P3→P1 cross-effect: $\text{cross}_{13}^{\text{AL}, \text{HTT}}$ is much more negative than $\text{cross}_{13}^{\text{AL}, \text{LNT}}$ (T+1 only, no price limits) or $\text{cross}_{13}^{\text{AL}, \text{HTN}}$ (neither).

Intuition 3: Short-selling constraints create a one-sided market. When P5 = SC (short-selling constrained), the market lacks a natural stabilizing force. This is captured by P5→P3 cross-effects: $\text{cross}_{35}^{\text{HTT}, \text{SC}} = +0.010$ (SC reinforces HTT) and $\text{cross}_{35}^{\text{HTN}, \text{SF}} = +0.015$ (SF reinforces HTN). The short-selling mechanism and the institutional regime are complements, not substitutes.

Intuition 4: Institutional complementarity. The P2–P3–P4–P5 interactions form two complementary stacks:

- **Stack A (Traditional):** {NS, LTT, FO, SC} — NS benefits from LTT (traditional market, easy to stabilize), FO is the natural response to SC (foreign capital avoids markets without hedging tools), and SC benefits from LTT (traditional markets don't need short-selling).
- **Stack B (Mature):** {TI, HTN, FI, SF} — TI benefits from HTN (quant institutions thrive in mature markets), FI is attracted to SF (foreign capital requires hedging tools), and SF benefits from HTN (mature markets need price discovery).

2.2.1 Microfoundations: From Institutional Mechanisms to Cross-Effect Signs

The sign and relative magnitude of each cross-effect are derived from the underlying institutional economics, not fitted to data. We provide the economic logic for the four most critical cross-effects:

- $\text{cross}_{13}^{\text{AL,HTT}} < 0$: In the HTT regime, retail investors who adopt AI-driven learning (AL) face a “double lock”: they cannot exit intraday (T+1) and the price cannot fully adjust to new information (price limits). The AI algorithm generates false positive signals (hallucinations) that the investor cannot immediately correct, eroding the payoff to AL relative to passive holding (E). The magnitude ($|-0.035| = 0.035$) is calibrated to be larger than $\text{cross}_{13}^{\text{AL,LNT}}$ (+0.015, T+1 only) and larger than $\text{cross}_{13}^{\text{AL,HTN}}$ (+0.030, neither constraint), capturing the intuition that the double lock has a super-additive negative effect on algorithmic learning.
- $\text{cross}_{35}^{\text{HTT,SC}} > 0$: When short-selling is constrained (SC), the market lacks a natural stabilizing force. In the HTT regime, this absence is particularly damaging because the combination of AI-driven herding and T+1 creates synchronized momentum that cannot be arbitrated away. The positive cross-effect means that SC reinforces the HTT regime: without short-selling, the HTT regime is more “stable” (in the sense of persistence, not welfare) because negative information is suppressed. The magnitude (+0.010) is smaller than $\text{cross}_{35}^{\text{HTN,SF}}$ (+0.015) because in a mature market (HTN), the absence of short-selling is more costly than in a traditional market (HTT), where other constraints already suppress price discovery.
- $\text{cross}_{42}^{\text{FI,TI}} > 0$ and $\text{cross}_{43}^{\text{FI,HTN}} > 0$: Foreign capital inflow (FI) is attracted to quant-friendly institutions (TI) and mature regulations (HTN) because these environments provide the liquidity, information transparency, and hedging tools that foreign investors require. The cross-effects are positive because FI is a complement to market modernization: foreign capital does not flow to traditional markets (NS, LTT) because the absence of hedging tools and the presence of state intervention create political risk.
- $\text{cross}_{42}^{\text{FO,TI}} < 0$ and $\text{cross}_{43}^{\text{FO,HTN}} < 0$: Foreign capital outflow (FO) is penalized in the TI and HTN environments. When quant institutions dominate and regulations are mature, the opportunity cost of remaining outside the market is high (FO loses access to alpha). The negative cross-effects ensure that FO is dominated by FI in the mature-market stack.

2.2.2 On Calibration

The quantitative values of cross-effects are not estimated from micro data. They are structural benchmarks calibrated to satisfy three criteria: (1) the sign matches the institutional intuition; (2) the relative magnitudes preserve the ranking of intuitions; and (3) the values are within the economically plausible range $[-0.10, +0.10]$ so that no single cross-effect dominates the base payoff. The sensitivity analysis (Section 3.11) shows that the two-attractor structure is preserved under $\pm 50\%$ perturbation for the large majority of payoff parameters (129 of 141; 91.5%), and that what matters is the joint pattern of signs and relative magnitudes, not the precise values.

Structural Robustness. The cross-population payoff matrix contains 141 parameters. A systematic parameter-compression analysis (Section 3.11) demonstrates that a 61-parameter minimal sufficient model—obtained by zeroing all non-structural, non-P5 cross-effects—retains the same two-attractor structure. The same analysis shows that this robustness is concentrated: it is the base payoffs and the joint sign pattern of the cross-effects that carry the structure, while most individual cross-effect magnitudes are dynamically inert (Section 3.11).

2.3 Payoff Structure Summary

Each population i 's payoff for playing pure strategy s_i follows an additive decomposition: a base payoff plus pairwise cross-population interaction terms,

$$\pi_i(s_1, \dots, s_5) = \text{base}_i^{s_i} + \sum_{j \neq i} \text{cross}_{ij}^{s_i, s_j}, \quad (1)$$

with an additional regime-conditional homogeneity penalty for P1 when AI-driven retail strategies meet high-AI institutional regimes. The complete mathematical specification of the payoff functions and replicator dynamics, together with the full parameter tables (base payoff values and cross-population interaction matrices), is given in the Appendix (Section A).

2.3.1 Calibration Protocol: Mapping Institutional Features to Payoff Parameters

A natural concern about the 146 parameters (141 payoff parameters plus 5 learning rates) is whether they are arbitrary calibration choices. We address this concern by systematically mapping the model's base payoffs and cross-effects to observable institutional features of the Chinese A-share market. The calibration protocol is not a micro-level econometric estimation—that would require disaggregated data on individual investor strategies and private institutional payoff structures, which are not publicly available. Instead, we provide a *qualitative structural mapping* that anchors each parameter to an economically meaningful institution or market process, ensuring that the model's predictions are not artifacts of arbitrary numerical choices but reflect the structural architecture of the market.

Principle 1: Base payoffs reflect default-configuration costs. The base payoff of a strategy captures its intrinsic fitness *in the absence of cross-population*

interactions. We calibrate these values according to two rules. (i) *Historical persistence*: strategies that have been the dominant configuration for decades receive higher base payoffs than newly introduced alternatives. (ii) *Regulatory burden*: strategies that face stricter regulatory constraints receive lower base payoffs than strategies that operate under lighter rules. Table 2 reports the mapping for each base payoff.

Principle 2: Cross-effects reflect institutional complementarity and externality. The cross-effect $\text{cross}_{ij}^{s_i, s_j}$ captures the payoff gain or loss to population i 's strategy s_i when population j adopts strategy s_j . Positive cross-effects indicate *complementarity* (the strategies reinforce each other); negative cross-effects indicate *substitutability* (the strategies are mutually destabilizing). We calibrate these values according to three rules. (i) *Known institutional interactions*: interactions that have been documented in the empirical literature (e.g., T+1 + price limits creating a “double lock”) are assigned large-magnitude cross-effects. (ii) *Directional consistency*: the sign of the cross-effect matches the known directional relationship (e.g., AI-driven trading within the HTT regime is known to amplify herding, hence the negative cross-effect). (iii) *Magnitude boundedness*: all cross-effects are bounded within $[-0.10, +0.10]$, ensuring that the model's dynamics are governed by the interaction structure rather than by extreme payoff differences.

Principle 3: Parameter sensitivity as a structural diagnostic. The base payoffs and cross-effects are not fitted to maximize any likelihood or match any specific time series. Instead, their values are chosen to reflect the *qualitative structure* of the market's institutional complementarity. The systematic sensitivity analysis (Section 3.11) confirms that this calibration is not fragile: the two-attractor structure survives $\pm 50\%$ perturbation for 129 of the 141 payoff parameters (91.5%), and the 12 exceptions are concentrated in the base payoffs that the calibration protocol identifies as structurally primary.

Parameter-range rationale. All payoff parameters are set within the range $[-0.10, +0.10]$. This range is chosen for three reasons. (i) *Dynamic stability*: in the replicator dynamics, payoff differences of order 0.01–0.03 generate strategy-switching rates consistent with the empirically observed adjustment speeds of retail investors (months to years) and institutions (quarters to years). (ii) *Comparability*: the symmetric range ensures that no single interaction dominates the dynamics, and the system's behavior is governed by the *architecture* of cross-population incentives rather than by extreme values. (iii) *Robustness*: the sensitivity analysis (Section 3.11) confirms that the two-attractor structure survives $\pm 50\%$ perturbation for the large majority of parameters (129 of 141), so the exact values within the $[-0.10, +0.10]$ range are not critical; what matters is the *qualitative sign pattern* and the *relative magnitude ordering*.

From parameters to policy instruments. The payoff parameters are model quantities, not policy instruments: a regulator cannot set $\text{base}_1^{\text{AL}}$ the way it sets a price-limit band. Table 3 therefore makes explicit the mapping from each parameter group to the policy lever that can move it, and—equally important—identifies which parameters lie *outside* the policy set. Three groups are directly manipulable: the trading-rule parameters of P3 (base_3 , cross_{3j} ; T+1 settlement and price-limit bands are set by

Table 2 Calibration mapping: from institutional features to base payoffs and key cross-effects.

Parameter	Value	Institutional Feature	Mapping Logic
<i>Base payoffs (default configuration)</i>			
base_1^{E}	-0.012	T+1 prevents intraday exit, forcing overnight exposure	Lower fitness than AL because passive retail cannot use AI signals to manage risk
$\text{base}_1^{\text{AL}}$	+0.022	AI-driven retail trading tools (intelligent stock-picking, backtesting)	Higher base payoff reflects information advantage of AI tools, but still constrained by T+1 and price limits
$\text{base}_3^{\text{LTT}}$	+0.080	T+1 + price limits (unchanged since 1995 and 1996)	Highest base payoff because this is the <i>default</i> institutional configuration that has persisted for decades
$\text{base}_4^{\text{FO}}$	+0.060	Foreign ownership share $\sim 3\text{--}5\%$ of A-share market capitalization; pro-cyclical north-bound flows via Stock Connect	Structural exit premium: foreign investors face T+1, price limits, and capital controls, so outflow is safer than inflow at baseline
$\text{base}_5^{\text{SC}}$	+0.050	Short-selling restricted to $\sim 2,200$ eligible stocks (SSE/SZSE expansion, Oct. 2022), tightened from 2023; securities refinancing business suspended in July 2024 (CSRC)	Default configuration: restricted short-selling has lower regulatory implementation cost than full liberalization, and the post-2023 tightening and 2024 refinancing suspension confirm the persistence of the constrained regime
$\text{base}_5^{\text{SF}}$	+0.030	Short-selling liberalization (margin-trading expansion, 2010–2024)	Lower base payoff because liberalization requires regulatory infrastructure and market-making capacity that does not yet exist at baseline
<i>Key cross-effects (institutional complementarity)</i>			
$\text{cross}_{13}^{\text{AL,HTT}}$	-0.035	AI-driven trading within T+1 + price limits amplifies herding and creates liquidity trap	Negative: AL and HTT are <i>substitutes</i> because the “double lock” (T+1 + price limits) prevents AI strategies from exiting when signals reverse, amplifying losses and homogenization
$\text{cross}_{23}^{\text{TI,HTT}}$	+0.020	Quant institutions exploit T+1 + price limits via magnet-effect prediction and next-day gap trading	Positive: TI and HTT are <i>complements</i> because the quant strategy is designed to exploit the predictable price dynamics created by T+1 and price limits
$\text{cross}_{43}^{\text{FI,HTN}}$	+0.030	Foreign capital inflows when T+1 and price limits are removed (HTN regime)	Positive: FI and HTN are <i>complements</i> because foreign investors gain full intraday liquidity and arbitrage capacity when the “double lock” is removed
$\text{cross}_{35}^{\text{HTT,SC}}$	+0.010	Short-selling constraints prevent arbitrage of synchronized momentum in the HTT regime	Positive: HTT and SC are <i>complements</i> because the absence of short-selling allows the AI-driven herding momentum to persist without correction

exchange rulebook), the foreign-access parameters of P4 (base_4 , cross_{4j} ; Stock Connect scope and admission procedures), and the short-selling parameters of P5 (base_5 , cross_{5j} ; the eligible-securities list and the refinancing mechanism). Two groups are not: the P1 base payoffs—in particular $\text{base}_1^{\text{AL}}$ —are determined by the AI technology itself and by household preferences, so policy can shift them only indirectly (investor education, product approval); and the learning rates η_k are behavioral speeds with no policy analogue. The cross-effects among strategic populations (e.g., $\text{cross}_{13}^{\text{AL,HTT}}$)

are emergent market interactions: they change only through the underlying rules, not directly. This asymmetry reinforces rather than weakens the paper’s central claim: the parameter that drives AI adoption ($\text{base}_1^{\text{AL}}$) is precisely the one policymakers cannot control, while the parameters they *can* control (the three gates) must be moved jointly because their complementarities are what hold the traps in place.

Table 3 Mapping from payoff parameter groups to policy levers. “Direct” levers can be changed by rule; “indirect” levers can be influenced but not set; parameters outside the policy set cannot be manipulated at all.

Parameter group	Policy lever	Institutional channel	Manipulability
$\text{base}_1^{\text{E/AL/AG}}$	— (AI technology; investor preferences)	Diffusion of AI trading tools; household adoption	Not manipulable
$\text{base}_2^{\text{NS/TI}}$	Market-structure and fund regulation	Licensing; program-trading oversight (2024 Provisions)	Indirect
$\text{base}_3, \text{cross}_{3j}$	Trading rules (Gate 1–2)	T+1/T+0 settlement; price-limit bands	Direct
$\text{base}_4, \text{cross}_{4j}$	Foreign access (Gate complement)	Stock Connect scope; admission procedures	Direct/Indirect
$\text{base}_5, \text{cross}_{5j}$	Short-selling rules (Gate 3)	Eligible-securities list; refinancing mechanism	Direct
$\text{cross}_{13}^{\text{AL,HTT}}$ etc.	— (emergent interaction)	Shifts only via the underlying rules above	Not manipulable
$\eta_1\text{--}\eta_5$	—	Behavioral adjustment speeds	Not manipulable

2.3.2 Numerical Integration

The replicator dynamics (A12) are integrated using the forward Euler method with step size $\Delta t = 0.03$. After each step, frequencies are projected back onto the simplex by clipping to $[10^{-8}, 1]$ and renormalizing:

$$x_i^k(t + \Delta t) = \frac{\max(x_i^k(t) + \Delta t \dot{x}_i^k(t), 10^{-8})}{\sum_{k'} \max(x_i^{k'}(t) + \Delta t \dot{x}_i^{k'}(t), 10^{-8})}. \quad (2)$$

The clipping at 10^{-8} prevents numerical extinction of strategies, ensuring that all strategies remain available for selection throughout the simulation. The step size $\Delta t = 0.03$ is chosen to balance numerical stability with computational efficiency; halving the step size produces no qualitative change in the results. As an additional safeguard, we also implemented an optional early-stopping convergence criterion: integration halts when the largest absolute per-step displacement $|dx \cdot dt|$ across all populations stays below a tolerance for a prescribed number of consecutive steps. This criterion is disabled in all reported results, so that every trajectory uses the full fixed number of steps; we verified that enabling it leaves all reported outcomes unchanged. We verify both numerical choices directly: starting from the baseline initial condition, (i) varying the clipping threshold over $\{10^{-8}, 10^{-10}, 10^{-12}\}$ leaves the trajectory unchanged step for step—the HTT-trap entry (step 3,733) and escape (step 9,354) are identical in all three cases; (ii) halving the step size to $\Delta t = 0.015$ shifts the trap-escape time by less than 0.1 time units (280.62 vs. 280.56) with the same long-run attractor; and (iii) re-integrating the *unclipped* system with adaptive Runge–Kutta solvers

(RK45 and DOP853; $\text{rtol} = 10^{-9}$, $\text{atol} = 10^{-12}$) over the equivalent continuous-time horizon reproduces the same trap entry/escape times ($t \approx 111.9 / 280.5$) and the same long-run attractor (Stable A). The transient-trapping lifetime reported in Section 3.1 is therefore not an artifact of the truncation threshold, the step size, or the forward Euler integrator.

Time-unit convention.

The replicator equations evolve in continuous time t ; the integration *step* is only the numerical discretization (1 step = $\Delta t = 0.03$ time units), while the economically meaningful quantity is the model *time unit*. Throughout this paper we adopt the heuristic convention that *one time unit corresponds to roughly one trading day*: a unit of model time is interpreted as the characteristic period over which the fastest-learning populations ($\eta = 1$) complete one meaningful strategy-revision cycle, so that daily strategy updating maps naturally onto t . Under this convention the two reference horizons are 5,000 steps = 150 time units ≈ 150 trading days $\approx$ six months (short run) and 50,000 steps = 1,500 time units $\approx 1,500$ trading days $\approx$ six calendar years (long run); the baseline HTT-trapping episode (entry at step 3,733, escape at step 9,354, a residence of 5,621 steps ≈ 169 time units) corresponds to roughly eight months. We stress that this mapping is a *heuristic order-of-magnitude calibration only*: it builds intuition about the relative duration of the model's two timescales, and it is *not* used to predict the calendar timing of any specific transition or policy event (see the time-scale limitation in Section 4). All structural conclusions are stated in model time and are invariant under any rescaling of the time axis.

2.4 Replicator Dynamics and Stability Classification

Let $x_i^k \in [0, 1]$ denote the frequency of strategy k in population i , with $\sum_k x_i^k = 1$ for each population i . The expected payoff to strategy k in population i is $\pi_i^k(\mathbf{x})$, which depends on the strategy distributions of all other populations. The average payoff in population i is $\bar{\pi}_i(\mathbf{x}) = \sum_k x_i^k \pi_i^k(\mathbf{x})$. Each population evolves according to the standard replicator equation [1, 26]:

$$\dot{x}_i^k = x_i^k (\pi_i^k(\mathbf{x}) - \bar{\pi}_i(\mathbf{x})), \quad i = 1, \dots, 5, \quad k \in \mathcal{S}_i. \quad (3)$$

The system is an 8-dimensional continuous-time dynamical system (accounting for simplex constraints: $x_1^E + x_1^{AL} + x_1^{AG} = 1$, etc.).

At each pure-strategy corner, the 8×8 Jacobian matrix J describes the linearized dynamics. At a pure corner s^* , the Jacobian admits an exact closed form: for each population i and each strategy $k \neq s_i^*$, the corresponding *invasion eigenvalue* is

$$\lambda_{i,k} = \eta_i \left(\pi_i^k(s^*) - \pi_i^{s_i^*}(s^*) \right), \quad (4)$$

the learning-rate-weighted invasion fitness of strategy k against the incumbent. The full Jacobian at a corner is (up to simultaneous permutation of the reduced coordinates) upper-triangular in these invasion coordinates, so Eq. (4) gives all eight eigenvalues exactly, with no finite-difference error. We classify each corner as:

- **Stable attractor:** All eigenvalues of J have negative real parts.
- **Unstable saddle:** At least one eigenvalue has a positive real part.

Classifying all 96 corners by the analytic criterion (4) yields the steady-state landscape: a map of which institutional configurations are stable, which are traps, and which basins lead to which attractors. Exactly two corners are stable under the baseline parameterization (Stable A and Stable B; Section 3.2); a central-finite-difference evaluation of the same corners (perturbation $\varepsilon = 10^{-6}$, evaluated at the simulated 50,000-step quasi-steady state) produces the identical stable set, cross-validating the analytic classification (Section 3.2).

2.5 Formal Results

We state three formal results that underlie the simulation analysis in Section 3. Complete proofs are available from the authors upon reasonable request (anonymized for double-blind review).

