Racing the Benchmark: Competition and Robust Risk-Taking in Financial Markets

Lei Pan*

Abstract

We develop a continuous-time game-theoretic model in which financial institutions jointly choose portfolio allocation and surplus risk transfer under regime-switching asset returns, ambiguity aversion, and relative performance concerns. We derive closed-form Nash equilibrium strategies in both the finite-n game and the mean field limit. The model shows that benchmarking incentives generate a competitive amplification mechanism: stronger relative performance concerns increase equilibrium risky investment and weaken the incentive to transfer surplus risk, producing an arms-race effect that becomes unstable as average competition approaches its critical upper bound. Ambiguity aversion works in the opposite direction by reducing effective risk tolerance, compressing equilibrium leverage, and strengthening prudent risk management. The regime-switching structure makes these effects strongly state dependent. Under the empirical calibration, bull-market conditions support aggressive risky investment, while bear-market conditions with negative excess returns can induce short positions rather than merely smaller long positions. The multi-asset calibration further shows that high cross-asset correlation can generate leveraged long–short spread positions when institutions exploit differences in risk-adjusted returns, while negative correlation and strong model confidence sharply magnify diversification-driven risk-taking. On the risk-control margin, highly competitive and risk-tolerant institutions may become net absorbers of surplus risk, implying that systemic fragility can arise not only from crowded asset portfolios but also from endogenous shifts in risk bearing. The results provide new insights into benchmarking, model confidence, leverage, long–short crowding, risk transfer, and systemic fragility in modern financial markets.

Keywords: ambiguity aversion; relative performance; mean field games; regime switching; risk transfer
JEL Codes: C73; D81; G11; G23; G32

1 Introduction

Every bull market creates a seductive narrative. Returns are high, volatility appears manageable, and institutions that load aggressively into risky positions are often celebrated as visionary rather than fragile. But the story changes quickly when the market turns. In bad states, what had looked like skill can be reinterpreted as leverage, what had looked like conviction can look like crowding, and what had looked like prudent optimization can suddenly appear to rest on a misspecified model. For financial institutions, this tension is especially sharp because they are judged not only by absolute wealth, but by how they perform relative to rivals. A manager can survive a modest loss; trailing the peer group is often harder to explain. The modern institution therefore operates under a double pressure: it must compete in a league table while simultaneously doubting whether the stochastic environment used to guide its decisions is the right one.

*Corresponding author. School of Accounting, Economics, and Finance, Curtin University, Kent Street, Bentley, Australia. Email: lei.pan@curtin.edu.au.

That double pressure motivates the present paper. We study a class of financial institutions that allocate wealth across a risk-free asset and multiple risky assets, while also managing non-tradable surplus risk through proportional risk transfer. Their objective is not purely absolute wealth maximization. Instead, each institution cares about terminal wealth relative to the industry average, so strategic interaction enters directly through the payoff function. At the same time, institutions are ambiguity averse: they suspect that the benchmark probability law governing returns and surplus shocks may be misspecified, and therefore act robustly against adverse drift distortions. Finally, the market environment itself is not static. Asset returns and volatilities change across bull and bear regimes governed by a finite-state Markov chain. This combination captures a familiar feature of real financial markets: competition is fiercest exactly when the reliability of the model is most questionable and when the economy is most likely to switch states.

The economics of ambiguity provides the first pillar of our analysis. The foundational insight is that decision makers may evaluate acts using a set of plausible priors rather than a single subjective probability measure (Gilboa and Schmeidler 1989). In intertemporal and asset-pricing environments, this insight has been extended to Knightian uncertainty, recursive multiple-priors preferences, and smooth ambiguity preferences, each offering a different formalization of how agents fear misspecification and process unresolved model uncertainty (Epstein and Wang 1994, Chen and Epstein 2002, Epstein and Schneider 2003, Klibanoff et al. 2005, 2009). A complementary robust-control tradition interprets ambiguity aversion as a concern about model misspecification disciplined by an entropy penalty, thereby yielding tractable continuous-time decision rules and clear comparative statics (Hansen and Sargent 2001, Anderson et al. 2003a, Hansen et al. 2006, Hansen and Sargent 2010). In financial applications, robust methods have been shown to alter portfolio rules, compress effective risk tolerance, generate under-diversification, reshape liquidity demand, and improve the empirical performance of asset-pricing models when agents entertain doubts about drift, volatility, jumps, or information quality (Maenhout 2004a, 2006, Routledge and Zin 2009, Branger and Larsen 2013, Jahan-Parvar and Liu 2014, Jeong et al. 2015, Pham et al. 2022). What emerges from this literature is a disciplined message: ambiguity does not simply add another source of risk; it changes the way institutions interpret all risks.