Proposition 1 (Boundary fixed points) *For any parameter configuration, every pure-strategy corner of the five-population game is a fixed point of the replicator dynamics (3).*

Conjecture 2 (Structural instability of interior equilibria) *For generic parameter values, the five-population game admits no asymptotically stable interior mixed equilibrium.*

Proposition 3 (AI penetration paradox; numerical claim) *Let $\lambda \in [0, 3]$ parameterize the intensity of AI-driven cross-effects (the AL payoff advantage; Section 3.13). From the baseline initial condition, there exists a threshold $\lambda^* > 0$ such that for all $\lambda > \lambda^*$, the system converges to a state with P3 HTT frequency exceeding 0.90, while P3 HTN remains below 0.10—where convergence is assessed by the quasi-steady-state (argmax) label at the 50,000-step horizon, so that slow transient episodes (such as the $\lambda = 1.6$ – 1.7 band of Section 3.13) are read as labels at a finite horizon, not as asymptotic classifications. Moreover, the shift is driven primarily by cross-population interactions, not by the base payoff $\text{base}_1^{\text{AL}}$.*

We label Proposition 3 a *numerical claim*: it summarizes a finite-horizon, grid-based simulation observation (50,000 steps over a discrete λ grid; Section 3.13) rather than an analytically proven statement.

3 Results

3.1 The Two-Timescale Dynamic Lock-in: Short-Run Trapping and Long-Run Decay

Definition 1 (Two-timescale dynamic lock-in) Let T_S and T_L denote two simulation horizons with $T_S < T_L$. A dynamic lock-in occurs when the system exhibits two distinct asymptotic regimes:

- At timescale T_S , the system converges to a locally attractive trapping state x_{trap} that appears stable to all practical observations within this horizon.

- At timescale T_L , x_{trap} reveals itself as a *transient* state: it decays toward a different long-run attractor x^* .

Crucially, x_{trap} is *not* a true stable attractor in the asymptotic sense; it is a saddle point with a low-dimensional unstable manifold that becomes visible only when the simulation horizon exceeds T_S . The lock-in is therefore a finite-horizon illusion of stability.

From 200 independently sampled random initial conditions—each population’s initial strategy distribution drawn i.i.d. from a Dirichlet(1) distribution (i.e., uniform on its strategy simplex), with the random seed fixed at 42 for exact reproducibility—the system exhibits a *two-timescale dynamic structure* that reconciles the apparent stability of the current stock market regime with its underlying institutional fragility. The 96 pure-strategy corners are *excluded* from the basin statistics and are instead reported separately below as a boundary-robustness check, because a pure corner is a fixed point of the replicator dynamics: pooling corners with random initial conditions mechanically distorts the basin proportions, since corners remain frozen at short horizons and would inflate their apparent basin share. To establish this two-timescale structure, we evaluate the system at four simulation horizons—5,000, 20,000, 50,000, and 100,000 steps—using 5,000 steps (short-run) and 50,000 steps (long-run) as the two reference horizons and the full sequence as a horizon-convergence check on the basin shares. The choice of 5,000 steps as the short-run horizon is motivated by the observation that the system rapidly converges to a locally attractive HTT regime within this timescale, creating the illusion of institutional stability. The choice of 50,000 steps as the long-run horizon is motivated by the need to capture the basin structure that emerges after the short-run dynamic lock-in has dissipated: at this horizon, the HTT trap reveals itself as a transient trapping state, and the two stable attractors (Stable B and Stable A) dominate the landscape. We recognize that the system converges asymptotically toward the true steady state without reaching it fully within any finite horizon, but the qualitative structure—the dominance of two stable attractors and the transient trapping nature of the HTT trap—is already well established at 50,000 steps. We therefore use 5,000 steps to characterize the short-run dynamic lock-in and 50,000 steps to characterize the long-run basin structure.

Time-scale reading. Under the time-unit convention of Section 2 (one time unit $\approx$ one trading day; one step = 0.03 time units), the short-run horizon of 5,000 steps corresponds to 150 time units ($\approx$ six months) and the long-run horizon of 50,000 steps to 1,500 time units ($\approx$ six calendar years). This reading is a heuristic order-of-magnitude calibration, used only for intuition: all qualitative conclusions depend on the existence of two distinct timescales and their relative ordering, not on the absolute duration. The two-timescale structure—attractor locations, basin boundaries, and the transient trapping nature of the HTT trap—is invariant under any rescaling of the time axis; only the calendar interpretation of the transient trapping lifetime would rescale with it.

Short-run dynamics ($\sim 5,000$ steps): The Dynamic Lock-in. From the baseline initial conditions, the system rapidly converges to the HTT regime within approximately 5,000 steps. This is the *dynamic lock-in* that market participants and short-term policymakers observe: the market appears stable in its current institutional

configuration. At this timescale, the HTT trap is locally attractive and represents a substantial basin category (the AL+TI+HTT+FO+SC state alone commands 29.0% of the random initial conditions, or 46.5% when all HTT-related states are grouped), though no single regime commands a majority. Table 4 shows the top 10 steady states observed at the 5,000-step horizon.

Table 4 Top 10 steady states by short-run basin size (5,000 steps; 200 Dirichlet(1) random initial conditions, seed 42).

Rank	Steady State	Short-Run Basin (%)
1	AL + TI + HTT + FO + SC (HTT Quant Trap)	29.0
2	AL + TI + HTN + FI + SF (Stable B: Mature)	15.0
3	AL + NS + LTT + FO + SC (Stable A: Traditional)	14.0
4	AL + TI + HTT + FO + SF (Short-Sell Hybrid)	8.5
5	E + NS + LTT + FO + SC (Pre-AI Traditional)	8.0
6	E + TI + HTT + FO + SC (Current Stock Market Trap)	5.0
7	E + NS + LNT + FO + SC (No-Price-Limit Hybrid)	4.0
8	AL + TI + LTT + FO + SC (Low-AI Quant Trap)	3.5
9	AL + NS + LNT + FO + SC (Low-AI Retail Hybrid)	2.5
10	AL + TI + HTN + FO + SF (HTN Short-Sell Hybrid)	2.5

Long-run dynamics (~50,000 steps): The Slow Transient Decay. When the simulation horizon extends to 50,000 steps, the system continues to evolve. The HTT trap, which appeared stable at 5,000 steps, reveals itself as a *transient trapping state*: it slowly decays toward the traditional-market stable attractor (Stable A). The two true long-run stable attractors are Stable A (traditional market, 36.5% basin) and Stable B (mature market, 18.0% basin) at this horizon. These shares are based on the 200 Dirichlet(1) random initial conditions; the horizon-convergence analysis below shows that the system is still evolving at 50,000 steps. Table 5 shows the steady states observed at the 50,000-step horizon.

Table 5 Steady states by long-run basin size (50,000 steps; 200 Dirichlet(1) random initial conditions, seed 42).

Rank	Steady State	Long-Run Basin (%)
1	AL + NS + LTT + FO + SC (Stable A: Traditional)	36.5
2	E + NS + LTT + FO + SC (Pre-AI Traditional)	20.0
3	AL + TI + HTN + FI + SF (Stable B: Mature)	18.0
4	E + TI + HTT + FO + SC (Current Stock Market Trap)	16.0
5	E + TI + HTT + FI + SC (AI-Quant-Foreign Hybrid)	8.5
6	E + NS + HTT + FO + SC (HTT No-AI Hybrid)	0.5
7	E + TI + LTT + FI + SC (LTT Foreign-Inflow Hybrid)	0.5

Note: The 7 states account for 100% of the 200 random initial conditions.

At the 50,000-step horizon the 200 random initial conditions distribute across exactly 7 argmax states (Table 5). The horizon-convergence analysis (5,000 / 20,000 / 50,000 / 100,000 steps) shows that the basin structure is still evolving at 50,000 steps: Stable A’s basin share rises from 14.0% (5,000) through 27.5% (20,000) and 36.5% (50,000) to 51.0% (100,000 steps), while Stable B’s share has already converged to 18.0% by 20,000 steps; the HTT-related transient states shrink from 46.5% (5,000) through 45.0% (20,000) and 25.0% (50,000) to 5.5% (100,000 steps), confirming their transient trapping nature. As a boundary-robustness check, the 96 pure-strategy corners were simulated separately over the same four horizons: at 5,000 steps they remain essentially frozen at their own vertices (95 distinct argmax states), confirming that corners contribute no genuine basin mass at short horizons, while by 100,000 steps the numerically integrated corner trajectories migrate to the same attractor set as the random sample—27 to Stable B, 24 to Stable A, 32 to the HTT trapping family (17 to E+TI+HTT+FO+SC, 7 to E+TI+HTT+FI+SC, 6 to E+NS+HTT+FO+SC, 2 to AL+TI+HTT+FO+SC), and 13 to E+NS+LTT+FO+SC, the passive-retail (E) variant of Stable A—corroborating the robustness of the two-attractor structure at the boundary of the state space. To assess seed dependence, we replicated the full protocol on two additional seeds (123 and 7; 200 Dirichlet(1) ICs each): at 50,000 steps Stable A’s basin share is 36.5%/38.0%/30.0% across seeds 42/123/7 (mean $\pm$ range: 34.8% $\pm$ 8.0 percentage points) and Stable B’s share is 18.0%/19.0%/22.0% (19.7% $\pm$ 4.0pp); the seed-42 Wilson 95% confidence intervals are [30.1, 43.4] for Stable A and [13.3, 23.9] for Stable B, so the qualitative basin ordering is robust to the sampling seed. The two-timescale structure (Fig. 1)—where the short-run landscape is dominated by the HTT trap while the long-run landscape approaches two stable attractors—is the central feature of the dynamic lock-in.

3.2 The Two Long-Run Stable Attractors: Institutional Complementarity

Among the 96 steady states, only two are true stable attractors in the long-run horizon ($\sim 50,000$ steps). Strikingly, both feature AI-driven retail ($P1 = AL$). The institutional difference lies entirely in the complementary stack of $P2$ – $P3$ – $P4$ – $P5$.

Stable B: AL + TI + HTN + FI + SF (Mature Market)

- Basin: 18.0% (50,000-step horizon, 200 Dirichlet(1) random ICs, seed 42; converged already by 20,000 steps; three-seed replication 18.0%/19.0%/22.0%, mean $\pm$ range 19.7% $\pm$ 4.0pp, Section 3.1)
- P1: AL = 1.000 (AI-driven retail)
- P2: TI = 1.000 (quant institutions)
- P3: HTN = 1.000 (high AI, no T+1, no price limits)
- P4: FI = 1.000 (foreign capital inflow)
- P5: SF = 1.000 (short-selling free)
- Jacobian (analytic invasion eigenvalues, Eq. (4)): max eig = -0.015 (all 8 eigenvalues < 0)

Stable A: AL + NS + LTT + FO + SC (Traditional Market)




Fig. 1 Basin-of-attraction shares from 200 Dirichlet(1) random initial conditions (seed 42; the 96 pure-strategy corners are excluded and reported separately as a boundary-robustness check, see text). Bars show the exact share of initial conditions whose trajectory ends at each argmax state; unlisted states are pooled as “Other”. **(a) Short-run (5,000 steps):** HTT-related states dominate (46.5% combined; the HTT Quant Trap AL+TI+HTT+FO+SC alone commands 29.0%), creating the *dynamic lock-in* observed by market participants. **(b) Long-run (50,000 steps):** the HTT-related states shrink to 25.0% (and to 5.5% by 100,000 steps) while Stable A (AL+NS+LTT+FO+SC, 36.5%) becomes the largest basin and Stable B (AL+TI+HTN+FI+SF, 18.0%) has already converged, revealing the transient trapping nature of the apparent short-run stability.

- Basin: 36.5% (50,000-step horizon, 200 Dirichlet(1) random ICs, seed 42; still rising—51.0% at 100,000 steps; three-seed replication 36.5%/38.0%/30.0%, mean $\pm$ range $34.8\% \pm 8.0\text{pp}$, Section 3.1)
- P1: AL = 1.000 (AI-driven retail)
- P2: NS = 1.000 (state-controlled institutions)
- P3: LTT = 1.000 (low AI, T+1, price limits)
- P4: FO = 1.000 (foreign capital outflow)
- P5: SC = 1.000 (short-selling constrained)
- Jacobian (analytic invasion eigenvalues, Eq. (4)): max eig = -0.006 (all 8 eigenvalues < 0)

Jacobian reporting convention. Throughout Section 3, corner stability is classified by the *analytic invasion eigenvalues* of Eq. (4), evaluated exactly at each pure corner; under the baseline parameterization exactly two of the 96 corners are stable (Stable A and Stable B). As a cross-validation, we also computed the Jacobian by central finite differences of the 8-dimensional reduced system (one component per simplex, perturbation $\varepsilon = 10^{-6}$), evaluated at the *simulated 50,000-step quasi-steady state*—which carries the 10^{-8} frequency floor of the integration loop—rather than at the exact pure corner. The finite-difference classification agrees with the analytic one for all 96 corners. Because the finite-difference values are evaluated at a quasi-steady state rather than the exact fixed point, they carry a residual error: for Stable A the finite-difference estimate (-0.003) systematically underestimates the analytic value (-0.006) by about half, while for Stable B the two agree (-0.015). We therefore report the analytic values throughout: Stable A, max eig = -0.006 ; Stable B, max eig = -0.015 . Stable A remains “weakly stable” in the qualitative sense that its

slowest decay mode is several times slower than Stable B’s, but the conclusions in this paper rely only on the sign of the eigenvalues (all negative, hence local attractors)—and with the analytic values the margin below zero is twice as large as the earlier finite-difference estimate suggested.

Key Insight: Both stable states have AI-driven retail. The difference is not whether retail adopts AI but what institutional environment surrounds AI adoption. In the traditional stack, AI is used by retail but is harvested by state-controlled institutions within a constrained market. In the mature stack, AI is complemented by quant institutions, foreign capital, and hedging tools.

3.3 The HTT Trap: Current Stock Market Regime

From the baseline initial condition, the system rapidly converges to an unstable saddle state on the short timescale ($\sim 5,000$ steps): P1:E + P2:TI + P3:HTT + P4:FO + P5:SC. This is the “current stock market trap” observed by market participants and short-term policymakers. Throughout the paper, the term “HTT trap” denotes a *family* of adjacent trapping states that share P3 = HTT; the exact P1/P4/P5 (and P2) identities vary with the parameterization and the horizon—for example, the baseline trap is E+TI+HTT+FO+SC at 5,000 steps, while the AL+TI+HTT+FO+SC variant appears in the parameter scans of Sections 3.5 and 3.13:

- Retail investors remain passive (E) while quant institutions dominate (TI)
- But T+1 and price limits remain (HTT)
- Foreign capital is absent (FO)
- Short-selling is constrained (SC)

On the long timescale ($\sim 50,000$ steps), however, this HTT trap reveals itself as a transient trapping state: it slowly decays toward the traditional-market stable attractor (Stable A: AL+NS+LTT+FO+SC). The counterfactual baseline simulation (Table 6) confirms that after 50,000 steps, the system converges to Stable A rather than remaining in the HTT trap. This state is locally attractive in the short run but globally transient: the system converges to it from the baseline mixed initial condition at 5,000 steps but diverges toward Stable A by 50,000 steps. The HTT trap’s attraction is primarily local and temporal, centered on the baseline IC and nearby mixed initial conditions at the short-run horizon.

3.4 Counterfactual Simulations: Single-Point Reforms Fail

We test seven counterfactual scenarios from the baseline initial condition (P1: 35% E, 35% AL, 30% AG; P2: 40% NS, 60% TI; P3: 30% LTT, 10% LNT, 50% HTT, 10% HTN; P4: 60% FI, 40% FO; P5: 80% SC, 20% SF).

Full Foreign Capital Open (reported separately). A seventh scenario is deliberately excluded from Table 6 because it is not a like-for-like counterfactual. The six scenarios above are genuine reform shocks: they perturb payoffs or rules starting from the common baseline initial condition. Full Foreign Capital Open, by contrast, is implemented by *replacing the initial condition*—setting P4 = FI with frequency

Table 6 Counterfactual simulation results (50,000 steps): six single-point reform scenarios, all launched from the baseline initial condition. The complementary strategies (P2 NS, P3 LTT, P4 FO, P5 SC) are implied by the sum-to-one constraint. The seventh scenario, Full Foreign Capital Open, is reported separately below the table because it replaces the initial condition rather than applying a reform shock from the baseline, and is therefore not comparable to the six scenarios listed here.

Scenario	P1 AL	P2 TI	P3 HTN	P4 FI	P5 SF
Baseline	0.999	0.000	0.000	0.000	0.000
Remove T+1	1.000	0.000	0.000	0.000	0.000
Ban AI	0.000	0.000	0.000	0.000	0.000
Short-Selling Liberalization	0.866	1.000	0.000	0.000	1.000
Retail Protection	0.000	0.000	0.000	0.000	0.000
Tighten Reg	0.000	0.000	0.000	0.000	0.000

Note: The complementary strategies (P2 NS, P3 LTT, P4 FO, P5 SC) are implied by the sum-to-one constraint.

1.0, i.e., a market already fully committed to foreign capital inflow—and then letting the dynamics run. From that starting point the system converges to Stable B (AL+TI+HTN+FI+SF), the mature market, with all five reported strategy shares equal to 1.000. The scenario is therefore an *existence proof*, not a policy experiment: it demonstrates that the mature-market attractor exists in the basin landscape and can be reached when complementary institutions align, but it says nothing about whether foreign-capital liberalization alone suffices from arbitrary (in particular, the baseline) initial conditions. We stress, more generally, that all “reform success” statements in this paper—here and in the transition-barrier analysis of Section 3.9—are *conditional on the reported families of initial conditions* (the baseline mixed IC, perturbed corners, and the Dirichlet-sampled ICs of Section 3.1), not unconditional probabilities of reform success.

Key Findings (long-run, 50,000 steps):

1. **Baseline:** After 50,000 steps, the system converges to Stable A (AL+NS+LTT+FO+SC), not the HTT trap (Fig. 2). This confirms the transient trapping nature of the HTT regime: what appears as a stable trap at 5,000 steps decays toward the traditional-market attractor in the long run.
2. **Remove T+1 alone:** Despite the removal of T+1, the system converges to Stable A (AL+NS+LTT+FO+SC) in the long run. T+1 removal alone does not provide the escape velocity to reach the mature market; it merely shifts the trajectory toward the traditional stable attractor rather than the HTT trap.
3. **Ban AI:** The system converges to the pre-AI traditional market (E+NS+LTT+FO+SC), the only counterfactual that eliminates AI-driven retail entirely. By removing the technology that destabilizes the traditional stack, the system settles into the pre-AI traditional state.
4. **Full Foreign Capital Open:** The system converges to Stable B (AL+TI+HTN+FI+SF), the mature market. This is the only counterfactual that reaches the mature market in the long run, but it requires the initial condition to be fully committed to foreign capital inflow ($P4 = FI$). The policy implication is that foreign capital liberalization can be a path to the mature market, but only

when combined with a simultaneous shift toward quant-friendly institutions (TI) and mature regulations (HTN).

5. **Short-Selling Liberalization:** The system settles in an unstable mixed state (AL+TI+HTT+FO+SF) with P1 AL = 0.866, P2 TI = 1.000, P3 HTN = 1.000, P4 FO = 1.000, P5 SF = 1.000. Short-selling liberalization alone does not push the system toward the mature market; instead, it creates a “quant-friendly but institutionally immature” unstable state where institutions can hedge but retail investors are still locked in T+1 + price limits [27, 28].
6. **Retail Protection and Tighten Reg:** Both scenarios converge to the pre-AI traditional market (E+NS+LTT+FO+SC), confirming that targeted retail protection or tighter regulation, without complementary institutional reform, locks the system into the traditional stable state.




Fig. 2 Counterfactual simulation results (50,000 steps). Each bar group represents one scenario; colored bars show final shares of key strategies (P1 AL, P2 TI, P3 HTN, P4 FI, P5 SF). Full Foreign Capital Open reaches the mature market (Stable B, AL+TI+HTN+FI+SF); Ban AI, Retail Protection, and Tighten Reg converge to the pre-AI traditional market (E+NS+LTT+FO+SC); Short-Selling Liberalization creates a new unstable trap (AL+TI+HTT+FO+SF). The Baseline and Remove T+1 scenarios both converge to Stable A (AL+NS+LTT+FO+SC). See the separate note following Table 6 for the existence-proof status of Full Foreign Capital Open.

3.5 Parameter Sensitivity: The HTT Trap is Robust

Scanning six key parameters from the baseline initial condition reveals a rich pattern of regime transitions (Fig. 3) that confirms the transient trapping nature of the HTT trap and identifies the narrow conditions under which the mature market becomes accessible:

Sensitivity to $\text{base}_1^{\text{AL}}$ (Retail AI Base Payoff):

- At low values ($\text{base}_1^{\text{AL}} \leq 0.011$), the system converges to the pre-AI traditional market (E+NS+LTT+FO+SC), with P1 dominated by E (≈ 1.0). This confirms that without sufficient AI base payoff, the system reverts to the pre-AI traditional state.

- At intermediate values ($\text{base}_1^{AL} = 0.016$ to 0.033), the system converges to Stable A (AL+NS+LTT+FO+SC). In this regime, AI-driven retail exists (P1 AL ≈ 0.87 – 0.99) but is constrained by traditional institutions (NS, LTT, FO, SC), creating a hybrid stable state.
- At high values, the system shifts to HTT-related states: at $\text{base}_1^{AL} \approx 0.038$ it converges to an HTT-related unstable state (E+TI+HTT+FO+SC), and from $\text{base}_1^{AL} \geq 0.044$ to the fully developed HTT trap (AL+TI+HTT+FO+SC), with P1 AL ≈ 0.99 and P2 TI = 1.0. High AI base payoff, when combined with quant institutions (TI), pushes the system toward the HTT regime rather than the mature market.
- The mature market (Stable B) is absent across the entire base_1^{AL} range, confirming that AI base payoff alone cannot drive institutional modernization.

Sensitivity to base_3 -HTT (HTT Base Payoff):

- At low values ($\text{base}_3^{HTT} = 0.04$ to 0.048), the system settles in unstable mixed states dominated by HTT (P3 HTT ≈ 0.50 – 0.86), with either E or AL as the dominant P1 strategy. The HTT regime is present but not fully consolidated.
- At intermediate values ($\text{base}_3^{HTT} = 0.057$ to 0.093), the system converges to Stable A (AL+NS+LTT+FO+SC). This is the most surprising finding: raising the HTT base payoff beyond a threshold (≈ 0.055) does not reinforce the HTT trap but instead pushes the system toward the traditional stable attractor. The mechanism is that higher HTT payoff strengthens the HTT regime in the short run, but on the long timescale (50,000 steps) the strengthened HTT regime decays more rapidly toward Stable A—a pattern consistent with the HTT saddle losing persistence relative to the traditional attractor, though we have not directly compared eigenvalues across the two parameterizations.
- At high values ($\text{base}_3^{HTT} \geq 0.102$), the system converges to the pre-AI traditional market (E+NS+LTT+FO+SC, unstable). The “unstable” label is a finite-horizon transient classification: the terminal argmax is this corner, but its Jacobian’s dominant eigenvalue is positive, so it is not a stable attractor in the sense of Section 2.4. Extreme HTT base payoff collapses the HTT trap entirely, leaving only the pre-AI market as the dominant trajectory.
- The mature market (Stable B) is absent across the entire base_3^{HTT} range, confirming that T+1 + price limit payoff alone cannot drive institutional modernization.

Sensitivity to base_5 -SF (Short-Selling Free Base Payoff):

- At low values ($\text{base}_5^{SF} = 0.01$ to 0.06), the system converges to Stable A (AL+NS+LTT+FO+SC). Short-selling constraints, even with low payoff, do not prevent the system from settling in the traditional stable attractor.
- Within a transitional band between the traditional trap and the HTT trap ($0.06 < \text{base}_5^{SF} < 0.08$), the system converges to Stable B (AL+TI+HTN+FI+SF), the mature market. This is the first parameter scan to reveal the mature market: at the only sampled point inside this band ($\text{base}_5^{SF} = 0.07$), the system crosses the basin boundary from Stable A to Stable B; the scan resolution does not support a sharper localization of the band edges. The mature market requires not just short-selling

freedom but a sufficient base payoff for SF to overcome the institutional inertia of the traditional stack.

- At high values ($\text{base}_5^{SF} \geq 0.08$), the system settles in unstable HTT-related states (AL+TI+HTT+FO+SF or AL+TI+HTT+FI+SF). Excessive short-selling liberalization, without complementary institutional reform (T+1 removal and price limit abolition), shifts the system toward the HTT trap rather than the mature market. The system creates a “quant-friendly but institutionally immature” unstable state.
- The mature market is accessible only from a narrow transitional band ($0.06 < \text{base}_5^{SF} < 0.08$; Table 7), confirming that institutional modernization requires reform intensity within a bounded window, not just directional adjustment.

Additional Findings from Full 6-Parameter Scan:

- **base_4^{FI} (Foreign Capital Inflow Base Payoff):** At low values (-0.01 to 0.043), the system converges to Stable A. At intermediate values (0.061 to 0.114), the system settles predominantly in unstable HTT-related mixed states, though a narrow stable window exists at approximately 0.079 ($\text{max_eig} \approx -0.0004$). At high values (0.132 to 0.15), the system converges to Stable B (mature market). This confirms that foreign capital liberalization, when sufficiently strong, can be a path to the mature market—but only when the base payoff is high enough to overcome the traditional institutional inertia.
- **$\text{cross}_{13}^{AL,HTT}$ (AL-HTT Cross-Effect):** At strongly negative values (-0.08 to -0.024), the system converges to Stable A. At weakly negative to positive values (-0.013 to 0.02), the system converges to the HTT trap. The transition threshold is approximately -0.013 : when the cross-effect is sufficiently negative (AL and HTT are strong substitutes), the system avoids the HTT trap and settles in Stable A; when the cross-effect is weakly negative or positive, AL and HTT become complements, reinforcing the HTT trap.
- **base_1^E (Retail Passive Base Payoff):** At low values (-0.03), the system settles in an unstable HTT state. At intermediate values (-0.021 to -0.007), the system converges to Stable A. At high values (≥ -0.003), the system converges to the pre-AI traditional market. The pattern is non-monotonic: extreme low and high values both avoid Stable A, but for different reasons (low favors HTT, high favors pre-AI).

Finer-Scale Perturbation. A $\pm 50\%$ perturbation test of the three structurally critical parameters (base_1^{AL} , base_3^{HTT} , base_5^{SF}) confirms that the two-attractor structure is not a fragile calibration artifact; we report this test together with the full perturbation battery in Section 3.11.

3.6 Bifurcation Analysis: Critical Thresholds

Scanning four key parameters across 96 corner initial conditions (ICs) + 10 representative mixed ICs reveals critical thresholds (Fig. 4) that define the policy safety corridor and confirm the Goldilocks nature of institutional reform:

Threshold 1: $\text{base}_3\text{-HTT} \approx 0.111$ — The Disappearance of the Mature Market. Across the entire tested range of base_3^{HTT} (0.04 to 0.12), the two-attractor structure (with both Stable A and Stable B present) is preserved for most values.




Fig. 3 Parameter sensitivity analysis (50,000 steps, 10 points per parameter). (Left) As base_1^{AL} increases from 0.0 to 0.05, the system transitions from Pre-AI traditional (E+NS+LTT+FO+SC) to Stable A (AL+NS+LTT+FO+SC) at ≈ 0.016 , then to HTT-related unstable states (E+TI+HTT+FO+SC) at ≈ 0.038 , and to the fully developed HTT trap (AL+TI+HTT+FO+SC) at ≥ 0.044 . The mature market (Stable B) is absent. (Center) As base_3^{HTT} increases from 0.04 to 0.12, the system transitions from HTT-dominated unstable states to Stable A at ≈ 0.055 , then to Pre-AI at ≈ 0.102 . Higher HTT payoff paradoxically pushes the system toward the traditional market in the long run, confirming metastability. (Right) As base_5^{SF} increases from 0.01 to 0.10, the system transitions from Stable A to Stable B (mature market) within the transitional band $0.06 < \text{base}_5^{SF} < 0.08$ (the only sampled point in the band is 0.07), then to HTT-related unstable states at ≥ 0.08 . The mature market is accessible only from this narrow band, demonstrating the Goldilocks nature of institutional reform.

Table 7 Parameter windows for reaching the mature market (Stable B). Values show the parameter range from the baseline initial condition at the long-run horizon (50,000 steps). The Stable B column for base_5^{SF} reports the *transitional band* between the traditional trap and the HTT trap, not a precise optimal point: the mature-market outcome is observed only within the band $0.06 < \text{base}_5^{SF} < 0.08$, and the scan resolution does not support a sharper localization.

Parameter	Traditional Trap (Stable A)	Transitional Band (Stable B)	HTT Trap
base_5^{SF}	0.01–0.06	$0.06 < \text{base}_5^{SF} < 0.08$	≥ 0.08
base_4^{FI}	–0.01–0.043	≥ 0.132	0.061–0.114

However, at the extreme high end ($\text{base}_3^{HTT} \geq 0.111$), the mature market (Stable B) disappears entirely: Stable B becomes unstable, and only one stable attractor remains (the HTT-related state E+TI+HTT+FI+SC, reached from Mixed ICs). This confirms that the mature market is not universally accessible: extreme HTT base-payoff values collapse the landscape to a single HTT-dominated regime. The policy implication is that T+1 + price limit reform, if pushed too aggressively (i.e., making the HTT regime too attractive), can paradoxically eliminate the mature market by strengthening the HTT trap to the point where it becomes the only stable state.

Threshold 2: $\text{base}_3\text{-HTT} \approx 0.084\text{--}0.093$ — The Disappearance of the Traditional Market. As base_3^{HTT} increases beyond the baseline (0.075), Stable A (the traditional market) gradually weakens. At $\text{base}_3^{HTT} \approx 0.084$, Stable A becomes unstable for some ICs (HTT_Trap and Pre_AI corners no longer converge to Stable A). By $\text{base}_3^{HTT} \approx 0.093$, Stable A has disappeared entirely as a stable attractor (has_trad

= false), and the HTT-dominated state (E+TI+HTT+FI+SC) becomes a new stable attractor (Table 8). The system retains Stable B (mature market) but gains an additional HTT-dominated attractor, creating a “mature-or-HTT” regime. This confirms that raising the HTT base payoff beyond a threshold destabilizes the traditional market but does not eliminate the HTT trap; instead, the HTT trap hardens into a permanent attractor.

Threshold 3: $\text{base}_5^{\text{SF}} \approx 0.075\text{--}0.08$ — The Disappearance of the Traditional Market (Short-Selling Channel). As $\text{base}_5^{\text{SF}}$ increases, the system transitions from a two-attractor landscape (Stable A + Stable B) to a one-attractor landscape (Stable B only). At $\text{base}_5^{\text{SF}} \approx 0.075$, Stable A is still present for some ICs. By $\text{base}_5^{\text{SF}} \approx 0.08$, Stable A has disappeared entirely ($\text{has_trad} = \text{false}$), and only Stable B remains as a stable attractor. This is the flip side of the sensitivity finding: short-selling liberalization, when sufficiently strong, can eliminate the traditional market trap and push the system toward the mature market. However, the transition is abrupt: at $\text{base}_5^{\text{SF}} = 0.07$, the system still has two stable attractors; at $\text{base}_5^{\text{SF}} = 0.08$, only one remains. This narrow transition window confirms the Goldilocks nature of reform: the parameter must be precisely calibrated to navigate between the two traps. The apparent tension with the sensitivity finding of Section 3.5—where $\text{base}_5^{\text{SF}} \geq 0.08$ leads to *unstable* HTT-related states rather than Stable B—resolves through the experimental design: the sensitivity scan launches from the baseline mixed initial condition and classifies trajectory endpoints, whereas the bifurcation scan evaluates attractor existence across 96 corner ICs plus 10 representative mixed ICs. Corner stability and mixed-IC trajectory endpoints answer different questions: at $\text{base}_5^{\text{SF}} = 0.08$, Stable B is the only stable attractor, but the baseline mixed IC does not lie in its basin.

Threshold 4: $\text{cross}_{13}^{\text{AL,HTT}} \approx -0.013$ — The HTT Trap Becomes Stable. When the cross-effect between AL and HTT is strongly negative ($\text{cross}_{13}^{\text{AL,HTT}} \leq -0.024$), AL and HTT are substitutes: the system avoids the HTT trap and settles in Stable A or Stable B. When the cross-effect weakens to -0.013 or above, AL and HTT become complements, and the HTT trap (AL+TI+HTT+FO+SC) becomes a stable attractor. The number of stable attractors increases from 2 to 3, creating a three-attractor landscape. This is a structurally important threshold: when AI-driven retail and HTT institutions are complementary rather than substitutive, the system gains an additional stable trap that competes with the traditional and mature markets. The policy implication is that the relationship between AI adoption and institutional reform is not monotonic: if AI and HTT are strong substitutes (negative cross-effect), reform can succeed; if they are weak substitutes or complements ($\text{cross-effect} \geq -0.013$), the HTT trap becomes self-reinforcing.

Threshold 5: $\text{base}_1^{\text{E}} \approx 0.001$ — The Traditional Market Disappears. As the retail passive base payoff increases from strongly negative (-0.03) to positive (0.01), the system undergoes a regime transition. At low values, Stable A and Stable B coexist. At $\text{base}_1^{\text{E}} \approx -0.003$, the pre-AI market (E+NS+LTT+FO+SC) becomes stable, creating a three-attractor landscape. At $\text{base}_1^{\text{E}} \geq 0.001$, Stable A disappears, and the system transitions to a two-attractor landscape dominated by the pre-AI market and Stable B. This confirms that excessive retail passivity (high base_1^{E}) eliminates the

traditional AI-hybrid market, leaving only the pre-AI and mature markets as stable options.

Observation: $\text{base}_5^{\text{SF}}$ — No Emergence Threshold Detected in Scan Range. Scanning $\text{base}_5^{\text{SF}}$ from 0.01 to 0.10 reveals that the mature-market stable state exists across the entire range. Unlike $\text{base}_3^{\text{HTT}}$, $\text{base}_5^{\text{SF}}$ does not trigger a bifurcation within the tested range. This confirms that short-selling liberalization is neither a necessary nor sufficient condition for reaching the mature market.




Fig. 4 Bifurcation analysis (50,000 steps, 10 points per parameter $\times$ 10 ICs). (Top-left) As base_1^F increases, the system transitions from a two-attractor landscape (Stable A + Stable B) to a three-attractor landscape (Pre-AI + Stable A + Stable B) at ≈ -0.003 , then to a two-attractor landscape (Pre-AI + Stable B) at ≥ 0.001 . (Top-right) As $\text{base}_3^{\text{HTT}}$ increases, the mature market (Stable B) disappears at ≈ 0.111 , and the traditional market (Stable A) disappears at ≈ 0.084 – 0.093 , creating a “mature-or-HTT” regime. (Bottom-left) As $\text{base}_5^{\text{SF}}$ increases, the traditional market disappears at ≈ 0.075 – 0.08 , leaving only the mature market. The transition is abrupt, confirming the Goldilocks nature of reform. (Bottom-right) As $\text{cross}_{13}^{\text{AL,HTT}}$ increases from strongly negative to weakly negative, the HTT trap becomes a stable attractor at ≈ -0.013 , increasing the number of stable attractors from 2 to 3. The gray dashed line marks the baseline value for each parameter.

3.7 Stochastic Dynamics: Robustness Under Noise

To test whether the stable attractors are robust to stochastic perturbations, we inject additive noise ($\sigma = 0.5\%$, 1% , 2%) into the dynamics from each stable state and run 30 independent replications for 50,000 steps each, with the random seed fixed (seed = 42) so that all results are exactly reproducible. The experiment uses the correct attractor

Table 8 Bifurcation threshold summary (50,000 steps, 96 corner initial conditions (ICs) + 10 mixed ICs).

Parameter	Threshold	n_{stable} (Before→After)	Description
base_1^E	-0.003	2 → 3	Pre-AI market emerges as stable attractor
base_1^E	0.001	3 → 2	Stable A disappears; traditional retail exits
base_3^{HTT}	0.093	2 → 2	HTT-dominated state replaces Stable A, creating a two-attractor landscape (Stable B + HTT-related mixed)
base_3^{HTT}	0.111	2 → 1	Mature market collapses; only HTT remains
base_5^{SF}	0.080	2 → 1	Stable A disappears; only Stable B remains
$\text{cross}_{13}^{AL,HTT}$	-0.013	2 → 3	HTT trap becomes a stable attractor

corners as identified by the deterministic stability analysis—in particular, Stable A is the AL+NS+LTT+FO+SC corner. An escape is recorded when the final argmax strategy combination deviates from the starting combination, and we report the distribution (min/median/max) of the first-escape time in steps. The noise is implemented as persistent frequency shocks: at every integration step, i.i.d. $\text{normal}(0, \sigma)$ perturbations are added to every strategy frequency in all five populations, the perturbed frequencies are clipped to $[10^{-8}, 1]$ and renormalized onto the simplex, and the deterministic replicator update is then applied to the perturbed state. This large-sample design allows us to characterize the *statistical stability* of the steady states: under noise, neither state is permanent; instead, each attractor has a *finite stochastic residence time*, and the system executes spontaneous transitions between basins.

Key Findings (30 replications, seeded):

1. **Stable A escapes at high rates at all noise levels (Fig. 5):** At $\sigma = 0.005$, the escape rate is 93.3% (28/30 trajectories), with first-escape times of min/median/max = 1,286/5,901/23,268 steps (the median, 5,901 steps ≈ 177 time units, reads as roughly 8–9 months under the time-unit convention of Section 2), indicating that the traditional-market attractor is statistically fragile even under modest perturbation. At $\sigma = 0.01$, the escape rate is 100% (30/30), with first-escape times of 278/910/4,087 steps. At $\sigma = 0.02$, the escape rate remains 100% (30/30), with first-escape times of 75/246/623 steps. Stable A is not a permanent equilibrium but a statistical steady state whose residence time collapses rapidly with noise intensity.
2. **Stable B is markedly more persistent at low noise, but equally fragile at moderate noise:** At $\sigma = 0.005$, the escape rate is only 40.0% (12/30)—a majority of trajectories (60%) remain at the mature-market attractor for the full 50,000 steps—and even the escaping trajectories persist long (first-escape min/median/max = 3,633/17,445/34,939 steps). At $\sigma = 0.01$, the escape rate rises to 100% (30/30), with first-escape times of 383/1,177.5/3,360 steps. At $\sigma = 0.02$, the escape rate is 100% (30/30), with first-escape times of 95/272.5/676 steps. The two attractors therefore differ *quantitatively, not symmetrically*: under low noise Stable B is substantially “stickier” than Stable A (40.0% vs. 93.3% escape), and the

claim that the two attractors are equally transient does not survive the corrected data—Stable B is a genuinely persistent low-noise state, while both states degrade comparably fast once $\sigma \geq 0.01$.

Interpreting σ and the deterministic–stochastic relationship. The noise term σ in our stochastic extension should not be confused with stock-market return volatility. σ is a *strategy-frequency noise*: the replicator dynamic is subject to persistent additive shocks (i.i.d. normal perturbations applied to every strategy frequency at every step) that push the trajectory away from the deterministic path. This is standard in evolutionary-game Monte Carlo experiments [29]. The cumulative effect of $\sigma = 0.005$ over the full horizon $T = 50,000$ steps is of order $\sigma\sqrt{T} \approx 0.005 \times \sqrt{50,000} \approx 1.12$; over the empirically observed median residence time of Stable A at this noise level (5,901 steps), the accumulated displacement is $\sigma\sqrt{t} \approx 0.005 \times \sqrt{5901} \approx 0.38$, which is comparable to the typical basin width observed in our bifurcation analysis (Section 3.5). Thus the stochastic simulation is not a separate “noise model” but a *stress test* of the deterministic basin structure: trajectories that are initially inside a basin can be pushed across the boundary by accumulated noise, and the jump rate measures the *mean first-passage time* of this process.

Deterministic basin size vs. stochastic jump rate. The deterministic analysis (Sections 3.2–3.6) measures *basin size*: the fraction of initial conditions that converge to a given attractor under noise-free dynamics. This is a “best-case” measure of accessibility. The stochastic analysis (this section) measures the *jump rate*: the fraction of trajectories that escape the attractor within 50,000 steps when subject to frequency noise. The two measures capture different aspects of institutional fragility: a large deterministic basin does not guarantee a long stochastic residence time, because a single well-placed perturbation can push the system into a different basin. Our results show that Stable A has the larger basin (36.5%) but also a high escape rate (93.3% at $\sigma = 0.005$, rising to 100% at $\sigma \geq 0.01$), confirming that the traditional market is easily reached yet fragile under noise. Stable B has a smaller basin (18.0%) but a much lower low-noise escape rate (40.0% at $\sigma = 0.005$, with a median residence of 17,445 steps among escapees): the mature-market basin is harder to reach but considerably *stickier* once reached. Stable B too becomes fragile at $\sigma = 0.01$, where all 30 trajectories escape with a median first-escape time of 1,177.5 steps. (These shares are based on 200 Dirichlet(1) random initial conditions, seed 42, with corners excluded; a pooled-corner sampling design would mechanically distort the basin proportions—see Section 3.1.) The policy-relevant insight is that the two attractors are not equally transient: under low noise Stable B is a genuinely persistent state while Stable A is already fragile, and only under moderate-to-strong noise do both converge to short-lived statistical steady states.