A second strand of literature studies the incentives created by delegated investment and relative performance evaluation. The central lesson is that portfolio choice changes materially once managers care about rankings, fund flows, compensation benchmarks, or career concerns rather than only absolute returns. Tournament-style incentives can induce managers to alter risk-taking dynamically over the evaluation horizon (Brown et al. 1996, Chevalier and Ellison 1997). Benchmarking and tracking-error concerns change the state dependence of optimal strategies and the effectiveness of formal risk constraints (Basak et al. 2006, Zhao 2007, Basak et al. 2008). The measurement of performance itself is nontrivial, because evaluation procedures can distort incentives and even invite manipulation if they are not robust to strategic behavior (Goetzmann et al. 2007, Chan et al. 2009). At the equilibrium level, relative-wealth concerns can amplify price pressure, fuel bubbles, and increase volatility when managers internalize how flows and reputation respond to peer comparison (DeMarzo et al. 2008, Guerrieri and Kondor 2012, Basak and Pavlova 2013, Basak and Makarov 2014, Buffa et al. 2022). In continuous-time settings, explicit portfolio games under relative performance concerns show that competition can have first-order effects on risky investment even when agents are otherwise standard expected-utility maximizers (Espinosa and Touzi 2015, Lacker and Zariphopoulou 2019a). Yet most of this literature takes the stochastic environment as given, or abstracts from ambiguity and from institutional risk transfer decisions outside the traded asset menu.

A third literature emphasizes that the investment opportunity set itself is regime dependent rather than stationary. Starting from Hamilton's seminal Markov-switching framework, macro-financial regimes

have become a workhorse device for modeling sudden and persistent changes in expected returns, volatilities, and correlations (Hamilton 1989, 1990). In finance, regime-switching methods have been used to study the connection between stock market volatility and the business cycle, the international allocation of risky assets, and the dynamic implications of multivariate state shifts for optimal portfolio choice (Hamilton and Lin 1996, Ang and Bekaert 2002, Guidolin and Timmermann 2007, Ang and Timmermann 2012). This literature is especially relevant here because benchmarking incentives are unlikely to have the same bite in a calm bull market as in a stressed bear market. When Sharpe ratios collapse and volatility spikes, the temptation to chase peers, de-risk, or transfer operational risk may all change simultaneously. A regime-switching structure is therefore not a cosmetic embellishment, but a natural way to embed macro-financial instability directly into the strategic problem.

There is also a related body of work on institutional risk control outside the pure portfolio margin. Financial institutions are not only investors; they are also carriers of liabilities, operating expenditures, underwriting exposure, and other non-tradable risks. Classical continuous-time studies show that optimal investment cannot be separated cleanly from ruin control, liability management, or insurance decisions (Browne 1995, 1997). In insurance and actuarial settings, optimal reinsurance and dividend policies reveal that external risk transfer is itself a strategic margin of adjustment rather than a passive background contract (Azcue and Muler 2005). This insight is crucial for our setting. Once institutions care about relative performance, choosing how much surplus risk to retain or transfer becomes part of competitive positioning, not merely a back-office hedging decision.

Despite substantial progress in each of these literatures, an important gap remains. Robust portfolio models are often single-agent and therefore mute the strategic escalation created by benchmarking. Relative-performance games usually take beliefs as known and rarely allow for entropy-penalized ambiguity aversion. Regime-switching portfolio models typically focus on asset allocation under exogenous preferences, without embedding large-population strategic interaction. Risk-transfer models, in turn, generally abstract from peer comparison and from the way market states affect the attractiveness of risky investment. As a result, we still lack a tractable theory of how ambiguity aversion, macro-financial regime shifts, peer competition, and risk transfer interact inside one unified continuous-time environment. That omission matters economically because these forces are complements, not substitutes. Institutions do not first solve a portfolio problem, then a benchmarking problem, then a model-uncertainty problem, and finally a risk-transfer problem. They confront all of them at once.

This paper fills that gap by developing a continuous-time game among competing financial institutions with four central ingredients. First, institutions trade in a multi-asset financial market with regime-dependent investment opportunities, so the same strategic agent faces different mean-variance tradeoffs in bull and bear states. Second, each institution manages surplus risk through proportional transfer, creating an explicit interaction between portfolio allocation and operational risk control. Third, ambiguity aversion is modeled through multiplier preferences, so robust decision making remains analytically tractable while preserving the core insight that model uncertainty shrinks effective risk tolerance. Fourth, institutions care about relative performance, which transforms a standard control problem into an n-agent non-zero-sum game and, in the large-population limit, a mean field game. The result is a framework in which traded risk, non-traded risk, competition, and model uncertainty are determined jointly rather than sequentially.

Our contribution is fivefold. First, we provide a unified theoretical framework that integrates ambiguity aversion, regime switching, multi-asset portfolio choice, relative performance concerns, and endogenous risk transfer. To the best of our knowledge, these features have not previously been combined in a single closed-form game-theoretic model. Second, we derive explicit constant Nash equilibrium

strategies in the finite-n game and show how they converge to a corresponding mean field equilibrium. This delivers an analytically transparent bridge between strategic interaction in a finite industry and the large-population limit. Third, the model identifies a sharp economic mechanism through which ambiguity aversion acts as a brake on competitive escalation: by reducing effective risk tolerance, ambiguity dampens risky investment and strengthens the incentive to transfer surplus risk, thereby partially offsetting the arms-race logic generated by relative performance concerns. Fourth, because the framework is multi-asset and regime-switching, it yields novel comparative statics on how correlations, bull-bear transitions, and heterogeneity in competition intensity shape equilibrium leverage, long–short exposure, and risk control. Fifth, we calibrate the model using Chinese equity-market data and show that the theoretical mechanisms are quantitatively meaningful: competition can generate explosive portfolio escalation, bear-market excess returns can reverse the sign of equilibrium investment, high empirical asset correlation can produce leveraged long–short spread positions, and highly competitive institutions may become net absorbers of surplus risk. These results speak directly to macroprudential debates about benchmarking, model governance, crowded trades, risk-transfer markets, and systemic fragility in delegated finance.

The economic intuition of the model is simple but powerful. Competition pushes institutions toward larger risky positions because each agent wants to keep pace with the industry average. Regime dependence makes that pressure state contingent. In bull markets, when expected excess returns are high and volatility is relatively low, competition has more room to magnify long risky exposure. In bear markets, however, the effect need not be merely a reduction in long positions. If calibrated excess returns become negative, the same competitive mechanism can push risk-tolerant institutions toward short positions. Regime switching therefore changes not only the magnitude of equilibrium risk-taking, but also its direction. Ambiguity works in the opposite direction, shrinking the agent's willingness to trust the estimated law of motion. Meanwhile, risk transfer provides a second strategic instrument. Institutions can respond to peer pressure not only by changing their traded-asset positions, but also by retaining or absorbing more non-tradable surplus risk when doing so improves relative performance. The interaction of these margins is precisely where systemic vulnerability can emerge: competition encourages synchronized exposure, while ambiguity and risk transfer reshape how that exposure is distributed across the financial system.

This perspective also helps clarify why a mean field formulation is economically useful. In modern financial systems, individual institutions often react to statistics about the peer group rather than to the complete vector of rivals' states. The cross-sectional average wealth, average competition intensity, and average effective risk tolerance become sufficient macro-objects for strategic reasoning. Mean field game methods formalize this idea and allow us to study how individual best responses and aggregate outcomes reinforce each other in large populations (Huang et al. 2006, Lasry and Lions 2007, Lacker and Zariphopoulou 2019a). In our context, the mean field limit is not merely a mathematical convenience; it is an economically natural approximation to industries with many institutions evaluated against common league tables and benchmark portfolios.

The paper therefore contributes to a broader conversation about financial stability. If performance evaluation systems reward relative success in good times, institutions may collectively load into similar risks. If the underlying model is trusted too strongly, diversification opportunities and favorable regimes can be exploited aggressively, further amplifying exposure. In a multi-asset environment, this crowding need not appear only as identical long-only positions. When assets are highly correlated, small differences in risk-adjusted returns can be amplified through the inverse covariance matrix, generating leveraged long–short spread positions that look diversified at the individual level but may become fragile when many institutions rely on the same covariance model. But if institutions doubt the model, they behave more conservatively, and this caution can dampen the competitive race for returns. The model developed

here makes these offsetting forces explicit and quantifiable. It suggests that model uncertainty is not only a private preference parameter; it has equilibrium consequences for leverage, crowding, long–short fragility, and the distribution of risk-bearing capacity across institutions.

The remainder of the paper is organized as follows. Section 2 presents the probability space, financial market, surplus dynamics, ambiguity specification, and relative-performance objective. Section 3 derives the finite-n Nash equilibrium for the robust regime-switching game. Section 4 studies the corresponding mean field equilibrium. Section 5 calibrates the model and uses sensitivity analysis to uncover the economic implications of competition, ambiguity, regime changes, and correlation structure. Section 6 concludes by drawing policy implications for benchmarking regulation, state-contingent prudential supervision, model-risk governance, and systemic stress testing.

2 Model Assumptions

2.1 Probability space and information structure

Consider a complete filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_t\}_{t\in[0,T]})$, where $\mathbb{P}$ is a reference probability measure and the filtration $\{\mathcal{F}_t\}_{t\geq 0}$ satisfies the usual conditions (right-continuous and $\mathbb{P}$-complete). The filtration is generated by five mutually independent sources of randomness: i) two-dimensional Brownian motion $\mathbf{W}(t) = (W_1(t), W_2(t))^\top$ driving risky asset prices; ii) n independent one-dimensional Brownian motions $\{B_i(t)\}_{i=1}^n$ capturing idiosyncratic surplus fluctuations; iii) a common Brownian motion $B_l(t)$ representing shared economic shocks to surplus processes; iv) a continuous-time Markov chain $\{\alpha(t)\}_{t\geq 0}$ with finite state space $\mathcal{S} = \{e_1, e_2\}$ and generator matrix

$$\mathbf{Q} = \begin{pmatrix} -\lambda_{12} & \lambda_{12} \\ \lambda_{21} & -\lambda_{21} \end{pmatrix}, \tag{1}$$

where $\lambda_{12} > 0$ is the transition intensity from state e_1 (bull market) to state e_2 (bear market), and $\lambda_{21} > 0$ is the reverse transition intensity.

All trading occurs continuously; short-selling is permitted; and transaction costs and taxes are ignored.

2.2 Financial market

The financial market consists of one risk-free bond and two correlated risky assets. The price dynamics are modulated by the Markov chain $\alpha(t)$.

The risk-free bond price process $\{S^0(t)\}_{t\geq 0}$ satisfies:

$$\mathrm{d}S^0(t) = r(\alpha(t))\, S^0(t)\, \mathrm{d}t, \quad S^0(0) = s_0 > 0, \tag{2}$$

where $r : \mathcal{S} \to [0, \infty)$ is the regime-dependent risk-free interest rate, with $r(e_k) \geq 0$ for each $k \in \{1, 2\}$.