3.8 Two-Dimensional Policy Space: The HTT Trap Dominates

To test whether any combination of two policy parameters can escape the HTT trap, we conduct a two-dimensional scan of $\text{base}_3^{\text{HTT}} \times \text{base}_5^{\text{SF}}$ from 64 systematically sampled parameter combinations (8×8 grid: $\text{base}_3^{\text{HTT}}$ from 0.040 to 0.120, $\text{base}_5^{\text{SF}}$ from 0.010 to 0.100, Fig. 6), evaluated at the long-run (50,000 steps) horizon. This




Fig. 5 Stochastic dynamics under persistent additive noise (30 seeded replications, seed = 42, 50,000 steps each), using the corrected Stable A corner (AL+NS+LTT+FO+SC). **(a)** Escape rates (final argmax combination deviates from the starting combination). At $\sigma = 0.005$, Stable A escapes at 93.3% (28/30) while Stable B escapes at only 40.0% (12/30), indicating that the mature-market attractor is substantially more persistent under low noise; at $\sigma = 0.01$ and $\sigma = 0.02$ both attractors escape at 100%. **(b)** Distribution of first-escape times (steps, log scale; boxes show the interquartile range, whiskers the min–max of escaping trajectories). Median first-escape times at $\sigma = 0.005/0.01/0.02$ are 5,901/910/246 steps for Stable A and 17,445/1,177.5/272.5 steps for Stable B. Reform must occur during the transient convergence phases between perturbations.

design reveals the *two-timescale dynamic lock-in* in the policy space: what appears inaccessible in the short run may become accessible in the long run, but only under sustained parameter commitment.

Key Findings:

- 1. Short-run (5,000 steps): Universal trapping.** At the 5,000-step horizon, all sampled parameter combinations remain trapped in the HTT regime or related unstable states. None reaches the mature market (Stable B). This confirms that the dynamic lock-in is operative in the policy space: even when two institutional levers are adjusted simultaneously, the system is captured by the short-run HTT attractor before any long-run restructuring can occur.
- 2. Long-run (50,000 steps): 9.4% escape to Stable B.** When the horizon extends to 50,000 steps, 6 out of 64 parameter combinations (9.4%) converge to Stable B (mature market). The remaining 90.6% are distributed across four outcome categories: HTT-related states (31.2% total, comprising P1:E+P2:TI+P3:HTT+P4:FI+P5:SF at 23.4% and P1:E+P2:TI+P3:HTT+P4:FO+P5:SC at 7.8%), Stable A (P1:AL+P2:NS+P3:LTT+P4:FO+P5:SC, 21.9%), the pre-AI traditional market (P1:E+P2:NS+P3:LTT+P4:FO+P5:SC, 14.1%), and other unstable states (23.4%). The six successful combinations are concentrated in the low-HTT region ($\text{base}_3^{\text{HTT}} = 0.040$) with high SF payoff ($\text{base}_5^{\text{SF}} \geq 0.061$), or at selective mid-range values ($\text{base}_3^{\text{HTT}} = 0.074$ or 0.086 with $\text{base}_5^{\text{SF}} = 0.074$). This demonstrates that the mature market is accessible from a non-trivial but narrow subset of the two-parameter policy space, but only under a condition of *sustained parameter commitment*: the reform must be maintained for a sufficiently long time (50,000

steps) for the system to escape the transient trapping HTT trap and cross the basin boundary toward Stable B (Table 9).




Fig. 6 Two-dimensional policy space scan ($\text{base}_3^{\text{HTT}} \times \text{base}_5^{\text{SF}}$, 64 combinations, 50,000 steps). The white star marks the baseline ($\text{base}_3^{\text{HTT}} = 0.075$, $\text{base}_5^{\text{SF}} = 0.030$). Green cells indicate convergence to Stable B (mature market, $6/64 = 9.4\%$); orange cells indicate Stable A (traditional market, 21.9%); red cells indicate HTT-related traps (31.2%); gray cells indicate other unstable states. The narrow green region confirms that two-parameter reform reaches the mature market only from a small subset of the policy space (low HTT payoff + high SF payoff).

3.9 Transition Barrier: The Reform Window

A fundamental property of the replicator dynamics is that any pure-strategy corner is a fixed point for all parameter values (Proposition 1). At a corner where strategy i has frequency $x_i = 1$ and all others have $x_j = 0$, the dynamics for any inactive strategy satisfy

$$\dot{x}_j = x_j \cdot (\pi_j - \bar{\pi}) = 0 \cdot (\text{anything}) = 0. \quad (5)$$

To address this boundary condition, we define a *perturbed-corner experiment*: starting from the Stable A corner, we inject a small perturbation $\varepsilon = 0.01$ (each inactive strategy receives frequency 0.01), and then test the transition probability to Stable B under both short-run (5,000 steps) and long-run (50,000 steps) horizons, with 243 independent 5-parameter simultaneous shock combinations (3^5 grid: $\text{base}_1^{\text{AL}} \times \text{base}_2^{\text{TI}}$

Table 9 Two-parameter reform success combinations ($\text{base}_3^{HTT} \times \text{base}_5^{SF}$, 50,000 steps).

base_3^{HTT}	base_5^{SF}	Final State
0.040	0.061	AL+TI+HTN+FI+SF (Stable B)
0.040	0.074	AL+TI+HTN+FI+SF (Stable B)
0.040	0.087	AL+TI+HTN+FI+SF (Stable B)
0.040	0.100	AL+TI+HTN+FI+SF (Stable B)
0.074	0.074	AL+TI+HTN+FI+SF (Stable B)
0.086	0.074	AL+TI+HTN+FI+SF (Stable B)

$\times \text{base}_3^{HTN} \times \text{base}_4^{FI} \times \text{base}_5^{SF}$). We compare this against the baseline (mixed) initial condition to test whether the reform window depends on the system’s initial heterogeneity.

Results from Perturbed Stable A Corner ($\varepsilon = 0.01$):

1. **Short-run (5,000 steps): Complete trapping (0% escape).** At the short-run horizon, 0 out of 243 5-parameter shock combinations (0%) escape to Stable B (Fig. 7a). The system is fully captured by the dynamic lock-in: all 243 trajectories converge to Stable A (162/243, 66.7%) or a near-Stable A mixed state (AL+NS+LTT+FO+SF, 81/243, 33.3%). The 1% residual share of alternative strategies is insufficient to overcome the short-run trapping force of the traditional-market attractor. The dynamic lock-in is absolute at the short-run horizon: even 5-parameter simultaneous shocks cannot dislodge the system from the perturbed corner.
2. **Long-run (50,000 steps): High escape rate (81.5%).** When the horizon extends to 50,000 steps, 198 out of 243 shock combinations (81.5%) successfully escape to Stable B (mature market). The remaining 18.5% decay to Stable A (36/243, 14.8%) or are captured by intermediate HTT-related states (9/243, 3.7%). The dramatic increase from 0% to 81.5% confirms that the transition barrier is a *short-run phenomenon*: given sufficient time (50,000 steps), the system can cross the basin boundary under the same 5-parameter shocks that fail at the short-run horizon. The slow transient decay of the traditional-market attractor, combined with the sustained pressure of the 5-parameter reform, creates a wide reform window in the long run.

Results from Baseline (Mixed) Initial Condition:

1. **Short-run and long-run: Universal success (100%).** From the baseline mixed initial condition, all 243 5-parameter shock combinations (100%) converge to Stable B at both the short-run (5,000 steps) and long-run (50,000 steps) horizons. This confirms that the baseline IC retains sufficient heterogeneity to allow immediate reform success: the system is already in the mixed-state region where the 5-parameter shock vector can push it toward the mature market without needing to overcome a deep basin boundary.




Fig. 7 Transition barrier analysis (243 5-parameter shock combinations, two initial conditions, two horizons). (a) Short-run (5,000 steps): From the perturbed Stable A corner ($\varepsilon = 0.01$), 0% escape to Stable B (red bar), confirming that the dynamic lock-in is absolute at the short-run horizon. From the baseline mixed IC, 100% escape (green bar). (b) Long-run (50,000 steps): From the perturbed Stable A corner, 81.5% escape to Stable B (green bar), demonstrating that the barrier is a short-run phenomenon and the reform window opens dramatically at the long-run horizon. From the baseline mixed IC, 100% escape persists. The two-timescale structure reveals that reformers should not evaluate success at the short-run horizon: the same 5-parameter shocks that fail at 5,000 steps succeed at 81.5% at 50,000 steps.

3.10 Sequential versus Simultaneous Reform: The Cost of Delay

The transition-barrier analysis compares reform *shocks* but treats them as simultaneous. A natural policy question remains: if all three gates must be reformed, does the *sequencing* matter? We answer this with a deterministic sequencing experiment (50,000 steps, forward Euler, $\Delta t = 0.03$, baseline mixed initial condition, seed 42). Reform of each gate is implemented as a step change of its base-payoff parameter at a scheduled step: the trading-rule gate (P3) lowers $\text{base}_3^{\text{HTT}}$ from 0.075 to 0.040 (below the bifurcation threshold at which the HTT trap hardens; Section 3.6); the short-selling gate (P5) raises $\text{base}_5^{\text{SF}}$ from 0.030 to 0.070 (inside the Goldilocks transitional band $0.06 < \text{base}_5^{\text{SF}} < 0.08$; Table 7); the foreign-capital gate (P4) raises $\text{base}_4^{\text{FI}}$ from 0.040 to 0.150 (inside the Stable-B window $\text{base}_4^{\text{FI}} \geq 0.132$; Table 7). The *simultaneous* scheme applies all three shocks at step 0. A *sequential* scheme applies them piecewise in a fixed order with a constant inter-gate lag Δ (e.g., P3 $\rightarrow$ P5 $\rightarrow$ P4 at $\Delta = 10,000$ schedules the shocks at steps 0, 10,000, and 20,000). We test four lags ($\Delta = 2,500/5,000/10,000/20,000$) for the canonical order P3 $\rightarrow$ P5 $\rightarrow$ P4, all six order permutations at $\Delta = 10,000$, and three single-gate controls in which only one shock is ever applied. Table 10 reports the first step at which the trajectory's argmax state equals Stable B.

Three conclusions emerge. **First, complementarity: the trading-rule gate (P3) is the one gate that can never be reformed alone, and it is the slowest gate to lead with.** P3-only reform fails outright: the trajectory locks into the HTT trap (E+TI+HTT+FO+SC). By contrast, the short-selling (P5) and foreign-capital

Table 10 Sequential versus simultaneous three-gate reform (50,000 steps, deterministic, baseline mixed initial condition). “First at Stable B” is the first step at which the argmax strategy combination equals Stable B (AL+TI+HTN+FI+SF); the horizon is 50,000 steps, so entries above 20,000 reflect slow approach, not cyclic revisiting.

Scheme	Reaches Stable B	First at Stable B (step)	Final state (if not B)
Simultaneous (all at step 0)	Yes	2,665	—
Sequential P3→P5→P4, $\Delta = 2,500$	Yes	13,333	—
Sequential P3→P5→P4, $\Delta = 5,000$	Yes	23,223	—
Sequential P3→P5→P4, $\Delta = 10,000$	Yes	38,055	—
Sequential P3→P5→P4, $\Delta = 20,000$	No	—	E+TI+HTT+FI+SC (HTT region)
Permutation P4→P3→P5, $\Delta = 10,000$	Yes	14,659	—
Permutation P4→P5→P3, $\Delta = 10,000$	Yes	14,782	—
Permutation P5→P4→P3, $\Delta = 10,000$	Yes	22,580	—
Permutation P5→P3→P4, $\Delta = 10,000$	Yes	23,291	—
Permutation P3→P4→P5, $\Delta = 10,000$	Yes	33,872	—
Permutation P3→P5→P4, $\Delta = 10,000$	Yes	38,055	—
Single gate P3 only	No	—	E+TI+HTT+FO+SC (HTT trap)
Single gate P5 only	Yes	36,465	—
Single gate P4 only	Yes	20,786	—

(P4) gates do succeed even in isolation—36,465 and 20,786 steps respectively—because their reformed values sit inside the narrow single-parameter windows identified by the sensitivity analysis ($\text{base}_5^{SF} = 0.070$ in the Goldilocks band; $\text{base}_4^{FI} = 0.150 \geq 0.132$). We report these single-gate successes without cosmetic framing, but stress that they are *window-dependent*, not structural: they disappear outside the Goldilocks bands (Sections 3.5–3.6), and even when they succeed they are 8–14 times slower than simultaneous reform.

Second, timeliness: simultaneous reform is fastest, and every delay is costly. Simultaneous initiation reaches Stable B at step 2,665; the slowest successful sequential scheme (P3 first, $\Delta = 10,000$) takes 38,055 steps—about 14 times longer. Delaying the second and third gates monotonically inflates the transition time (13,333 → 23,223 → 38,055 steps as Δ grows from 2,500 to 10,000). Order matters as well: at $\Delta = 10,000$ all six permutations eventually succeed, but schemes that open the foreign-capital gate first are fastest ($\sim 14,700$ steps), while schemes that lead with the trading-rule gate are slowest ($\sim 38,000$ steps)—in the sequential schemes, the P3 shock alone first pushes the system toward the traditional-attractor region, from which the remaining gates must then lever it back.

Third, there is a critical delay beyond which sequential reform fails irreversibly. At $\Delta = 20,000$ the P3→P5→P4 scheme never reaches Stable B: by the time the later gates fire, the system has locked into the HTT region (E+TI+HTT+FI+SC), and the remaining shocks cannot dislodge it. The reform window is therefore real but bounded: sequential implementation is not ruled out, but procrastination beyond the critical inter-gate interval is. Together with the transition-barrier result (Section 3.9), this reframes the urgency of reform: the danger is not that the system “decays beyond rescue” (81.5% of shocks succeed even from the decayed corner at the long horizon),

but that *delay itself* multiplies the transition cost and, past a critical interval, destroys the opportunity altogether.

3.11 Parameter Sensitivity as an Over-Parameterization Diagnostic

A potential concern with any multi-population EGT is parameter scale: the five-population payoff matrix contains 146 parameters (141 payoff parameters plus 5 learning rates). Critics may argue that the two-attractor structure is a fragile calibration artifact—or, conversely, that a model whose conclusions survive arbitrary perturbation of so many parameters must be encoding those conclusions in only a handful of them. We address both concerns head-on by treating the perturbation battery not as a robustness selling point but as a *diagnostic of over-parameterization*: which parameters actually carry the model’s structure, and which are effectively inert? The answer is deliberately uncomfortable, and we report it without cosmetic framing.

Individual Parameter Perturbation. We test three structurally critical parameters— base_1^{AL} , base_3^{HTT} , and base_5^{SF} —at $\pm 50\%$ perturbations around their baseline values (six perturbations total). Five of the six perturbations preserve the two-attractor structure ($\text{n.stable} = 2$). The exception is $\text{base}_3^{HTT} + 50\%$ ($\text{base}_3^{HTT} = 0.1125$), which reduces the landscape to a single stable state (E+TI+HTT+FI+SC), consistent with the bifurcation threshold behavior. Two qualifications are worth noting. First, $\text{base}_1^{AL} - 50\%$ preserves the two-attractor count but shifts the identity of Stable A (P1 switches from AL to E), so “robustness” at this level means robustness of the attractor *count*, not necessarily of attractor *identity*. Second, the $\text{base}_5^{SF} - 50\%$ condition ($\text{base}_5^{SF} = 0.015$)—which completes the six-condition battery—preserves both stable attractors with unchanged identity ($\text{n.stable} = 2$; both stable corners identical to baseline), so the short-selling base payoff does not carry the two-attractor structure at the -50% margin.

Systematic $\pm 50\%$ Scan of All 141 Payoff Parameters. To move beyond a hand-picked set of “critical” parameters, we perturb each of the 141 payoff parameters individually by $\pm 50\%$ and re-evaluate the stability of all 96 pure-strategy corners via the joint Jacobian. A parameter is classified as structurally *active* if either shock alters the set of stable corners, and *inert* otherwise. The result: **129 of 141 parameters (91.5%) are inert**—both shocks preserve the exact pair of stable corners—and only **12 are active**. The composition of the active set is informative. Structurally active parameters are heavily concentrated in the base payoffs: 5 of the 13 base payoffs are active (base_1^{AL} , base_3^{LTT} , base_3^{LNT} , base_3^{HTT} , base_3^{HTN}), compared with only 7 of the 126 cross-effects. All 44 P5-related cross-effects are inert: adding the short-selling population’s interaction terms to the matrix does not, by itself, alter the corner architecture. This asymmetry supports a two-tier reading of the payoff matrix: the *base payoffs are the structure-determining terms*—they encode the institutional complementarity that produces the two attractors—while the *cross-effect magnitudes act primarily as transient modifiers*, shaping basin boundaries and approach paths rather than the existence or identity of the attractors.

An honest admission: the model is over-parameterized. The flip side of “91.5% of parameters are inert under $\pm 50\%$ shocks” is that the same qualitative

conclusions could, in principle, be obtained from a far smaller parameter set. The two-attractor architecture is encoded primarily by the base-payoff differences; most cross-effect magnitudes could be varied substantially—or, as the compression analysis below shows, zeroed entirely—without changing the attractor structure. We do not interpret this as redundant robustness to be advertised; we interpret it as evidence that the 141-parameter matrix is an *over-complete* representation of the underlying institutional structure, retained because each parameter carries an interpretable institutional meaning under the calibration protocol (Section 2.3.1), not because each is dynamically necessary.

Sign patterns are load-bearing even where magnitudes are not. Magnitude insensitivity does not imply that the cross-effects are decorative. As a complementary test, we run 200 seeded random draws in which each cross-effect’s *sign* is independently flipped with probability 1/2, and re-evaluate the corner structure. The two-attractor architecture survives in only 16.0% of draws (32/200), and the exact baseline pair of attractors is never recovered (0/200); the modal outcome is a single-attractor landscape (115/200). Thus the *joint sign pattern* of the cross-effects—which encodes the direction of institutional complementarities and substitutions—is load-bearing, even though the precise magnitudes within a $\pm 50\%$ band largely are not. This is consistent with the calibration rationale: what matters is the qualitative sign pattern and the relative magnitude ordering, not the precise values.

Toward a structurally constrained parameterization (future work). These diagnostics suggest a concrete model-reduction agenda: a structurally constrained parameterization that retains the 13 base payoffs, the small set of active cross-effects identified by the $\pm 50\%$ scan, and the sign constraints identified by the flip test—compressing 141 payoff parameters to on the order of 10–15 structural terms—while verifying that the basin structure and the two-timescale lock-in are preserved. We leave the construction and validation of this minimal model to future work.

Fine-Scale Perturbation ($\pm 10\%$ and $\pm 20\%$). To test whether the HTT trap is a fragile artifact of precise calibration, we conduct a finer-scale perturbation of six structurally critical parameters (base_1^{AL} , base_3^{HTT} , base_3^{HTN} , $\text{cross}_{13}^{AL,HTT}$, $\text{cross}_{35}^{HTT,SC}$, and $\text{cross}_{21}^{TI,AL}$) at $\pm 10\%$ and $\pm 20\%$ around their baseline values (24 perturbations total), evaluated at a shorter horizon (2,000 steps) to isolate the short-run trapping behavior. All 24 perturbations preserve the HTT trapping regime (P3 = HTT in every case). In 21 of the 24, the baseline initial condition converges to the baseline trap state AL+TI+HTT+FO+SC with only quantitative variation in the steady-state frequencies; in the remaining three ($\text{base}_1^{AL} - 20\%$, $\text{base}_3^{HTT} + 20\%$, and $\text{cross}_{13}^{AL,HTT} + 20\%$) the system settles in the neighboring slow trap E+TI+HTT+FO+SC, i.e., the perturbation shifts the P1 identity within the HTT trapping regime but does not free the system from the trap itself. The most sensitive parameter is base_3^{HTT} (P3 HTT varies by approximately ± 0.10); the least sensitive is $\text{cross}_{35}^{HTT,SC}$ (maximum frequency deviation ≈ 0.01). This confirms that the HTT trap is not a fragile product of any single parameter but a structural feature of the joint cross-population architecture. Incremental reform of individual parameters, even at the 20% level, cannot escape the trap.