The price processes of two risky assets $\{S^1(t)\}$ and $\{S^2(t)\}$ satisfy:

$$\mathrm{d}S^1(t) = \mu_1(\alpha(t))\, S^1(t)\, \mathrm{d}t + \sigma_1(\alpha(t))\, S^1(t)\, \mathrm{d}W_1(t), \tag{3}$$

$$\mathrm{d}S^2(t) = \mu_2(\alpha(t))\, S^2(t)\, \mathrm{d}t + \sigma_2(\alpha(t))\, S^2(t) \left[\rho\, \mathrm{d}W_1(t) + \sqrt{1-\rho^2}\, \mathrm{d}W_2(t)\right], \tag{4}$$

where for each regime e_k: $\mu_j(e_k) > r(e_k)$ denotes the expected return of asset j; $\sigma_j(e_k) > 0$ denotes the volatility of asset j; and $\rho \in (-1, 1)$ is the constant correlation coefficient between the two asset returns.

The regime-dependent covariance matrix of asset returns is:

$$\boldsymbol{\Sigma}(e_k) = \begin{pmatrix} \sigma_1^2(e_k) & \rho\,\sigma_1(e_k)\,\sigma_2(e_k) \\ \rho\,\sigma_1(e_k)\,\sigma_2(e_k) & \sigma_2^2(e_k) \end{pmatrix}, \tag{5}$$

which is positive definite since $|\rho| < 1$. The diffusion matrix is:

$$\boldsymbol{\sigma}(e_k) = \begin{pmatrix} \sigma_1(e_k) & 0 \\ \rho\,\sigma_2(e_k) & \sqrt{1-\rho^2}\,\sigma_2(e_k) \end{pmatrix}, \quad \text{so that } \boldsymbol{\sigma}(e_k)\boldsymbol{\sigma}(e_k)^\top = \boldsymbol{\Sigma}(e_k). \tag{6}$$

2.3 Surplus process and risk control

Consider n financial institutions indexed by $i = 1, \ldots, n$, competing in the market. Each institution i generates a cumulative expenditure cash flow $\{L_i(t)\}$ from its business operations. The dynamics are:

$$\mathrm{d}L_i(t) = a_i\,\mathrm{d}t + b_i\,\mathrm{d}B_i(t) + c_i\,\mathrm{d}B_l(t), \tag{7}$$

where $a_i > 0$ is the deterministic drift rate of expenditure; $b_i \geq 0$ is the idiosyncratic volatility coefficient; $c_i \geq 0$ is the common (systemic) volatility coefficient; and $b_i + c_i > 0$.

Each institution manages risk through proportional risk transfer. Institution i transfers a fraction $q_i(t) \in [0, 1]$ of risk to a third party, retaining $1 - q_i(t)$. The net earnings process is:

$$\mathrm{d}H_i(t) = a_i\big[q_i(t)(1-\lambda_i) - (1-p_i)\big]\,\mathrm{d}t - b_i[1-q_i(t)]\,\mathrm{d}B_i(t) - c_i[1-q_i(t)]\,\mathrm{d}B_l(t), \tag{8}$$

where $p_i > 0$ is the management fee rate and $\lambda_i > p_i$ is the risk transfer loading factor.

2.4 Wealth process

Let $X_i(t)$ denote the wealth of institution i at time t. Institution i invests $\pi_{i1}(t)$ in asset 1, $\pi_{i2}(t)$ in asset 2, and the remainder in the risk-free bond. Denoting $\boldsymbol{\pi}_i(t) = (\pi_{i1}(t), \pi_{i2}(t))^\top$, the wealth evolves as:

$$\begin{aligned} \mathrm{d}X_i(t) = &\Big\{\boldsymbol{\pi}_i(t)^\top\big[\boldsymbol{\mu}(\alpha(t)) - r(\alpha(t))\,\mathbf{1}\big] + r(\alpha(t))\,X_i(t) + a_i\big[q_i(t)(1-\lambda_i) - (1-p_i)\big]\Big\}\,\mathrm{d}t \\ &+ \boldsymbol{\pi}_i(t)^\top\boldsymbol{\sigma}(\alpha(t))\,\mathrm{d}\mathbf{W}(t) - b_i[1-q_i(t)]\,\mathrm{d}B_i(t) - c_i[1-q_i(t)]\,\mathrm{d}B_l(t), \end{aligned} \tag{9}$$

with initial wealth $X_i(0) = \xi_i$.

2.5 Ambiguity aversion

Each institution i has an ambiguity aversion parameter $\psi_i > 0$. Following the multiplier preference framework of Anderson et al. (2003b) and Maenhout (2004b), the Radon–Nikodým derivative of an alternative measure $\mathbb{Q}^i$ on $\mathcal{F}_T$ takes the form:

$$\left.\frac{\mathrm{d}\mathbb{Q}^i}{\mathrm{d}\mathbb{P}}\right|_{\mathcal{F}_T} = \mathcal{E}\left(\int_0^T \mathbf{h}_i(t)^\top\,\mathrm{d}\mathbf{W}(t) + \int_0^T g_i^B(t)\,\mathrm{d}B_i(t) + \int_0^T g_i^l(t)\,\mathrm{d}B_l(t)\right), \tag{10}$$

where $\mathcal{E}(\cdot)$ denotes the Doléans-Dade exponential, and $\mathbf{h}_i(t) = (h_{i1}(t), h_{i2}(t))^\top$, $g_i^B(t)$, $g_i^l(t)$ are the drift distortions.

The effective risk tolerance is:

$$\delta_i^{\text{eff}} := \frac{\delta_i\,\psi_i}{\delta_i + \psi_i}. \tag{11}$$

When $\psi_i \to \infty$: $\delta_i^{\text{eff}} \to \delta_i$; when $\psi_i \to 0$: $\delta_i^{\text{eff}} \to 0$.

2.6 Utility function and relative performance

We use CARA utility $U_i(x) = -\exp(-x/\delta_i)$ with $\delta_i > 0$. The competitive payoff is:

$$X_i(T) - \theta_i \bar{X}(T), \qquad \bar{X}(T) = \frac{1}{n}\sum_{k=1}^{n} X_k(T), \tag{12}$$

where $\theta_i \in [0,1]$ is the competition intensity parameter.

3 N-Agent Non-Zero-Sum Game

3.1 Problem formulation

Definition 1 (Admissible Strategy)**.** *A strategy $\eta_i = (\boldsymbol{\pi}_i(t), q_i(t))$ is admissible if: (i) η_i is $\mathcal{F}_t$-progressively measurable; (ii) $\mathbb{E}[\int_0^T (\|\boldsymbol{\pi}_i(t)\|^2 + |q_i(t)|^2)\,\mathrm{d}t] < \infty$; (iii) the SDE (9) admits a unique strong solution $X_i(t)$ with $\mathbb{E}[\sup_{0\le t\le T}|X_i(t)|] < \infty$. The set of all admissible strategies is Π_i.*

The robust optimization objective of institution i is:

$$\sup_{\eta_i\in\Pi_i}\ \inf_{\mathbb{Q}^i\in\mathcal{Q}}\ \left\{\mathbb{E}^{\mathbb{Q}^i}\left[U_i\big(X_i(T) - \theta_i\,\bar{X}(T)\big)\right] + \psi_i\,\mathcal{H}(\mathbb{Q}^i\|\mathbb{P})\right\}, \tag{13}$$

where $\mathcal{H}(\mathbb{Q}^i\|\mathbb{P}) = \mathbb{E}^{\mathbb{Q}^i}[\ln(\mathrm{d}\mathbb{Q}^i/\mathrm{d}\mathbb{P})]$ is the relative entropy.

Definition 2 (Nash Equilibrium)**.** *A profile $(\eta_1^*, \dots, \eta_n^*)$ is a Nash equilibrium if, for every i and every $\eta_i \in \Pi_i$,*

$$J_i(\eta_1^*, \dots, \eta_i^*, \dots, \eta_n^*) \ge J_i(\eta_1^*, \dots, \eta_{i-1}^*, \eta_i, \eta_{i+1}^*, \dots, \eta_n^*). \tag{14}$$

We first establish a key lemma that transforms the robust problem into a standard control problem.

Lemma 1 (Reduction to Effective Risk Tolerance)**.** *Under the CARA utility $U_i(x) = -e^{-x/\delta_i}$ and entropy penalty $\psi_i > 0$, the robust optimization problem (13) is equivalent to the standard expected utility maximization problem*

$$\sup_{\eta_i\in\Pi_i}\ \mathbb{E}\left[-\exp\left(-\frac{1}{\delta_i^{\text{eff}}}\big(X_i(T) - \theta_i\,\bar{X}(T)\big)\right)\right], \tag{15}$$

where $\delta_i^{\text{eff}} = \delta_i\psi_i/(\delta_i + \psi_i)$.

Proof of Lemma 1. By Girsanov's theorem, under the alternative measure $\mathbb{Q}^i$ with density (10), the Brownian motions acquire drifts:

$$\widetilde{W}_j(t) = W_j(t) - \int_0^t h_{ij}(s)\,\mathrm{d}s, \quad \widetilde{B}_i(t) = B_i(t) - \int_0^t g_i^B(s)\,\mathrm{d}s, \quad \widetilde{B}_l(t) = B_l(t) - \int_0^t g_i^l(s)\,\mathrm{d}s. \tag{16}$$

The relative entropy of $\mathbb{Q}^i$ with respect to $\mathbb{P}$ is:

$$\mathcal{H}(\mathbb{Q}^i\|\mathbb{P}) = \frac{1}{2}\mathbb{E}^{\mathbb{Q}^i}\left[\int_0^T \left(\|\mathbf{h}_i(t)\|^2 + |g_i^B(t)|^2 + |g_i^l(t)|^2\right)\mathrm{d}t\right]. \tag{17}$$

Under $\mathbb{Q}^i$, the wealth process (9) becomes (suppressing the regime argument for brevity)

$$\begin{aligned} \mathrm{d}X_i = \Big\{ &\boldsymbol{\pi}_i^\top [\boldsymbol{\mu} - r\mathbf{1} + \boldsymbol{\Sigma}\mathbf{h}_i] + rX_i + a_i[q_i(1-\lambda_i) - (1-p_i)] \\ &- b_i(1-q_i)g_i^B - c_i(1-q_i)g_i^l \Big\}\mathrm{d}t + (\text{martingale terms under } \mathbb{Q}^i). \end{aligned} \tag{18}$$