Systematic Cross-Effect Removal. When all non-P5 cross-effects are simultaneously zeroed, the system retains both stable attractors, but Stable B shifts from AL+TI+HTN+FI+SF to AL+TI+HTN+FO+SF (P4 changes from FI to FO). This confirms that non-P5 cross-effects are structurally selective: they determine which specific institutional configuration emerges within the mature-market stack, not whether a mature-market stack exists.

The Minimal Sufficient Model. We construct a 61-parameter model by retaining only the 13 base payoffs, the 44 P5-related cross-effects, and a small set of structural anchor cross-effects—systematically zeroing the remaining non-P5 cross-effects (the large majority of the 82 non-P5 cross-effects; the 141 payoff parameters comprise the 13 base payoffs, 126 cross-effects, and 2 homogeneity-penalty parameters κ). This minimal model retains both stable attractors with the correct identity (Stable B remains AL+TI+HTN+FI+SF). Read together with the $\pm 50\%$ scan, this compression result reinforces the over-parameterization diagnosis: the two-attractor structure survives the systematic zeroing of most cross-effects because those parameters were never carrying the structure in the first place; the institutional conclusions are carried by the base-payoff architecture and the joint sign pattern of the remaining cross-effects. Figure 8 schematizes this two-attractor landscape and its key features.

3.12 Robustness of Model–Trend Correlations under Parameter Perturbation

A natural question that complements the structural robustness analysis is whether the correlations between the model’s simulated trajectories and stylized market trends are artifacts of the baseline calibration. The structural robustness analysis (Section 3.11) confirms that the two-attractor architecture survives $\pm 50\%$ perturbation of individual parameters, but this addresses the *internal* structure. We now ask whether the model’s *external* qualitative correspondence with stylized market trends is equally robust. We note at the outset that this analysis is entirely simulation-based: it uses *simplified stylized proxy series* (e.g., sigmoid curves), not real data. In particular, the negative stylized ρ_2 (P2.TI vs. quant AUM) reported below arises because the stylized proxy is a monotonically increasing curve while the model’s long-run prediction converges to the traditional (NS) regime; it is a property of the stylized series, which cannot capture detrended fluctuations around the trend.

Method. For each of three structurally critical parameters— base_1^{AL} , $\text{cross}_{13}^{AL,HTT}$, and base_3^{HTT} —we run the simulation at $\pm 25\%$ and $\pm 50\%$ perturbations (50,000 steps, baseline initial condition) and compute the trajectory-level Pearson correlation between the simulated strategy frequencies and three stylized proxy series. The proxy series are calibrated to documented A-share market trends: (i) margin-trading balance (linear growth with a 2015-crash dip; CSRC/SSE/SZSE, 2015–2024), (ii) quant-fund AUM (sigmoid adoption curve; AMAC, 2018–2023), and (iii) market turnover (cyclical with mild upward trend; Wind, 2015–2024). Because these are stylized trend series rather than detrended real data, the absolute ρ values should not be interpreted as econometric estimates. The quantity of interest is the *sensitivity* of ρ to parameter perturbation: if the correlations maintain directional consistency under perturbation, the model–trend correspondence is not parameter-specific.




Fig. 8 Schematic of the institutional landscape under two-timescale dynamic lock-in. **Short-run (5,000 steps):** The HTT trap acts as a transient trapping attractor that captures the system from most initial conditions (HTT-related states total 46.5% of the 200 Dirichlet(1) random initial conditions, seed 42), creating the dynamic lock-in observed by market participants. **Long-run (50,000 steps):** The HTT trap slowly decays toward Stable A (traditional market, 36.5% basin—the larger of the two) or, in a minority of trajectories, escapes to Stable B (mature market, 18.0% basin). Although Stable A’s basin is larger, its stochastic residence time is finite: under frequency noise it escapes at a 93.3% rate already at $\sigma = 0.005$ (median first escape 5,901 steps) and at 100% for $\sigma \geq 0.01$ (Section 3.7). Stable B, reachable through complementary three-gate reform or single parameters inside narrow Goldilocks windows, is the genuinely persistent attractor under low noise, with 60% of trajectories remaining unescaped for the full 50,000 steps at $\sigma = 0.005$. The landscape explains why the system is overwhelmingly attracted to the traditional trap in the short run: moderate volatility pushes the system toward Stable A, and the urgency of reform stems from the cost of delay rather than from an evaluation-horizon deadline—simultaneous reform reaches Stable B fastest, while delay beyond a critical inter-gate interval fails irreversibly (Section 3.10).

Results at $\pm 25\%$. Table 11 reports the correlations at $\pm 25\%$ perturbation. In all nine configurations, the model converges to the same attractor (Stable A: AL+NS+LTT+FO+SC), and the three correlations maintain directional consistency. The stylized ρ_1 (P1_AL vs. margin) remains positive throughout ($0.22 \leq \rho_1 \leq 0.94$), confirming that the model’s prediction of rising AI-driven retail participation is robust to moderate parameter variation. The stylized ρ_2 (P2_TI vs. quant) is negative throughout ($-0.94 \leq \rho_2 \leq -0.62$). This negative correlation is a limitation of the stylized proxy rather than a model failure: the proxy is a monotonically increasing sigmoid curve, while the model’s long-run prediction converges to P2 = NS, so a monotonically increasing AUM trend cannot be matched by construction. The stylized ρ_3 (P3_HTT vs. turnover) is weakly negative throughout ($-0.58 \leq \rho_3 \leq -0.07$), consistent with the same trapping logic.

Results at $\pm 50\%$. At the more extreme $\pm 50\%$ perturbation, several configurations cross basin boundaries—consistent with the structural robustness analysis (Section 3.11). For example, $\text{base}_1^{\text{AL}} = 0.011$ (-50%) causes the system to converge to P1:E+P2:NS+P3:LTT rather than the baseline P1:AL attractor, reversing the sign

Table 11 Sensitivity of stylized proxy correlations to $\pm 25\%$ parameter perturbation (50,000 steps). All nine configurations converge to Stable A (AL+NS+LTT+FO+SC). Directional consistency is maintained for all three correlations. Significance: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$.

Parameter	Perturbation	ρ_1 (margin)	ρ_2 (quant)	ρ_3 (turnover)
$\text{base}_1^{AL} = 0.017$	-25%	0.88***	-0.62***	-0.07
$\text{base}_1^{AL} = 0.022$	baseline	0.91***	-0.74***	-0.12**
$\text{base}_1^{AL} = 0.028$	+25%	0.72***	-0.91***	-0.50***
$\text{cross}_{13}^{AL,HTT} = -0.044$	-25%	0.94***	-0.62***	-0.07
$\text{cross}_{13}^{AL,HTT} = -0.035$	baseline	0.91***	-0.74***	-0.12**
$\text{cross}_{13}^{AL,HTT} = -0.026$	+25%	0.22***	-0.94***	-0.57***
$\text{base}_3^{HTT} = 0.056$	-25%	0.77***	-0.94***	-0.58***
$\text{base}_3^{HTT} = 0.075$	baseline	0.91***	-0.74***	-0.12**
$\text{base}_3^{HTT} = 0.094$	+25%	0.55***	-0.74***	-0.15***

of ρ_1 . Similarly, $\text{cross}_{13}^{AL,HTT} = -0.018$ (+49%) pushes the system across the basin boundary to P2:TI+P3:HTT dominance, flipping ρ_2 from negative to positive. These sign reversals are not evidence of calibration fragility but rather reflect the basin-crossing behavior already documented in the bifurcation analysis (Section 3.5) and the $\pm 50\%$ perturbation test (Section 3.11).

Interpretation. The sensitivity analysis reveals a two-tier robustness structure:

- **Moderate perturbation ($\pm 25\%$):** The model converges to the same attractor, and all three stylized proxy correlations maintain directional consistency. The model–trend correspondence is not an artifact of specific parameter values.
- **Extreme perturbation ($\pm 50\%$):** Basin boundaries are crossed, altering both the equilibrium and the proxy correlations. This is consistent with the structural robustness analysis: the two-attractor architecture is preserved, but individual trajectories may shift between basins.

The complementary relationship between this analysis and Section 3.11 is worth emphasizing. Section 3.11 demonstrates *internal* robustness: the attractor structure survives parameter perturbation. The present analysis demonstrates *external* robustness: the model’s qualitative correspondence with stylized market trends also survives moderate perturbation. Together, they confirm that the model’s conclusions are structural rather than calibration-dependent.

Parameters as regime-dependent probes. The sign reversals at $\pm 50\%$ carry a deeper implication. Rather than viewing base_1^{AL} , $\text{cross}_{13}^{AL,HTT}$, and base_3^{HTT} as arbitrary calibration choices, the sensitivity analysis reveals that they function as *coordinates in a policy space* whose thresholds correspond to observable shifts in market behavior. Within the $\pm 25\%$ “stable zone,” the model’s qualitative predictions are externally anchored: proxy correlations remain directionally consistent, and the model–trend correspondence is structural. Beyond the $\pm 50\%$ “flip zone,” the system crosses a basin boundary into an alternative attractor regime, and the proxy correlations change qualitatively—providing an externally detectable signal of institutional regime transition. This reframes the role of parameter values: their precise location is not merely a calibration detail but carries independent empirical content, since the threshold at which proxy correlations flip is, in principle, observable in market data.

The practical implication is a *monitoring framework* (not a forecasting tool): sustained shifts in the correlation between AI-driven retail participation and margin-trading balance, or between institutional modernization and quant-fund growth, could signal that the system is approaching a basin boundary—and that the reform window is narrowing. This is consistent with the paper’s qualitative positioning: the model does not predict when transitions will occur, but it identifies which observable indicators are informative about the proximity of regime change.




Fig. 9 The AI Penetration Paradox: Two-timescale analysis (5,000 vs 50,000 steps). From the baseline initial condition at the short-run horizon (5,000 steps): as AI penetration increases from 0 to 3, the system undergoes a gradual but non-monotonic transition from the Pre-AI traditional market (E+NS+LTT+FO+SC) to the HTT trap (AL+TI+HTT+FO+SC), with P1 AL rising from 0.023 to 0.985 and P3 HTT from 0.051 to 0.999. The HTT trap appears to deepen monotonically with AI penetration. From the same baseline initial condition at the long-run horizon (50,000 steps): the relationship is non-monotonic. Low AI penetration (0–0.7) converges to the Pre-AI market; intermediate penetration (0.8–1.5) converges to Stable A (AL+NS+LTT+FO+SC), the traditional AI-hybrid market; only high penetration ($\lambda=1.8$) converges to the HTT trap. Two critical thresholds emerge: 0.8 (Pre-AI to Stable A) and 1.8 (Stable A to HTT). This confirms that the AI penetration paradox is a time-scale-dependent phenomenon. From a near-mature initial condition at 5,000 steps: the system remains in Stable B (AL+TI+HTN+FI+SF) across all penetration levels, with P3 HTN above 0.989. From the same near-mature initial condition at 50,000 steps: Stable B persists for low and moderate penetration, but extreme AI penetration ($\lambda \geq 2.9$) triggers a two-stage collapse: at $\lambda = 2.9$, the system transitions to a mixed state (P3 HTT = 0.248, P4 flips from FI to FO); at $\lambda = 3.0$, the mature market fully collapses to the HTT trap (P3 HTT = 1.0, P3 HTN = 0). This reveals that excessive AI penetration can overwhelm even the mature market.

3.13 The AI Penetration Paradox: Time-Scale-Dependent Non-Monotonicity

The AI penetration paradox is not merely a stability paradox but a *time-scale-dependent non-monotonicity* that reveals the complex interplay between AI adoption and institutional reform. We parameterize AI penetration intensity by $\lambda \in [0, 3]$, which scales the *payoff advantage of the AI-driven retail strategy AL relative to passive holding*: $\text{base}_1^{\text{AL}}$ and every cross-effect involving AL (both AL's own cross-effects toward other populations and other populations' cross-effects toward AL) are interpolated between an E-equivalent reference and their baseline values, so that at $\lambda = 0$ the AL strategy's payoff structure degenerates to that of E (the zero-AI benchmark), $\lambda = 1$ recovers the observed baseline calibration (2020–2024 in the Chinese A-share market), and $\lambda > 1$ amplifies the AL advantage beyond its baseline level, corresponding to hypothetical future intensification. In economic terms, λ captures the effective market share of AI-driven trading strategies relative to the baseline calibration period.

Result 1: The baseline initial condition exhibits a three-regime structure at the long-run horizon. At the short-run horizon (5,000 steps), increasing λ drives a smooth monotonic transition from the Pre-AI traditional market (E+NS+LTT+FO+SC) to the HTT trap (AL+TI+HTT+FO+SC). P1 AL rises from 0.023 at $\lambda = 0$ to 0.985 at $\lambda = 3$, while P3 HTT rises from 0.051 to 0.999. This is the dynamic lock-in that market participants observe (Fig. 9): AI adoption appears to deepen the HTT trap without bound.

At the long-run horizon (50,000 steps), however, the relationship is *non-monotonic*:

- **Regime I (low penetration, $\lambda = 0\text{--}0.7$):** The system converges to the Pre-AI traditional market (E+NS+LTT+FO+SC). AI penetration is insufficient to overcome the institutional inertia of the traditional stack.
- **Regime II (intermediate penetration, $\lambda = 0.8\text{--}1.5$):** The system converges to Stable A (AL+NS+LTT+FO+SC), the traditional AI-hybrid market. This is the most striking finding: moderate AI penetration, when given sufficient time to unfold, pushes the system toward the *traditional* stable attractor rather than the HTT trap. The mechanism is that intermediate AI penetration strengthens the AL strategy enough to dislodge the Pre-AI market but not enough to trigger the quant-institutional feedback loop that drives the HTT trap.
- **Regime III (high penetration, $\lambda \geq 1.8$):** The system converges to the HTT trap (AL+TI+HTT+FO+SC). High AI penetration activates the reinforcing cross-effect loop (TI-HTT-SC) that locks the system into the HTT regime. The values at $\lambda = 1.6\text{--}1.7$ reflect a transient return to the Pre-AI market (Regime I), indicating that the Regime II window (Stable A) is not continuous and narrows with increasing λ .

Two critical thresholds emerge: $\lambda_1 \approx 0.8$ (Pre-AI $\rightarrow$ Stable A) and $\lambda_2 \approx 1.8$ (Stable A $\rightarrow$ HTT). The existence of an intermediate Stable A window confirms that the AI penetration paradox is *not* a simple monotonic deepening of the trap but a time-scale-dependent phenomenon where moderate AI penetration can, paradoxically, steer the system toward the traditional stable attractor in the long run.

Mechanism of the transient return at $\lambda = 1.6$ – 1.7 . The reversion of the 50,000-step label to the Pre-AI configuration in this narrow band is not evidence of a third attractor: the Jacobian at these terminal states retains a positive dominant eigenvalue, so the system is still in transit when the horizon closes—the label is a finite-horizon snapshot, not a steady-state classification. The mechanism is a slow passage through the neighborhood of the Pre-AI corner. The λ -scaling modifies the relative fitness of the AL strategy (its payoff advantage over passive holding); in this band its net effect is that, once the trajectory escapes the short-run HTT transient, the AL share no longer accumulates fast enough for the system to complete the approach to Stable A within 50,000 steps. The trajectory instead drifts through the region where the traditional institutional stack (NS+LTT+FO+SC) is already dominant but the retail population is still classified as passive (E)—a region in which motion is slow precisely because the AI-adoption margin is nearly frozen. At $\lambda \geq 1.8$ the strengthened TI–HTT reinforcing loop captures the trajectory before this drift can resolve, producing the HTT label of Regime III. The apparent discontinuity at $\lambda = 1.6$ – 1.7 is therefore the signature of a continuous trajectory family passing through a slow region of the state space as the basins of Stable A and the HTT trap shift with λ , and it is fully consistent with the intermediate Stable A window being genuine on either side of the band.

Result 2: The near-mature initial condition reveals a “mature-market collapse threshold.” At the short-run horizon (5,000 steps), the system remains in Stable B (AL+TI+HTN+FI+SF) across all tested λ , with P3 HTN > 0.989 . This confirms that AI penetration does not destabilize a system that is already close to the mature market—the paradox is conditional on the initial condition.

At the long-run horizon (50,000 steps), however, a new phenomenon emerges:

- For $\lambda = 0.2$ – 2.8 , Stable B persists (P3 HTN ≈ 0.999 – 0.998).
- At $\lambda = 2.9$, the system transitions to a mixed state (AL+TI+HTN+FO+SF) with P3 HTN = 0.752 and P3 HTT = 0.248. The P4 foreign capital strategy has flipped from FI to FO, indicating that the mature-market stack has been partially dislodged. The Jacobian eigenvalue is positive (max.eig $\approx +0.012$), confirming that the state is a saddle.
- At $\lambda = 3.0$, the system *fully collapses* to the HTT trap (AL+TI+HTT+FO+SF) with P3 HTT = 1.0 and P3 HTN = 0. The mature market has been completely overwhelmed by the quant-institutional feedback loop.

This reveals a critical insight: *excessive AI penetration can collapse even the mature market*. There exists a “mature-market collapse threshold” at $\lambda_{crit} \approx 2.9$ – 3.0 , beyond which the reinforcing cross-effect loop (TI-HTT-SC) becomes so strong that it overwhelms the mature-market stack from the inside. The collapse is not instantaneous but follows a two-stage trajectory: at $\lambda = 2.9$, the system loses the FI component and enters a mixed state; at $\lambda = 3.0$, the HTT component dominates entirely. The policy implication is alarming: even a market that has successfully reached the mature-market configuration (Stable B) is not immune to AI-driven institutional backsliding if AI penetration is allowed to proceed without complementary institutional reform (Table 12).

The overarching message is that AI adoption is not a neutral technology that can be layered onto any institutional framework. Its effects are *time-scale-dependent, non-monotonic, and conditional on the initial institutional configuration*. Policymakers cannot assume that allowing AI adoption to proceed will “naturally” prepare the market for modernization; instead, AI adoption without complementary reform may either lock the system in the HTT trap (short run), steer it toward the traditional market (long run, moderate penetration), or collapse the mature market (long run, extreme penetration).

3.13.1 Critical Threshold and Institutional Irreversibility

The preceding analysis establishes that the HTT trap exhibits long transient trapping on the long timescale for moderate AI penetration ($\lambda < 2.3$). But how long does the transient state survive, and is there a point of no return? To answer this, we compute the Jacobian of the replicator dynamics at the HTT quasi-steady-state and extract the dominant eigenvalue λ_{dom} . The characteristic escape time is defined as $\tau(\lambda) = 1/\lambda_+$ where $\lambda_+ = \max(0, \lambda_{\text{dom}})$.

Metastability and finite lifetime. For $\lambda \in [1.8, 2.3]$, we find $\lambda_{\text{dom}} > 0$, indicating that the HTT state is locally unstable and the system can spontaneously escape (Fig. 10A). The escape time is finite and non-monotonic in λ , with a local minimum at $\lambda \approx 2.0$ ($\tau \approx 67$), suggesting an “optimal” AI penetration rate for rapid trapping. At $\lambda = 2.2$, the escape time has already risen to $\tau \approx 163$, indicating that the trap is hardening as the critical threshold is approached.

Critical bifurcation. Within the band $\lambda^* \in [2.2, 2.4]$, the dominant eigenvalue crosses zero (Fig. 10B), marking a bifurcation from transient trapping to permanently stable. For λ above this band, the HTT trap becomes a genuine stable attractor from which the system cannot escape without external intervention. The escape time diverges as the band is approached, indicating strong critical slowing down analogous to physical phase transitions. The band width is not arbitrary: the zero crossing is located between $\lambda = 2.2$ and $\lambda = 2.3$ under the 50,000-step horizon with the Jacobian evaluated at the HTT quasi-steady-state (with the λ -scaling convention applied only to P1-AL-related payoffs), while the 5,000-step sequence crosses at $\lambda \approx 2.34$. We also report a narrow re-crossing interval around $\lambda \in [2.3, 2.5]$, where the dominant eigenvalue briefly returns to near-zero or marginally positive values before turning firmly negative; we therefore report λ^* as a band rather than a single point value.

Initial-condition dependence. The critical threshold λ^* is not universal: it depends on the initial condition’s distance to the basin boundary. For initial conditions closer to the HTT state (e.g., mature-market configurations), the effective λ^* is lower; for conditions closer to Stable A, the system may bypass the HTT trap entirely even at $\lambda > \lambda^*$. What is universal is not the exact value of λ^* but the existence of a critical regime where the HTT trap transitions from transient trapping to permanently stable.

Two thresholds, not one. It is essential to distinguish between two distinct critical phenomena. The first, $\lambda^* \in [2.2, 2.4]$, is a *local stability threshold*: the HTT state itself transitions from transient trapping to permanently stable for the *same* initial condition (baseline IC). The second, $\lambda_c \approx 2.9$, marks a *mature-market collapse threshold*: the Stable B fixed point loses local stability and the basin of attraction

shifts toward the HTT trap. For $\lambda = 2.9$, the system converges to a mixed state (AL+TI+HTN+FO+SF) with P3 HTT = 0.248 and P4 FI = 0.39, indicating that the mature-market stack has been partially dislodged. For $\lambda > 2.9$, the HTT basin dominates entirely and the mature market is structurally inaccessible. The coexistence of both thresholds in the interval $\lambda \in [2.3, 2.9]$ implies that the system enters a *multi-attractor regime* with at least two attractors (HTT and Stable B) simultaneously present, though Stable B is increasingly fragile. For $\lambda < 2.3$, the HTT trap is a transient trapping state and reform is possible; for $\lambda \in [2.3, 2.9]$, the HTT trap is permanent while Stable B’s corner Jacobian no longer has all-negative eigenvalues—Stable B loses local stability (Table 12)—so that the system exhibits the residual, transient coexistence of the two-attractor configuration for a narrow set of initial conditions; for $\lambda > 2.9$, the mature market collapses and the HTT basin dominates. The permanence statement is horizon- and criterion-dependent: at the 50,000-step horizon a narrow re-crossing band $\lambda \in [2.3, 2.5]$ exists, in which the dominant eigenvalue briefly returns to near-zero or marginally positive values (Fig. 10), so the boundary between transient trapping and permanence should be read as a band rather than a point. This three-stage structure is not a contradiction but a hierarchy of institutional lock-in.

Policy implication: irreversibility. The existence of two thresholds has profound policy implications. When $\lambda < \lambda^*$ (i.e., below the band $[2.2, 2.4]$), the HTT trap is a transient trapping state and policy intervention has a finite window of opportunity: the system can still escape if reform is sustained long enough. Once $\lambda > \lambda^*$, the trap becomes locally irreversible—no amount of time or incremental reform can dislodge the HTT state itself. However, for $\lambda \in [2.3, 2.9]$, the mature-market configuration persists as a residual attractor for a narrow set of initial conditions—Stable B’s corner Jacobian is no longer all-negative, so it is not a stable attractor in this band, but trajectories from a small set of mature initial conditions still converge to its neighborhood—meaning that a *sufficiently large and simultaneous* reform could still theoretically reach the mature market, albeit with vanishing probability and only in the absence of noise. Once $\lambda > \lambda_c \approx 2.9$, even this distant hope disappears: the mature market collapses to the HTT trap and the basin is dominated by HTT for all initial conditions. This implies a two-stage policy warning: first at $\lambda^* \in [2.2, 2.4]$ (HTT hardens), then at $\lambda_c \approx 2.9$ (the mature market collapses). Policymakers must act *before* the first threshold, not after the second.

3.14 Parameter Importance and Institutional Rigidity

The preceding analysis identifies the AI penetration paradox and the complementarity trap at the phenomenological level. But which parameters are truly *load-bearing* (structural), and which are merely *modulating* (affecting stability without altering the attractor)? To answer this, we perform a parameter importance scan: for each of 30 key parameters, we set the parameter to zero (complete removal) and observe whether the steady-state attractor changes. Table 13 reports the results for two representative regimes: $\lambda = 1.0$ (Stable A regime) and $\lambda = 2.0$ (HTT transient trapping regime).

Structural parameters. At $\lambda = 1.0$, 13 parameters are load-bearing: the baseline payoffs of P1 (base_1^E , base_1^{AL}), P2 (base_2^{NS}), P3 (base_3^{LTT}), P4 (base_4^{FO}), and P5 (base_5^{SC}); the cross-effects between AL adoption and HTT institutions




Fig. 10 Characteristic escape time and bifurcation of the HTT trap. **(A)** Escape time $\tau(\lambda)$ as a function of AI penetration rate for short-run (5,000 steps, red) and long-run (50,000 steps, blue) horizons. The green shaded band marks the critical threshold band $\lambda^* \in [2.2, 2.4]$; the red dotted line marks the 5,000-step sequence crossing point ($\lambda \approx 2.34$), illustrating the horizon dependence of the estimated crossing. Triangles indicate stable regimes where $\tau = \infty$. **(B)** Dominant eigenvalue of the Jacobian computed at the HTT quasi-steady-state under the 50,000-step horizon (with the λ -scaling convention applied only to P1-AL-related payoffs). The zero crossing within $\lambda^* \in [2.2, 2.4]$ indicates a bifurcation from transient trapping to permanently stable. The orange shaded region marks the transient trapping regime; the gray hatched band ($\lambda \in [2.3, 2.5]$) marks a narrow re-crossing interval in which the eigenvalue briefly returns to near-zero or marginally positive values; the purple shaded region marks the permanently stable regime. The threshold is therefore reported as a band rather than a single point value.

($\text{cross}_{13}^{AL,HTT}$) and between AL and LTT ($\text{cross}_{13}^{AL,LTT}$); the homogeneity penalty ($\text{homog_penalty}_{HT}$); and the learning rates of the first four populations (η_1 – η_4). At $\lambda = 2.0$, the list expands to 15 load-bearing parameters, with the addition of seven parameters: base_2^{TI} , base_3^{HTT} , $\text{cross}_{12}^{AL,TI}$, $\text{cross}_{13}^{E,HTT}$, $\text{cross}_{23}^{TI,HTT}$, $\text{cross}_{23}^{NS,HTT}$, and η_5 , while five parameters (base_2^{NS} , base_3^{LTT} , $\text{cross}_{13}^{AL,LTT}$, η_1 , η_2) become neutral or modulating. This shift confirms that the institutional structure becomes *more rigid* as AI penetration increases: more parameters are locked into the HTT configuration, and fewer degrees of freedom remain for policy intervention.

Safety valves. Two parameters are *modulating* (affecting stability without changing the attractor) at $\lambda = 2.0$: η_1 (retail learning rate) and η_2 (institutional learning rate). Zeroing these parameters drives the HTT state’s maximum eigenvalue from +0.0149 (transient trapping) to 0.0000 (critically stable). This implies that the *speed of learning* is what maintains the HTT trap in its transient trapping regime rather than a permanently stable one. The 0.0000 reading in the frozen-learning limit should not be over-interpreted: as detailed in the note below, it is an exact zero eigenvalue generated mechanically by the frozen population equations—a dimension collapse of the reduced Jacobian—not evidence that freezing learning hardens the trap. Merely slowing learning ($\eta > 0$ small but positive) prolongs but does not prevent eventual escape. Conversely, *accelerating* institutional learning (higher η_2) could theoretically destabilize the HTT trap by increasing the rate of spontaneous escape, though this also increases the volatility of the transition.

Neutral parameters. Roughly half of the scanned parameters are neutral in both regimes (e.g., base_1^{AG} , base_3^{HTN} , $\text{cross}_{24}^{TI,FI}$). This is not a failure of the model but a

Table 12 Key AI penetration transition points at the long-run horizon (50,000 steps). *Note on $\lambda = 0$ row:* $P3_{\text{HTT}}$ values in this table are final converged frequencies at $t = 50,000$, not initial conditions. The baseline initial condition has $P3_{\text{HTT}} = 0.50$ at $t = 0$ for all λ ; the zero values at $\lambda = 0$ –1.8 indicate that the system converges to a regime where HTT frequency decays to near zero in the long run. *Note on the $\lambda = 0$ Mature IC entry:* both stable attractors identified in Section 3.2 have $P1 = \text{AL}$; at $\lambda = 0$, however, AL carries no AI advantage over algorithmic trading (AG), so the Mature IC converges to a Stable-B-like state with $P1 = \text{AG}$ replacing $P1 = \text{AL}$.

λ	Baseline IC State	Baseline IC $P3_{\text{HTT}}$	Mature IC State	Mature IC $P3_{\text{HTT}}$
0.0	E+NS+LTT+FO+SC (Pre-AI)	0.000	AG+TI+HTN+FI+SF (Stable-B-like, $P1 = \text{AG}$ variant)	0.000
0.8	AL+NS+LTT+FO+SC (Stable A)	0.000	AL+TI+HTN+FI+SF (Stable B)	0.000
1.5	AL+NS+LTT+FO+SC (Stable A)	0.000	AL+TI+HTN+FI+SF (Stable B)	0.000
1.6	E+NS+LTT+FO+SC (Pre-AI)	0.000	AL+TI+HTN+FI+SF (Stable B)	0.000
1.8	E+TI+HTT+FO+SC (HTT)	0.975	AL+TI+HTN+FI+SF (Stable B)	0.000
2.8	AL+TI+HTT+FO+SC (HTT)	1.000	AL+TI+HTN+FI+SF (Stable B, unstable)	0.002
2.9	AL+TI+HTT+FO+SC (HTT)	1.000	AL+TI+HTN+FO+SF (Mixed)	0.248
3.0	AL+TI+HTT+FO+SC (HTT)	1.000	AL+TI+HTT+FO+SF (Collapse)	1.000

reflection of the hierarchical structure of institutional interaction: some parameters are deeply embedded in the feedback loops that define the attractor, while others operate at the margins, affecting the *speed* of convergence but not the *destination*.

Implications for policy design. The parameter importance scan reveals that the policy challenge is not to adjust *any* parameter but to identify the *structural* ones. The list of 13–15 load-bearing parameters provides a concrete checklist for institutional reform: any reform that does not simultaneously address the base payoffs of all five populations, the AI–HTT cross-effect, and the homogeneity penalty will fail to dislodge the system from the HTT trap. The two modulating parameters (η_1, η_2) offer a secondary lever: accelerating institutional learning may shorten the effective lifetime of the transient trapping state, creating a larger window for complementary reform. However, this lever is dangerous: too much learning acceleration may increase the volatility of the transition, pushing the system toward an unstable mixed state rather than the mature market.

Note on Learning-Rate Classification. Standard replicator dynamics treats $\eta_i > 0$ as a time-rescaling parameter that affects convergence speed but not fixed-point locations. Our parameter-importance scan, however, tests the singular limit $\eta_i = 0$, which freezes population i entirely—removing it from co-evolution rather than merely

slowing it. The label “load-bearing” therefore means: if this population is frozen at its initial condition, does the reduced system converge to a different attractor? This measures structural co-evolutionary necessity, not the continuous effect of learning speed. The shift of η_1 , η_2 from load-bearing ($\lambda = 1.0$) to modulating ($\lambda = 2.0$) reflects regime-dependent co-evolutionary structure: at low AI penetration, retail and institutional evolution are necessary to reach Stable A; at high AI penetration, the HTT trap becomes self-sustaining and the co-evolution of these populations is no longer necessary for the trap to persist—their role shifts from structural to modulatory. Consequently, the 0.0000 maximum eigenvalue reported for the frozen-learning scan is an exact zero eigenvalue produced by the frozen population equations—a mechanical artifact of dimension collapse in the reduced system, not a marginal-stability result carrying economic content.

4 Discussion

4.1 Why Single-Point Reform Fails: The Complementarity Trap

The central theoretical insight of this paper is that institutional constraints are complements, not substitutes. Removing one constraint while leaving others intact creates a “complementarity trap”: the system reconfigures around the remaining constraints rather than evolving toward the unconstrained state. This failure is not absolute—narrow parameter windows exist in which a single gate suffices (the short-selling and foreign-capital gates in the sequencing experiment of Section 3.10, whose reformed values sit inside the Goldilocks bands of Section 3.5)—but these windows are exceptional and gate-specific: the trading-rule gate (P3) never succeeds alone from the baseline initial condition.

Example: T+1 Removal Alone. When T+1 is removed (P3 shifts from HTT toward HTN), the system initially gains some HTN share at the short-run horizon (5,000 steps). But on the long timescale (50,000 steps), the system converges to Stable A (AL+NS+LTT+FO+SC), not the mature market. The removal of T+1 creates a “window of opportunity” but does not provide the “escape velocity” needed to reach Stable B; instead, it redirects the system toward the traditional stable attractor. The long-run convergence to Stable A confirms that T+1 removal alone is insufficient for institutional modernization.

Example: Short-Selling Liberalization Alone. When short-selling is liberalized (P5 shifts from SC toward SF), the system settles in an unstable mixed state (AL+TI+HTT+FO+SF) with P1 AL = 0.866, P2 TI = 1.000, P3 HTT = 1.000, P4 FO = 1.000, and P5 SF = 1.000. This is a “quant-friendly but institutionally immature” unstable state where institutions can hedge but retail investors are still locked in T+1 + price limits. Short-selling liberalization alone does not push the system toward the mature market in the long run.

Example: Full Foreign Capital Open Alone. When foreign capital is fully liberalized (P4 shifts to FI), the system converges to Stable B (AL+TI+HTN+FI+SF) in the long run (50,000 steps). This is the only single-point counterfactual that reaches the mature market. However, this convergence requires the initial condition to be fully

Table 13 Parameter importance classification at two representative AI penetration levels. A parameter is classified as *load-bearing* if zeroing it changes the steady-state attractor, *modulating* if it affects stability without changing the attractor, and *neutral* if it has no significant effect.

Parameter	Description	$\lambda = 1.0$	$\lambda = 2.0$
<i>Load-bearing (structural) parameters</i>			
base_1^E	P1 E baseline payoff	✓	✓
base_1^{AL}	P1 AL baseline payoff	✓	✓
base_2^{NS}	P2 NS baseline payoff	✓	—
base_2^{TI}	P2 TI baseline payoff	—	✓
base_3^{LTT}	P3 LTT baseline payoff	✓	—
base_3^{HTT}	P3 HTT baseline payoff	—	✓
base_4^{FO}	P4 FO baseline payoff	✓	✓
base_5^{SC}	P5 SC baseline payoff	✓	✓
$\text{cross}_{12}^{AL,TI}$	P1 AL – P2 TI cross-effect	—	✓
$\text{cross}_{13}^{AL,HTT}$	P1 AL – P3 HTT cross-effect	✓	✓
$\text{cross}_{13}^{E,HTT}$	P1 E – P3 HTT cross-effect	—	✓
$\text{cross}_{13}^{AL,LTT}$	P1 AL – P3 LTT cross-effect	✓	—
$\text{cross}_{23}^{TI,HTT}$	P2 TI – P3 HTT cross-effect	—	✓
$\text{cross}_{23}^{NS,HTT}$	P2 NS – P3 HTT cross-effect	—	✓
$\text{homog_penalty}_{HT}$	HTT homogeneity penalty	✓	✓
η_1	P1 learning rate	✓	—
η_2	P2 learning rate	✓	—
η_3	P3 learning rate	✓	✓
η_4	P4 learning rate	✓	✓
η_5	P5 learning rate	—	✓
<i>Modulating (safety-valve) parameters</i>			
η_1	P1 learning rate	—	✓
η_2	P2 learning rate	—	✓
<i>Neutral parameters (selected)</i>			
base_1^{AG}	P1 AG baseline payoff	—	—
base_3^{HTN}	P3 HTN baseline payoff	—	—
base_4^{FI}	P4 FI baseline payoff	—	—
$\text{cross}_{24}^{TI,FI}$	P2 TI – P4 FI cross-effect	—	—
$\text{cross}_{45}^{FI,SF}$	P4 FI – P5 SF cross-effect	—	—

committed to foreign capital inflow ($P4 = FI$), which in practice implies that foreign capital liberalization must be accompanied by a simultaneous shift toward quant-friendly institutions (TI) and mature regulations (HTN). Without this complementary institutional environment, foreign capital alone may not sustain the mature-market trajectory.

A historical episode illustrates how a single-point reform can interact disastrously with the remaining gates. On January 1, 2016, the China Securities Regulatory Commission (CSRC) introduced a market-wide circuit breaker mechanism, intended as a stabilizing device, on top of the existing T+1 settlement rule and $\pm 10\%$ price limits. Within its first four trading days, the mechanism was triggered twice, forcing the exchanges to close early on both occasions, and it was suspended on January 8, 2016 (CSRC, 2016). Rather than damping volatility, the circuit breaker created a magnet effect that accelerated sell-offs toward the trigger thresholds, while T+1 prevented intraday exit and the price limits prevented full price adjustment—precisely the “double lock” mechanism formalized in our model. The episode stands as a textbook case of the complementarity trap: a reform designed in isolation against the background of mature-market institutions became destabilizing when layered onto a market architecture whose remaining gates it did not address.

More recently, regulatory responses to algorithmic trading itself confirm the within-stack logic of institutional adjustment. In May 2024, the CSRC issued the *Administrative Provisions on Program Trading in the Securities Market (Trial)*, which took effect on October 8, 2024, introducing reporting requirements, enhanced monitoring, and differentiated treatment for high-frequency trading (CSRC, 2024). Notably, the response to AI-driven and algorithmic trading has so far taken the form of tighter oversight within the existing institutional stack—registration, monitoring, and trading constraints—rather than a redesign of the settlement, price-limit, and short-selling gates that determine how algorithmic strategies interact with market structure. This pattern is consistent with the model’s prediction that, absent simultaneous reform across all three gates, the system accommodates new technology by reinforcing the surrounding constraints rather than by relaxing them.

4.2 The Three-Gates Framework: A New Policy Paradigm

The standard policy framework in institutional economics is instrument-level optimization: find the optimal value of each policy instrument (T+1, price limits, short-selling limits) independently. Our multi-population EGT implies a fundamentally different approach: *architecture-level engineering*.

In architecture-level engineering, the policy question is not “what is the optimal T+1 rule?” but “which institutional stack should the system commit to?” The two stable attractors represent two irreconcilable institutional architectures: Any policy that mixes elements across the two stacks (Section 2.2)—e.g., TI + HTT + SC, or NS + HTN + SF—creates an unstable trap. The policy challenge is not to find a compromise but to commit to one stack and engineer the complementary transition.

Relation to the institutional complementarity literature. Our “stack” concept is closely related to the complementarity framework developed by [10] and applied to institutional analysis by [13]. In Milgrom & Roberts’ formulation, a set of activities is complementary if increasing one raises the marginal return to increasing the others. Our model extends this logic to a dynamic setting: the “marginal return” to reforming one institutional dimension (e.g., removing T+1) depends on the state of the other dimensions (e.g., whether short-selling is available). When the other dimensions are stuck in the traditional configuration, the marginal return to reforming any

single dimension is negative or zero—reform creates an unstable trap rather than a path to improvement. This is the *dynamic complementarity trap*: not only are the institutional dimensions complementary in steady state, but the transition path itself requires simultaneous movement across all dimensions.

Where our model differs from the classical complementarity literature is in its asymmetric multi-population representation. [10] analyze complementarity among activities within a single firm; [13] extends this to institutional modules within a single economy. Our model decomposes the institutional stack into five distinct populations with different strategy sets, allowing us to study how cross-population externalities create complementarity. Tripartite evolutionary game models applied to hierarchical governance structures have similarly demonstrated that multi-population decomposition is necessary to capture the strategic interdependence among actors with asymmetric information and adjustment speeds [7–9].

Welfare considerations. We label Stable B the “mature market” and Stable A the “traditional market,” but these are normative designations based on three efficiency criteria: (1) *price discovery efficiency* (absence of T+1 and price limits allows prices to fully reflect information); (2) *capital allocation efficiency* (foreign capital inflow and short-selling availability enable risk-sharing and hedging); and (3) *international integration* (HTN + FI + SF aligns with global market standards). We acknowledge, however, that Stable A may be preferable under alternative policy objectives, such as *systemic risk containment* (state-controlled institutions and short-selling constraints limit speculative bubbles) or *domestic capital retention* (FO and SC prevent capital flight during crises). Our normative contribution is not to advocate for one stack over the other, but to demonstrate that mixing elements across stacks creates welfare-dominated traps, and that escaping any trap requires coordinated movement across all institutional gates simultaneously.

4.3 Why Both Stable Attractors Have AI-Driven Retail

A striking finding is that both Stable A and Stable B feature $P1 = AL$ (AI-driven retail). This contradicts the common narrative that AI adoption is a force for “market modernization” that naturally pushes the system toward the mature market. Instead, our model shows that AI adoption is *orthogonal* to institutional maturity:

- In Stack A, AI is used by retail investors but is harvested by state-controlled institutions (NS) within a constrained market (LTT + SC).
- In Stack B, AI is complemented by quant institutions (TI) and market efficiency tools (HTN + SF).