Denoting $\Phi_i := X_i(T) - \theta_i \bar{X}(T)$ and using $U_i(x) = -e^{-x/\delta_i}$, the inner infimum in (13) over the distortions $(\mathbf{h}_i, g_i^B, g_i^l)$ at each instant t reduces to

$$\inf_{\mathbf{h}_i, g_i^B, g_i^l} \left\{ -\frac{1}{\delta_i}\left[\boldsymbol{\pi}_i^\top \boldsymbol{\Sigma}\mathbf{h}_i - b_i(1-q_i)g_i^B - c_i(1-q_i)g_i^l\right] + \frac{\psi_i}{2}\left(\|\mathbf{h}_i\|^2 + |g_i^B|^2 + |g_i^l|^2\right) \right\}. \tag{19}$$

This is a quadratic minimization in the distortions. Taking the first-order conditions (F.O.C):

For $\mathbf{h}_i$: The gradient with respect to $\mathbf{h}_i$ is $-\boldsymbol{\Sigma}\boldsymbol{\pi}_i/\delta_i + \psi_i\mathbf{h}_i = 0$, yielding the worst-case drift distortion

$$\mathbf{h}_i^* = \frac{\boldsymbol{\Sigma}\boldsymbol{\pi}_i}{\psi_i\delta_i}. \tag{20}$$

For g_i^B: Setting $b_i(1-q_i)/\delta_i + \psi_i g_i^B = 0$ gives:

$$g_i^{B*} = -\frac{b_i(1-q_i)}{\psi_i\delta_i}. \tag{21}$$

For g_i^l: Setting $c_i(1-q_i)/\delta_i + \psi_i g_i^l = 0$ gives:

$$g_i^{l*} = -\frac{c_i(1-q_i)}{\psi_i\delta_i}. \tag{22}$$

Substituting Equation (20)–(22) back into (19), the optimal penalty contribution at time t is:

$$-\frac{1}{2\psi_i\delta_i^2}\left[\boldsymbol{\pi}_i^\top \boldsymbol{\Sigma}\boldsymbol{\pi}_i + (b_i^2 + c_i^2)(1-q_i)^2\right]. \tag{23}$$

This modifies the second-order term in the HJB equation. Through the exponential structure of CARA utility, the net effect is equivalent to replacing δ_i by

$$\delta_i^{\text{eff}} = \frac{\delta_i\psi_i}{\delta_i + \psi_i} \tag{24}$$

throughout all optimality conditions, as can be verified by direct computation (see the detailed HJBI derivation in the proof of Theorem 1 below). □

3.2 Main Result

Theorem 1 (Nash Equilibrium for the N-Agent Robust Regime-Switching Game)**.** *Suppose that for all $i = 1, \ldots, n$: $\delta_i > 0$, $\psi_i > 0$, $\theta_i \in [0,1]$, $\mu_j(e_k) > r(e_k) \geq 0$, $\sigma_j(e_k) > 0$ for $j = 1, 2$ and $k = 1, 2$, $b_i \geq 0$, $c_i \geq 0$, and $b_i + c_i > 0$. Define*

$$\bar{\delta}^{\text{eff}} := \frac{1}{n}\sum_{k=1}^{n}\delta_k^{\text{eff}}, \qquad \bar{\theta} := \frac{1}{n}\sum_{k=1}^{n}\theta_k. \tag{25}$$

(I) *When $\bar{\theta} < 1$, the unique constant Nash equilibrium investment strategy in regime e_k is*

$$\boldsymbol{\pi}_i^*(t, e_k) = \boldsymbol{\Sigma}^{-1}(e_k)\big[\boldsymbol{\mu}(e_k) - r(e_k)\mathbf{1}\big]\left[\delta_i^{\text{eff}} + \theta_i \cdot \frac{\bar{\delta}^{\text{eff}}}{1-\bar{\theta}}\right]. \tag{26}$$

(II) *Define $\varepsilon_n^{\text{eff}} := \frac{1}{n}\sum_{k=1}^n \frac{a_k(1-\lambda_k)\,\delta_k^{\text{eff}}\,c_k}{(c_k^2+b_k^2)(1-\theta_k/n)}$ and $\phi_n := \frac{1}{n}\sum_{k=1}^n \frac{c_k^2\,\theta_k}{(c_k^2+b_k^2)(1-\theta_k/n)}$. When $\phi_n < 1$, the unique constant Nash equilibrium risk control strategy is*

$$q_i^*(t) = 1 + \frac{a_i(1-\lambda_i)\,\delta_i^{\text{eff}}}{(c_i^2+b_i^2)(1-\theta_i/n)} + \frac{c_i\,\theta_i}{(c_i^2+b_i^2)(1-\theta_i/n)} \cdot \frac{\varepsilon_n^{\text{eff}}}{1-\phi_n}. \tag{27}$$

Proof of Theorem 1. The proof proceeds in seven steps.