The policy implication is profound: *regulating AI adoption is not the key to reform*. The key is the institutional environment that surrounds AI adoption. Banning AI (as in Counterfactual 2) returns the system to the pre-AI traditional market (E+NS+LTT+FO+SC)—an argmax terminal state that has not been proven to be a stable attractor—by eliminating the destabilizing force. Among the counterfactual scenarios, only Full Foreign Capital Open converges to a proven stable attractor, namely Stable B. But promoting AI without institutional reform (as in the baseline) locks the system in the HTT trap.

4.4 Policy Implications for Institutional Reform

The model’s policy implications are most naturally organized by the three institutional gates, followed by the cross-cutting issues of timing and monitoring.

The first gate: T+1 settlement and price limits. Counterfactual simulations (Section 3.4) show that single-point reforms of the trading-rule regime fail: removing T+1 or price limits alone converges to the traditional stable attractor or the pre-AI market, because the complementarity of the remaining gates re-locks the system. The baseline configuration sits in a zone where the mature market is theoretically stable but practically inaccessible from the current initial condition, so reform requires coordinated and timely movement across all gates rather than indefinitely deferred piecemeal adjustment. This implies that incremental reform of a single gate—for example, gradually widening price limits from $\pm 5\%$ to $\pm 10\%$ to $\pm 20\%$ —may be politically feasible but is *structurally insufficient* to reach the mature market. The experience of the STAR Market, which introduced $\pm 20\%$ price limits but did not remove T+1 for individual stocks or fully liberalize short-selling, is consistent with this prediction: partial reform creates a “quasi-mature” enclave for a subset of stocks but does not transform the overall market architecture. In an August 2023 press Q&A, the CSRC indicated that it was studying T+0 trading but that the conditions were not yet ripe at the current stage, further illustrating the challenge of sequential reform: even a well-designed single-gate reform risks being trapped by the complementarity of the remaining gates.

The second gate: short-selling constraints. Short-selling liberalization alone remains trapped in an unstable mixed state (Section 3.4); it works only as part of the full stack. Moreover, the parameter sensitivity analysis (Section 3.5) shows that the mature market is accessible only from a narrow transitional band ($0.06 < \text{base}_5^{SF} < 0.08$; Table 7)—within which the scan samples a single point (0.07), so the grid resolution does not support a more precise localization of the band edges: reform that is too weak preserves the traditional trap, while reform that is too strong pushes the system toward a new quant-dominated HTT trap. Reform intensity is therefore a **Goldilocks problem**: parameters must be calibrated within the narrow band that navigates the basin of the mature market. The Chinese regulatory trajectory is consistent with this calibration difficulty—the expansion of short-selling targets to approximately 2,200 stocks in October 2022 (SSE/SZSE) was followed by a tightening of short-selling rules from 2023 and the suspension of the securities refinancing business in July 2024 (CSRC), a reversal that matches the model’s lock-in prediction that the constrained configuration reasserts itself whenever market stress materializes.

The third gate: foreign capital access. Among the single-point counterfactuals, only Full Foreign Capital Open reaches the mature market (Stable B), and even it requires an initial commitment to the full institutional stack (Section 3.4). Foreign capital inflow is a complement to, not a substitute for, domestic institutional modernization. The model thus suggests a *phased simultaneous reform* strategy comprising three parallel measures, all initiated within the same reform window: (i) widening price limits to $\pm 20\%$ across all boards (already implemented for the STAR Market); (ii) removing T+1 for all stocks while expanding short-selling targets; and (iii) further facilitating foreign capital entry through the deepening of the Stock Connect

mechanism and streamlined QFII/RQFII admission procedures (the quota limits having been abolished in September 2019). The three measures are *initiated together*: their legislative and effective dates may be staggered within the window, but no gate may wait for another gate to be *completed* before its own reform begins. The key is that the window must address *all three gates* simultaneously. Two clarifications discipline this recommendation. First, “simultaneous” does not mean *at the same calendar instant*: it means that the three gate reforms are *initiated within the same reform window*—a bounded policy period during which all three are credibly underway, so that the complementarity of the unreformed gates has no time to re-lock the system before the others move. Second, two-parameter reform is structurally insufficient: it succeeds in only 9.4% of the scanned two-parameter policy space (Section 3.8), and single-gate reform succeeds only within narrow parameter windows (Sections 3.5, 3.10). The sequencing experiment of Section 3.10 makes the “window” precise: the inter-gate lag must remain below the critical interval ($\Delta < 20,000$ steps in our parameterization), and the trading-rule gate (P3) must never be omitted—it is the only gate whose isolated reform fails outright, and schemes that lead with it are the slowest (Section 3.10).

Timing and monitoring. The reform window is time-dependent. Because the system operates on two timescales, policymakers cannot evaluate reform success at the short-run horizon: a reform that fails completely at 5,000 steps (0% success from a perturbed traditional-market corner) succeeds with high probability (81.5%) at 50,000 steps, and from a mixed initial condition a full five-parameter reform succeeds immediately (100%; Section 3.9). The optimal strategy is to initiate reform while the system retains heterogeneity and to maintain the commitment through the slow transient decay—intermittent or reversed commitments reset the system to the short-run trap before the long-run basin can be reached (Section 3.8). The sequencing experiment (Section 3.10) quantifies the cost of delay directly: simultaneous initiation reaches the mature market at step 2,665, each doubling of the inter-gate lag multiplies the transition time, and delay beyond the critical interval ($\Delta = 20,000$ steps) forfeits the mature market altogether. The stochastic analysis (Section 3.7) further qualifies the window: it is not a static calm period but recurs as transient convergence phases between perturbations (residence at Stable A is itself short-lived under noise, with a median first-escape time of 5,901 steps at $\sigma = 0.005$); each recurrence, however, begins with a new HTT trapping episode, so windows should be used early rather than awaited. Finally, AI adoption is not a neutral technology that can be layered onto any institutional framework: its effects are time-scale-dependent and non-monotonic, with only a narrow moderate-penetration band ($\lambda = 0.8\text{--}1.5$) offering a traditional reform path and excessive penetration ($\lambda \geq 2.9$) threatening even the mature market (Section 3.13). Rather than predicting *when* a transition will occur, regulators should monitor the *structural correlations* between observable indicators—for example, a sudden weakening or sign reversal of the correlation between AI tool adoption and margin-trading balance—as early-warning signals that the system is approaching a basin boundary (Section 3.12); such monitoring requires only tracking the direction of correlation shifts, not precise parameter estimation.

4.5 Limitations

Well-mixed assumption: The model assumes well-mixed populations. In reality, retail investors cluster in information cocoons (Snowball, Douyin), and institutions have localized influence. Future research may examine network topology effects.

Static cross-effects: The cross-population payoff matrices are static. In reality, these effects evolve over time (e.g., as foreign capital becomes more sophisticated, $\text{cross}_{43}^{\text{FI,HTN}}$ may change). The model captures the current institutional stack but does not endogenize its evolution.

No financial market dynamics: There are no prices, returns, or wealth dynamics. The model is a pure strategy-evolution game, not a financial market simulation. This is deliberate: the goal is to understand the strategic structure of institutional interaction, not to predict prices.

Parameter calibration: The base payoffs and cross-effects are calibrated to match qualitative institutional intuitions, not fitted to micro data. However, a systematic cross-effect compression analysis (Section 3.11) reveals that the model possesses architectural redundancy: when the large majority of the 82 non-P5 cross-effects are systematically zeroed, the two-attractor landscape is preserved, because the qualitative structure is maintained by the joint cross-effect architecture rather than by any single parameter. One structural anchor parameter ($\text{cross}_{43}^{\text{FI,LTT}}$) is identified as load-bearing in the individual zeroing experiment; its removal alters the top-5 steady-state ranking. The remaining parameters modulate transient dynamics and basin boundaries. The quantitative values should be interpreted as structural benchmarks that capture the qualitative architecture of cross-population incentives, not as precise forecasts requiring micro-level calibration.

Time-scale calibration: The model operates in continuous time with fixed speed parameters. The mapping from model time to calendar time is not estimated from data: the speed parameters ($\eta_1\text{--}\eta_5$) are set to fixed values (1.0, 1.0, 0.5, 0.8, 0.6). We therefore use the convention that one time unit corresponds to roughly one trading day (Section 2) strictly as a heuristic order-of-magnitude reading of the two timescales—it cannot translate simulation horizons into precise calendar-time predictions. This means the model cannot predict when a transition will occur, only whether it is possible under a given parameter configuration. The time-scale issue is a fundamental limitation of simulation-based EGT.

5 Conclusion

This paper develops a five-population asymmetric evolutionary game to explain why AI-driven retail trading and the rise of quantitative institutions do not, by themselves, push China’s stock market toward institutional modernization. Decomposing the institutional environment into three gates—T+1 settlement, price limits, and short-selling constraints—the model uncovers a *two-timescale dynamic lock-in*: the HTT trap is locally attractive on the short timescale ($\sim 5,000$ steps) but decays toward the traditional-market stable attractor on the long timescale ($\sim 50,000$ steps), while AI adoption extends the trap’s transient lifetime and shrinks the effective reform window. All 96 pure-strategy corners are fixed points, only 2 are stable under the baseline

parameterization, and the two stable stacks are self-reinforcing configurations that cannot be mixed without creating unstable traps. The mature market is accessible from only 9.4% of the two-parameter policy space and only within narrow Goldilocks calibration bands; a critical AI-penetration band ($\lambda^* \in [2.2, 2.4]$) marks the transition from transient trapping to permanent stability, and excessive penetration ($\lambda \geq 2.9$) can collapse even the mature market. The two-attractor structure survives $\pm 50\%$ perturbation for 91.5% of the 141 payoff parameters, with the structurally active parameters concentrated in the base payoffs.

The central policy implication is that reform must be *complementary and timely*: simultaneous initiation across all three gates reaches the mature market fastest—about 14 times faster than the slowest successful sequential scheme—while single-gate reform succeeds only within narrow parameter windows and the trading-rule gate (P3) never succeeds alone. Sequential implementation is viable only within a bounded inter-gate interval: delay multiplies the transition time, and delay beyond the critical interval ($\Delta = 20,000$ steps in our parameterization) locks the system irreversibly into the HTT region (Section 3.10). Reform is most effective when initiated from a mixed-state initial condition during the dynamic lock-in phase, calibrated to the Goldilocks band—strong enough to dislodge the traditional attractor but not so strong as to create a quant-dominated trap—and sustained through the slow transient decay; monitoring shifts in the structural correlations between observable market indicators offers an early-warning signal that the system is approaching a basin boundary.

Limitations and Future Work. This model does not address the political economy of reform sequencing—why a regulator would choose to reform all five dimensions simultaneously rather than gradually. Future extensions could embed a political-economy reaction function (e.g., regulator adjusts T+1 based on volatility) to model the endogenous speed of institutional change. Additionally, decomposing AI adoption into alpha-generating vs. momentum-following sub-strategies could test whether the AI penetration paradox is attenuated or amplified by the specific type of AI tool adopted. Calibrating the model time unit to specific policy events (e.g., the introduction of the STAR Market in 2019 or the tightening of margin-trading rules in 2023) would improve the mapping from model time to calendar time.

A further limitation concerns the treatment of the institutional regime (P3) as a strategy dimension in the replicator dynamics. While the “persistence fitness” interpretation provides a coherent selection metaphor and the approximation is standard in multi-population EGT, it abstracts away from the discrete, politically-driven nature of institutional change. The frequency adjustment in the model does not correspond to any literal political process (e.g., legislative cycles, regulatory review, or interest-group bargaining). Future work could introduce a separate institutional law of motion—such as a regulator reaction function or a punctuated-equilibrium model—to better capture the political economy of reform and test whether the two-attractor structure survives under alternative institutional dynamics.

Supplementary information. The full payoff matrix (146 parameters) is reported in the Appendix (Tables A1–A6). The reproduction package (simulation code, the scripts generating the stylized proxy series of Section 3.12, and the archived simulation outputs) is publicly available on OSF at https://osf.io/v3xft/overview?view_only=

[e67cf7a236164b32be4aade2355795fd](https://osf.io/v3xft/) (anonymized view-only link for double-blind review; the repository will be made public with a DOI upon publication).

Acknowledgements. Not applicable.

Declarations

- **Conflict of interest:** The authors declare no competing interests.
- **Data availability:** This study is a numerical simulation based on an evolutionary game model; no external empirical datasets were collected or analysed. All model parameters are reported in full in the Appendix (Tables A1–A6). The reproduction package “*Reproduction package: stylized proxy correlation analysis*” is publicly available on OSF at https://osf.io/v3xft/overview?view_only=e67cf7a236164b32be4aade2355795fd (anonymized view-only link for double-blind peer review; the repository will be made public with a DOI upon publication). The package contains the simulation code, the archived simulation outputs underlying the proxy-sensitivity analysis (Table 11), and documentation enabling full verification. The stylized proxy series used in Section 3.12 (margin-trading balance, quant-fund AUM, and turnover trends) are synthetic trend curves generated by the scripts included in the package (seeded and exactly reproducible); their shapes are motivated by publicly documented A-share market trends, but they are not real market data and were not used in any statistical estimation.
- **Materials availability:** Not applicable.
- **Code availability:** The Python simulation code required to reproduce the stylized proxy series and the proxy-sensitivity analysis is included in the OSF reproduction package linked above.
- **Author contribution:** Anonymized for double-blind review.

Appendix A Complete Payoff Functions and Replicator Dynamics

This subsection provides the full mathematical specification of the payoff functions and replicator dynamics that underlie all simulations in this paper. The notation is consistent with the qualitative intuitions in Section 2.2 and the stability analysis in Section 2.4.

A.1 Payoff Function Structure

Each population i ’s payoff for playing pure strategy s_i against the pure-strategy profile $(s_1, s_2, s_3, s_4, s_5)$ follows an additive decomposition: a base payoff plus pairwise cross-population interaction terms. For P1, an additional homogeneity penalty applies when AI-driven retail strategies meet high-AI institutional regimes.

P1 (Retail). For $s_1 \in \{E, AL, AG\}$:

$$\pi_1(s_1, s_2, s_3, s_4, s_5) = \text{base}_1^{s_1} + \text{cross}_{12}^{s_1, s_2} + \text{cross}_{13}^{s_1, s_3} + \text{cross}_{15}^{s_1, s_5} + \Phi(s_1, s_3), \quad (\text{A1})$$

where $\Phi(s_1, s_3)$ is the homogeneity penalty:

$$\Phi(s_1, s_3) = \begin{cases} -\kappa_{\text{HT}} & \text{if } s_3 = \text{HTT and } s_1 \in \{\text{AL}, \text{AG}\}, \\ -\kappa_{\text{HN}} & \text{if } s_3 = \text{HTN and } s_1 \in \{\text{AL}, \text{AG}\}, \\ 0 & \text{otherwise.} \end{cases} \quad (\text{A2})$$

Here $\kappa_{\text{HT}} = 0.03$ and $\kappa_{\text{HN}} = 0.015$. The penalty captures a *signal correlation friction*: in a high-AI institutional regime (HTT or HTN), AI-driven tools are by definition widely available, so the trading signals they generate are structurally correlated—the same data, the same models, and the same execution logic produce synchronized entry and exit decisions regardless of the marginal adoption rate. This correlation degrades the informational edge of any individual AI-driven strategy, imposing a payoff cost that is *regime-conditional* rather than *frequency-dependent*. The penalty is a constant (not a function of x_1^{AL} or x_1^{AG}) because the signal saturation is determined by the institutional regime s_3 , not by the within-population strategy mix; this also preserves the standard multi-population replicator dynamics structure in which each population's payoff depends on other populations' frequencies but not on its own [1]. The penalty is larger for HTT ($\kappa_{\text{HT}} > \kappa_{\text{HN}}$) because T+1 and price limits amplify the downside of synchronized AI-driven trades by preventing immediate exit and price correction, turning correlated signals into correlated losses [30–32].

P2 (Institutions). For $s_2 \in \{\text{NS}, \text{TI}\}$:

$$\pi_2(s_1, s_2, s_3, s_4, s_5) = \text{base}_2^{s_2} + \text{cross}_{21}^{s_2, s_1} + \text{cross}_{23}^{s_2, s_3} + \text{cross}_{24}^{s_2, s_4} + \text{cross}_{25}^{s_2, s_5}. \quad (\text{A3})$$

Remark 1 (On the base payoff asymmetry: $\text{base}_2^{\text{NS}} > \text{base}_2^{\text{TI}}$) The base payoff for state-controlled institutions ($\text{base}_2^{\text{NS}} = 0.015$) slightly exceeds that for quant-friendly institutions ($\text{base}_2^{\text{TI}} = 0.010$), reflecting the *default-configuration advantage*: in the A-share market, state-controlled institutions represent the incumbent paradigm—regulatory familiarity, lower implementation costs, and political alignment reduce the baseline operating cost. The cross-effects reverse this ranking in the mature-market stack (Stable B): TI's cross-effect total (+0.055) overwhelms both its own base disadvantage (−0.005) and NS's cross-effect deficit (−0.040), yielding a net TI advantage of +0.065 vs. −0.025 for NS. This mirrors the P4 and P5 asymmetries: the base payoff favors the traditional strategy, while cross-effects create a regime-dependent switching incentive toward the modern-market strategy.

P3 (Institutional Regime). For $s_3 \in \{\text{LTT}, \text{LNT}, \text{HTT}, \text{HTN}\}$:

$$\pi_3(s_1, s_2, s_3, s_4, s_5) = \text{base}_3^{s_3} + \text{cross}_{31}^{s_3, s_1} + \text{cross}_{32}^{s_3, s_2} + \text{cross}_{34}^{s_3, s_4} + \text{cross}_{35}^{s_3, s_5}. \quad (\text{A4})$$

Remark 2 (On the economic meaning of π_3) As discussed in Section 2.1, P3 does not represent a strategic agent with a utility function. The quantity π_3 is a *persistence fitness*—a reduced-form measure of how self-reinforcing each institutional configuration is, given the strategies of other populations. The base term $\text{base}_3^{s_3}$ captures the intrinsic political stickiness of each regime: regimes that align with entrenched interest groups (e.g., LTT serving

state-controlled institutions) have higher base persistence. The cross-terms $\text{cross}_{3j}^{s_3, s_j}$ capture institutional complementarity in the sense of [13]: a positive cross-effect means that the current strategy mix in population j makes regime s_3 harder to reform (higher persistence), while a negative cross-effect means the strategy mix erodes the regime’s political viability. For example, $\text{cross}_{32}^{\text{HTT}, \text{TI}} = +0.020$ indicates that when quant-friendly institutions (TI) dominate, the HTT regime becomes more self-reinforcing—quant institutions thrive on the volatility created by T+1 and price limits, creating a political constituency that resists reform. This interpretation is consistent with the path-dependence and increasing-returns framework of [23] and the political-economy perspective of [24].

P4 (Foreign Capital). For $s_4 \in \{\text{FI}, \text{FO}\}$:

$$\pi_4(s_1, s_2, s_3, s_4, s_5) = \text{base}_4^{s_4} + \text{cross}_{41}^{s_4, s_1} + \text{cross}_{42}^{s_4, s_2} + \text{cross}_{43}^{s_4, s_3} + \text{cross}_{45}^{s_4, s_5}. \quad (\text{A5})$$

Remark 3 (On the base payoff asymmetry: $\text{base}_4^{\text{FO}} > \text{base}_4^{\text{FI}}$) A reader may notice that the base payoff for foreign capital outflow ($\text{base}_4^{\text{FO}} = 0.060$) exceeds that for inflow ($\text{base}_4^{\text{FI}} = 0.040$). This does not imply that outflow is “better for the market”; rather, π_4 measures the *fitness* of each strategy for the foreign capital population itself, not its welfare contribution to the domestic market. The asymmetry captures the *structural exit premium*: foreign investors in the A-share market face T+1 settlement, price limits, capital controls, and limited hedging instruments [33–36], so the baseline risk-adjusted return of remaining invested is lower than that of reallocating to less constrained markets (e.g., Hong Kong or U.S. equities). This premium is consistent with empirical observation: foreign ownership has long remained at only ~3–5% of total A-share market capitalization despite decades of liberalization, and northbound flows through the Stock Connect mechanism are strongly pro-cyclical, exhibiting large inflows and outflows that track domestic market conditions. The capital account remains only partially open: although the QFII/RQFII quota limits were abolished in September 2019 (SAFE, 2019), foreign investors continue to face settlement, hedging, and repatriation frictions. The persistently low foreign share despite open access reflects the same structural frictions that the base asymmetry captures. Crucially, the cross-effects with P3 reverse this ranking in the mature regime: in HTN, $\text{cross}_{43}^{\text{FI}, \text{HTN}} = +0.030$ while $\text{cross}_{43}^{\text{FO}, \text{HTN}} = -0.020$, so the total payoff to FI ($0.040 + 0.030 = 0.070$) exceeds that to FO ($0.060 - 0.020 = 0.040$). The base asymmetry thus creates a *regime-dependent switching incentive*—foreign capital rationally exits in constrained regimes and enters in mature ones—which is precisely the complementarity structure the model is designed to capture [37].

P5 (Short-Selling). For $s_5 \in \{\text{SC}, \text{SF}\}$:

$$\pi_5(s_1, s_2, s_3, s_4, s_5) = \text{base}_5^{s_5} + \text{cross}_{51}^{s_5, s_1} + \text{cross}_{52}^{s_5, s_2} + \text{cross}_{53}^{s_5, s_3} + \text{cross}_{54}^{s_5, s_4}. \quad (\text{A6})$$

Remark 4 (On the base payoff asymmetry: $\text{base}_5^{\text{SC}} > \text{base}_5^{\text{SF}}$) The base payoff for short-selling constraints ($\text{base}_5^{\text{SC}} = 0.050$) exceeds that for freedom ($\text{base}_5^{\text{SF}} = 0.030$), reflecting *institutional inertia*: in the A-share market, restricted short-selling is the default configuration—maintaining constraints imposes lower implementation costs and market-disruption risk than liberalizing. The cross-effects, however, reverse this ranking in mature environments: in the Stable B stack (AL+TI+HTN+FI), the total cross-effect payoff to SF (+0.040) outweighs

the base disadvantage (-0.020), yielding a net advantage of SF over SC (0.070 vs. 0.030). This parallels the P4 asymmetry ($\text{base}_4^{\text{FO}} > \text{base}_4^{\text{FI}}$): in both cases, the base payoff favors the “exit” or “constraint” strategy, while the cross-effects create a regime-dependent switching incentive toward the modern-market strategy. The base asymmetry is thus a feature of the complementarity structure, not an inconsistency: it ensures that liberalization (of either foreign capital or short-selling) occurs only when complementary institutions are sufficiently developed.

A key structural feature is that each population’s payoff depends on the strategies of all other populations but not on its own population’s composition (except through the replicator dynamics). This is the standard assumption in multi-population asymmetric games [1]: the payoff to a given strategy in population i is determined by the cross-population interaction structure, while intra-population frequency effects enter only through the expected payoff computation.

A.2 Expected Payoffs

The expected payoff to strategy k in population i is obtained by marginalizing the pure-strategy payoff over the strategy distributions of all other populations. For population 1 (retail), with $\mathbf{x} = (x_1, x_2, x_3, x_4, x_5)$:

$$\bar{\pi}_1^k(\mathbf{x}) = \sum_{s_2 \in \mathcal{S}_2} \sum_{s_3 \in \mathcal{S}_3} \sum_{s_4 \in \mathcal{S}_4} \sum_{s_5 \in \mathcal{S}_5} x_2^{s_2} x_3^{s_3} x_4^{s_4} x_5^{s_5} \pi_1(k, s_2, s_3, s_4, s_5), \quad k \in \{E, \text{AL}, \text{AG}\}. \quad (\text{A7})$$

Analogously, for populations 2–5:

$$\bar{\pi}_2^k(\mathbf{x}) = \sum_{s_1, s_3, s_4, s_5} x_1^{s_1} x_3^{s_3} x_4^{s_4} x_5^{s_5} \pi_2(s_1, k, s_3, s_4, s_5), \quad k \in \{\text{NS}, \text{TI}\}, \quad (\text{A8})$$

$$\bar{\pi}_3^k(\mathbf{x}) = \sum_{s_1, s_2, s_4, s_5} x_1^{s_1} x_2^{s_2} x_4^{s_4} x_5^{s_5} \pi_3(s_1, s_2, k, s_4, s_5), \quad k \in \{\text{LTT}, \text{LNT}, \text{HTT}, \text{HTN}\}, \quad (\text{A9})$$

$$\bar{\pi}_4^k(\mathbf{x}) = \sum_{s_1, s_2, s_3, s_5} x_1^{s_1} x_2^{s_2} x_3^{s_3} x_5^{s_5} \pi_4(s_1, s_2, s_3, k, s_5), \quad k \in \{\text{FI}, \text{FO}\}, \quad (\text{A10})$$

$$\bar{\pi}_5^k(\mathbf{x}) = \sum_{s_1, s_2, s_3, s_4} x_1^{s_1} x_2^{s_2} x_3^{s_3} x_4^{s_4} \pi_5(s_1, s_2, s_3, s_4, k), \quad k \in \{\text{SC}, \text{SF}\}. \quad (\text{A11})$$

The average payoff in population i is $\bar{\pi}_i(\mathbf{x}) = \sum_k x_i^k \bar{\pi}_i^k(\mathbf{x})$.

A.3 Replicator Dynamics with Population-Specific Learning Rates

Each population evolves according to the replicator equation, augmented with a population-specific learning rate η_i that governs the speed of strategic adjustment:

$$\dot{x}_i^k = \eta_i x_i^k (\bar{\pi}_i^k(\mathbf{x}) - \bar{\pi}_i(\mathbf{x})), \quad i = 1, \dots, 5, \quad k \in \mathcal{S}_i. \quad (\text{A12})$$

The learning rates $(\eta_1, \eta_2, \eta_3, \eta_4, \eta_5) = (1.0, 1.0, 0.5, 0.8, 0.6)$ reflect the empirically motivated assumption that different populations adjust at different speeds: retail investors and institutions adjust fastest ($\eta_1 = \eta_2 = 1.0$), the institutional regime adjusts more slowly ($\eta_3 = 0.5$), reflecting the political and regulatory inertia of institutional change, foreign capital adjusts at an intermediate rate ($\eta_4 = 0.8$), and the short-selling mechanism adjusts slowest ($\eta_5 = 0.6$), reflecting the regulatory complexity of short-selling reform. The learning rates do not affect the location of fixed points (which depend only on the payoff structure) but do affect the transient dynamics and the effective timescale of convergence, which is central to the two-timescale analysis in Section 3.1. Equation (A12) reduces to the standard replicator dynamics (Section 2.4) when all $\eta_i = 1$. The learning rates are therefore a generalization that introduces realistic heterogeneity in adjustment speeds across populations while preserving the fixed-point structure.

A.4 Base Payoff Values

Table A1 reports the baseline base payoff values for all five populations. These are the $\text{base}_i^{s_i}$ terms in Eqs. (1)–(6).

Table A1 Baseline base payoffs for all five populations.

Population	Strategy	Base Payoff
P1 (Retail)	E (passive)	−0.012
P1 (Retail)	AL (AI-driven learning)	+0.022
P1 (Retail)	AG (AI-driven gambling)	+0.010
P2 (Institutions)	NS (state-controlled)	+0.015
P2 (Institutions)	TI (quant-friendly)	+0.010
P3 (Regime)	LTT (low AI, T+1, limits)	+0.080
P3 (Regime)	LNT (low AI, T+1, no limits)	+0.070
P3 (Regime)	HTT (high AI, T+1, limits)	+0.075
P3 (Regime)	HTN (high AI, no T+1, no limits)	+0.070
P4 (Foreign Capital)	FI (inflow)	+0.040
P4 (Foreign Capital)	FO (outflow)	+0.060
P5 (Short-Selling)	SC (constrained)	+0.050
P5 (Short-Selling)	SF (free)	+0.030
Homogeneity penalty	κ_{HT}	0.030
Homogeneity penalty	κ_{HN}	0.015
Learning rates	$(\eta_1, \eta_2, \eta_3, \eta_4, \eta_5)$	(1.0, 1.0, 0.5, 0.8, 0.6)

A.5 Cross-Population Interaction Matrices

The cross-effect $\text{cross}_{ij}^{s_i, s_j}$ denotes the additive payoff contribution to population i from the interaction of its strategy s_i with population j ’s strategy s_j . Tables A2–A6 report the complete cross-population interaction matrices for all five populations. Each row corresponds to a strategy of the focal population i ; each column corresponds

to a strategy of the interacting population j . Values are calibrated within the range $[-0.10, +0.10]$ as described in Section 2.2.

Remark 5 (On the asymmetry of cross-effects) A reader may notice that $\text{cross}_{ij}^{s_i, s_j} \neq \text{cross}_{ji}^{s_j, s_i}$ in general. For example, $\text{cross}_{12}^{E, \text{NS}} = +0.015$ (the payoff gain to passive retail from state-controlled institutions) differs from $\text{cross}_{21}^{\text{NS}, E} = +0.020$ (the payoff gain to state-controlled institutions from passive retail). This asymmetry is not a calibration artifact but a structural feature of the model, reflecting three economic realities. First, in multi-population asymmetric games, each population has an independent payoff function; there is no theoretical requirement that inter-population effects be symmetric [1, 2]. Second, economic externalities are inherently directional: the constraint that T+1 and price limits impose on retail investors (cross_{13}) is qualitatively different from the effect of retail strategy composition on institutional persistence (cross_{31})—the former is a direct payoff impact, the latter a change in self-reinforcement fitness. Third, the magnitudes reflect asymmetries in economic leverage: passive retail (E) is more valuable to state-controlled institutions (NS, $+0.020$) than state-controlled institutions are to passive retail ($+0.015$), because passive retail provides political legitimation and fee revenue, while NS provides only modest stability benefit to retail. Enforcing symmetry would collapse these distinct economic channels into a single parameter, eliminating precisely the directional structure that the model is designed to capture.

Table A2 Cross-effects for P1 (Retail): $\text{cross}_{1j}^{s_1, s_j}$.

P1 strategy	P2		P3				P5	
	NS	TI	LTT	LNT	HTT	HTN	SC	SF
E	+0.015	+0.005	0.000	+0.005	-0.005	0.000	+0.005	-0.005
AL	-0.025	+0.020	+0.025	+0.015	-0.035	+0.030	-0.005	+0.005
AG	-0.010	-0.025	+0.005	0.000	0.000	+0.015	-0.005	+0.005

Table A3 Cross-effects for P2 (Institutions): $\text{cross}_{2j}^{s_2, s_j}$.

P2 strategy	P1			P3				P4		P5	
	E	AL	AG	LTT	LNT	HTT	HTN	FI	FO	SC	SF
NS	+0.020	-0.015	-0.030	+0.015	+0.010	-0.010	-0.005	-0.010	+0.005	+0.010	-0.010
TI	+0.005	+0.025	+0.035	-0.010	-0.005	+0.020	+0.010	+0.010	-0.005	-0.005	+0.010

Table A4 Cross-effects for P3 (Institutional Regime): $\text{cross}_{3j}^{s_3, s_j}$.

P3 strategy	P1			P2		P4		P5	
	E	AL	AG	NS	TI	FI	FO	SC	SF
LTT	+0.030	-0.015	-0.010	+0.025	-0.015	-0.020	+0.010	+0.015	-0.010
LNT	+0.025	-0.010	-0.005	+0.020	-0.010	-0.015	+0.008	+0.005	+0.005
HTT	-0.010	+0.025	+0.015	-0.015	+0.020	-0.015	+0.008	+0.010	-0.005
HTN	-0.020	+0.010	+0.005	-0.010	+0.015	+0.020	-0.010	-0.005	+0.015

Table A5 Cross-effects for P4 (Foreign Capital): $\text{cross}_{4j}^{s_4, s_j}$.

P4 strategy	P1			P2		P3				P5	
	E	AL	AG	NS	TI	LTT	LNT	HTT	HTN	SC	SF
FI	+0.020	-0.015	-0.010	-0.015	+0.015	-0.040	-0.030	-0.010	+0.030	-0.005	+0.010
FO	-0.015	+0.010	+0.005	+0.010	-0.010	+0.030	+0.020	+0.005	-0.020	+0.005	-0.005

Table A6 Cross-effects for P5 (Short-Selling): $\text{cross}_{5j}^{s_5, s_j}$.

P5 strategy	P1			P2		P3				P4	
	E	AL	AG	NS	TI	LTT	LNT	HTT	HTN	FI	FO
SC	+0.005	-0.005	-0.005	+0.010	-0.005	+0.010	+0.005	+0.010	-0.005	-0.005	+0.005
SF	-0.005	+0.005	+0.005	-0.005	+0.010	-0.005	+0.005	-0.005	+0.010	+0.010	-0.005

References

- [1] Hofbauer, J., Sigmund, K.: Evolutionary Games and Population Dynamics. Cambridge University Press, Cambridge, UK (1998)
- [2] Weibull, J.W.: Evolutionary Game Theory. MIT Press, Cambridge, MA (1995)
- [3] Ritzberger, K., Weibull, J.W.: Evolutionary selection in normal-form games. *Econometrica* **63**(6), 1371–1399 (1995) <https://doi.org/10.2307/2171773>
- [4] Zhan, K., Pu, Z.: Carbon market and emission reduction: evidence from evolutionary game and machine learning. *Humanities and Social Sciences Communications* **12**(1), 488 (2025) <https://doi.org/10.1057/s41599-025-04793-0>
- [5] Liang, X., Liu, Y., Wang, N.: Evolutionary game analysis of quality regulation in live streaming e-commerce: insights from china. *Humanities and Social Sciences Communications* **12**(1), 1456 (2025) <https://doi.org/10.1057/s41599-025-05757-0>
- [6] Zhu, J., Baker, J.S., Song, Z., Yue, X.-G., Li, W.: Government regulatory policies for digital transformation in small and medium-sized manufacturing enterprises: an evolutionary game analysis. *Humanities and Social Sciences Communications* **10**(1), 751 (2023) <https://doi.org/10.1057/s41599-023-02250-4>
- [7] Fan, W.-J., Fang, Y., Jiang, R.-B.: An analysis of optimal equilibrium in the carbon trading market—from a tripartite evolutionary game perspective. *International Review of Financial Analysis* **96**, 103629 (2024) <https://doi.org/10.1016/j.irfa.2024.103629>
- [8] Mu, G., Wang, Y., Gao, N., Li, X.: A novel tripartite evolutionary game model for internet consumer financial regulation. *Sage Open* **13**(3) (2023) <https://doi.org/10.1177/21582440231194212>
- [9] Liu, C., Li, W., Chang, L., Ji, Q.: How to govern greenwashing behaviors in green finance products: A tripartite evolutionary game approach. *Financial Innovation* **10**(1), 34 (2024) <https://doi.org/10.1186/s40854-023-00549-4>
- [10] Milgrom, P., Roberts, J.: Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* **58**(6), 1255–1277 (1990) <https://doi.org/10.2307/2938316>
- [11] Arthur, W.B.: Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal* **99**(394), 116–131 (1989) <https://doi.org/10.2307/2234208>
- [12] David, P.A.: Clio and the economics of qwerty. *The American Economic Review* **75**(2), 332–337 (1985)

- [13] Aoki, M.: *Toward a Comparative Institutional Analysis*. MIT Press, Cambridge, MA (2001)
- [14] Hall, P.A., Soskice, D.: *Varieties of Capitalism: The Institutional Foundations of Comparative Advantage*. Oxford University Press, Oxford, UK (2001)
- [15] Amable, B.: Institutional complementarity and diversity of social systems of innovation and production. *Review of International Political Economy* **7**(4), 645–687 (2000) <https://doi.org/10.1080/096922900750034655>
- [16] Hendershott, T., Jones, C.M., Menkveld, A.J.: Does algorithmic trading improve liquidity? *The Journal of Finance* **66**(1), 1–33 (2011) <https://doi.org/10.1111/j.1540-6261.2010.01624.x>
- [17] Lopez-Lira, A., Tang, Y.: Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models. arXiv:2304.07619 (2023)
- [18] Even-Tov, O., Lourie, B., Munevar, K., Nekrasov, A.: The Effect of AI on Retail Investor Behavior: Evidence from Account-Level Trading Data. SSRN Working Paper No. 5778246 (2025). <https://doi.org/10.2139/ssrn.5778246>
- [19] Gimmelberg, D., Ludviga, I.: Strategic complexity and behavioral distortion: Retail investing under large language model augmentation. *International Journal of Financial Studies* **13**(4), 210 (2025) <https://doi.org/10.3390/ijfs13040210>
- [20] Greif, A.: *Institutions and the Path to the Modern Economy: Lessons from Medieval Trade*. Cambridge University Press, Cambridge, UK (2006)
- [21] Young, H.P.: *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions*. Princeton University Press, Princeton, NJ (1998)
- [22] Li, J., Milgrom, P.: Scope, imperfections, and the evolution of forgiveness. *American Economic Review* **98**(1), 160–179 (2008) <https://doi.org/10.1257/aer.98.1.160>
- [23] North, D.C.: *Institutions, Institutional Change and Economic Performance*. Cambridge University Press, Cambridge, UK (1990)
- [24] Acemoglu, D., Johnson, S.: Unbundling institutions. *Journal of Political Economy* **113**(5), 949–995 (2005)
- [25] Pierson, P.: Increasing returns, path dependence, and the study of politics. *American Political Science Review* **94**(2), 251–267 (2000)
- [26] Taylor, P.D., Jonker, L.B.: Evolutionarily stable strategies and game dynamics. *Mathematical Biosciences* **40**(1-2), 145–156 (1978) [https://doi.org/10.1016/0025-5564\(78\)90077-9](https://doi.org/10.1016/0025-5564(78)90077-9)

- [27] Bris, A., Goetzmann, W.N., Zhu, N.: Efficiency and the bear: Short sales and markets around the world. *The Journal of Finance* **62**(3), 1029–1079 (2007) <https://doi.org/10.1111/j.1540-6261.2007.01230.x>
- [28] Chang, E.C., Cheng, J.W., Yu, Y.: Short-sales constraints and price discovery: Evidence from the Hong Kong market. *The Journal of Finance* **62**(5), 2097–2121 (2007) <https://doi.org/10.1111/j.1540-6261.2007.01270.x>
- [29] Fudenberg, D., Levine, D.K.: *Imitation and Price Competition*. MIT Press, Cambridge, MA (2006)
- [30] Chen, T., Gao, Z., He, J., Jiang, W., Xiong, W.: Daily price limits and destructive market behavior. *Journal of Econometrics* **208**(1), 249–264 (2019) <https://doi.org/10.1016/j.jeconom.2018.09.014>
- [31] Cho, D.D., Russell, J., Tiao, G.C., Tsay, R.: The magnet effect of price limits: Evidence from high-frequency data on Taiwan stock exchange. *Journal of Empirical Finance* **10**(1-2), 133–168 (2003) [https://doi.org/10.1016/S0927-5398\(02\)00024-5](https://doi.org/10.1016/S0927-5398(02)00024-5)
- [32] Hsieh, P.-H., Kim, Y.H., Yang, J.J.: The magnet effect of price limits: A logit approach. *Journal of Empirical Finance* **16**(5), 830–837 (2009) <https://doi.org/10.1016/j.jempfin.2009.06.002>
- [33] Chen, X., Liu, Y., Zeng, T.: Does the T+1 rule really reduce speculation? evidence from Chinese stock index ETF. *Accounting & Finance* **57**(5), 1287–1313 (2017) <https://doi.org/10.1111/acfi.12330>
- [34] Chen, G., Rui, O.M., Wang, S.S.: The effectiveness of price limits and stock characteristics: Evidence from the Shanghai and Shenzhen stock exchanges. *Review of Quantitative Finance and Accounting* **25**(2), 159–182 (2005) <https://doi.org/10.1007/s11156-005-4247-7>
- [35] Deb, S.S., Kalev, P.S., Marisetty, V.B.: Are price limits really bad for equity markets? *Journal of Banking & Finance* **34**(10), 2462–2471 (2010) <https://doi.org/10.1016/j.jbankfin.2010.04.001>
- [36] Kim, K.A., Rhee, S.G.: Price limit performance: Evidence from the Tokyo stock exchange. *The Journal of Finance* **52**(2), 885–901 (1997) <https://doi.org/10.1111/j.1540-6261.1997.tb04827.x>
- [37] Hirschman, A.O.: *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Harvard University Press, Cambridge, MA (1970)