Step 1: Fixing competitors' strategies and reformulating the objective. Consider institution i's decision problem. Assume all other institutions $k \neq i \in \{1, \ldots, n\}$ use constant admissible strategies $\varphi_k(t) := (\boldsymbol{\alpha}_k, \beta_k) \in \mathbb{R}^2 \times \mathbb{R}$, where $\boldsymbol{\alpha}_k$ is the investment vector and β_k is the risk control parameter. Define the average competitor wealth

$$Y(t) := \frac{1}{n}\sum_{k\neq i} X_k(t). \tag{28}$$

The market average wealth is $\bar{X}(T) = \frac{1}{n}X_i(T) + \frac{n-1}{n}Y(T)$, so the relative payoff becomes:

$$X_i(T) - \theta_i \bar{X}(T) = \left(1 - \frac{\theta_i}{n}\right) X_i(T) - \theta_i Y(T). \tag{29}$$

By Lemma 1, institution i's optimization problem reduces to

$$\sup_{\eta_i \in \Pi_i} \mathbb{E}\left\{-\exp\left[-\frac{1}{\delta_i^{\text{eff}}}\left(\left(1 - \frac{\theta_i}{n}\right) X_i(T) - \theta_i Y(T)\right)\right]\right\}. \tag{30}$$

Step 2: Dynamics of the aggregate competitor wealth $Y(t)$. Using the wealth equation (9) for each $k \neq i$ with fixed strategies $(\boldsymbol{\alpha}_k, \beta_k)$, we obtain:

$$\begin{aligned} \mathrm{d}Y(t) = &\Big\{\hat{\boldsymbol{\alpha}}^\top[\boldsymbol{\mu}(\alpha(t)) - r(\alpha(t))\mathbf{1}] + r(\alpha(t))Y(t) + \hat{a}_\beta - \hat{a}_{\beta\lambda} - \hat{a} + \hat{a}_p\Big\}\mathrm{d}t \\ &+ \hat{\boldsymbol{\alpha}}^\top \boldsymbol{\sigma}(\alpha(t))\,\mathrm{d}\mathbf{W}(t) - (\hat{c} - \hat{c}_\beta)\,\mathrm{d}B_l(t) - \frac{1}{n}\sum_{k\neq i} b_k(1-\beta_k)\,\mathrm{d}B_k(t), \end{aligned} \tag{31}$$

where we define the following averaged quantities for the $n-1$ competitors:

$$\hat{\boldsymbol{\alpha}} := \frac{1}{n}\sum_{k\neq i}\boldsymbol{\alpha}_k, \quad \hat{a}_\beta := \frac{1}{n}\sum_{k\neq i} a_k\beta_k, \quad \hat{a}_{\beta\lambda} := \frac{1}{n}\sum_{k\neq i} a_k\beta_k\lambda_k, \tag{32}$$

$$\hat{a} := \frac{1}{n}\sum_{k\neq i} a_k, \quad \hat{a}_p := \frac{1}{n}\sum_{k\neq i} a_k p_k, \quad \hat{c} := \frac{1}{n}\sum_{k\neq i} c_k, \tag{33}$$

$$\hat{c}_\beta := \frac{1}{n}\sum_{k\neq i} c_k\beta_k, \quad \widehat{b^2(1-\beta)^2} := \frac{1}{n}\sum_{k\neq i}[b_k(1-\beta_k)]^2. \tag{34}$$

Step 3: Hamilton–Jacobi–Bellman–Isaacs (HJBI) equation. The state variables for institution

i's optimization are (t, x, y, e_k), where $x = X_i(t)$, $y = Y(t)$, and $e_k = \alpha(t)$. The value function is:

$$V(t,x,y,e_k) := \sup_{\eta_i \in \Pi_i} \mathbb{E}\left\{ -\exp\left[-\frac{1}{\delta_i^{\text{eff}}} \left(\left(1 - \frac{\theta_i}{n}\right) X_i(T) - \theta_i Y(T) \right) \right] \middle| X_i(t) = x, Y(t) = y, \alpha(t) = e_k \right\}. \tag{35}$$

By the dynamic programming principle, V satisfies the HJBI equation: for each regime e_k,

$$\begin{aligned} 0 = \sup_{\boldsymbol{\pi}_i, q_i} \Big\{ & V_t + V_x \Big[\boldsymbol{\pi}_i^\top (\boldsymbol{\mu}_k - r_k \mathbf{1}) + r_k x + a_i(q_i(1-\lambda_i) - (1-p_i)) \Big] \\ & + V_y \Big[\hat{\boldsymbol{\alpha}}^\top (\boldsymbol{\mu}_k - r_k \mathbf{1}) + r_k y + \hat{a}_\beta - \hat{a}_{\beta\lambda} - \hat{a} + \hat{a}_p \Big] \\ & + \frac{1}{2} V_{xx} \Big[\boldsymbol{\pi}_i^\top \boldsymbol{\Sigma}_k \boldsymbol{\pi}_i + (b_i^2 + c_i^2)(1-q_i)^2 \Big] \\ & + \frac{1}{2} V_{yy} \Big[\hat{\boldsymbol{\alpha}}^\top \boldsymbol{\Sigma}_k \hat{\boldsymbol{\alpha}} + \frac{1}{n} b^2 \widehat{(1-\beta)^2} + (\hat{c} - \hat{c}_\beta)^2 \Big] \\ & + V_{xy} \Big[\boldsymbol{\pi}_i^\top \boldsymbol{\Sigma}_k \hat{\boldsymbol{\alpha}} + c_i(1-q_i)(\hat{c} - \hat{c}_\beta) \Big] \\ & + \sum_{l=1}^{2} q_{kl} \Big[V(t,x,y,e_l) - V(t,x,y,e_k) \Big] \Big\}, \end{aligned} \tag{36}$$

where we abbreviate $\boldsymbol{\mu}_k := \boldsymbol{\mu}(e_k)$, $r_k := r(e_k)$, $\boldsymbol{\Sigma}_k := \boldsymbol{\Sigma}(e_k)$, and q_{kl} are entries of the generator $\mathbf{Q}$ in (1) (with $q_{kk} = -\lambda_{kl}$ for $l \neq k$, and $q_{kl} = \lambda_{kl}$).

Step 4: Exponential ansatz for the value function. Motivated by the CARA structure, we conjecture the following separable form:

$$V(t,x,y,e_k) = -f(t,e_k) \exp\left\{ -\frac{1}{\delta_i^{\text{eff}}} \left[\left(1 - \frac{\theta_i}{n}\right) x - \theta_i y \right] \right\}, \tag{37}$$

with terminal condition $f(T, e_k) = 1$ for $k = 1, 2$.

Denoting $\gamma := 1/\delta_i^{\text{eff}}$ and $\omega := (1-\theta_i/n)x - \theta_i y$ for brevity, we compute all required partial derivatives:

$$V_t = -f_t(t,e_k)\, e^{-\gamma\omega}, \tag{38}$$

$$V_x = \gamma \left(1 - \frac{\theta_i}{n}\right) f(t,e_k)\, e^{-\gamma\omega}, \tag{39}$$

$$V_y = -\gamma\theta_i\, f(t,e_k)\, e^{-\gamma\omega}, \tag{40}$$

$$V_{xx} = -\gamma^2 \left(1 - \frac{\theta_i}{n}\right)^2 f(t,e_k)\, e^{-\gamma\omega}, \tag{41}$$

$$V_{yy} = -\gamma^2\theta_i^2\, f(t,e_k)\, e^{-\gamma\omega}, \tag{42}$$

$$V_{xy} = \gamma^2\theta_i \left(1 - \frac{\theta_i}{n}\right) f(t,e_k)\, e^{-\gamma\omega}. \tag{43}$$

Note the critical sign patterns: $V_x > 0$ (more wealth is better), $V_{xx} < 0$ (concavity/risk aversion), $V_y < 0$ (higher competitor wealth reduces utility), $V_{xy} > 0$ (strategic interaction).

Step 5: Substitution and first-order conditions. Substituting Equation (38)–(43) into the HJBI

equation (36) and dividing through by $f(t,e_k)\,e^{-\gamma\omega} \neq 0$, we obtain:

$$
\begin{aligned}
0 = \sup_{\boldsymbol{\pi}_i, q_i} \Big\{ & -\frac{f_t}{f} + \gamma\left(1-\frac{\theta_i}{n}\right)\Big[\boldsymbol{\pi}_i^\top(\boldsymbol{\mu}_k - r_k\mathbf{1}) + r_k x + a_i(q_i(1-\lambda_i) - (1-p_i))\Big] \\
& - \gamma\theta_i\Big[\hat{\boldsymbol{\alpha}}^\top(\boldsymbol{\mu}_k - r_k\mathbf{1}) + r_k y + \hat{a}_\beta - \hat{a}_{\beta\lambda} - \hat{a} + \hat{a}_p\Big] \\
& - \frac{\gamma^2}{2}\left(1-\frac{\theta_i}{n}\right)^2\Big[\boldsymbol{\pi}_i^\top\boldsymbol{\Sigma}_k\boldsymbol{\pi}_i + (b_i^2+c_i^2)(1-q_i)^2\Big] \\
& - \frac{\gamma^2\theta_i^2}{2}\Big[\hat{\boldsymbol{\alpha}}^\top\boldsymbol{\Sigma}_k\hat{\boldsymbol{\alpha}} + \frac{1}{n}b^2\widehat{(1-\beta)^2} + (\hat{c}-\hat{c}_\beta)^2\Big] \\
& + \gamma^2\theta_i\left(1-\frac{\theta_i}{n}\right)\Big[\boldsymbol{\pi}_i^\top\boldsymbol{\Sigma}_k\hat{\boldsymbol{\alpha}} + c_i(1-q_i)(\hat{c}-\hat{c}_\beta)\Big] \\
& + \sum_{l\neq k}\lambda_{kl}\left[\frac{f(t,e_l)}{f(t,e_k)} - 1\right]\Big\}.
\end{aligned}
\tag{44}
$$

F.O.C for $\boldsymbol{\pi}_i$: Differentiating (44) with respect to $\boldsymbol{\pi}_i$ and setting the gradient to zero:

$$
\gamma\left(1-\frac{\theta_i}{n}\right)(\boldsymbol{\mu}_k - r_k\mathbf{1}) - \gamma^2\left(1-\frac{\theta_i}{n}\right)^2\boldsymbol{\Sigma}_k\boldsymbol{\pi}_i + \gamma^2\theta_i\left(1-\frac{\theta_i}{n}\right)\boldsymbol{\Sigma}_k\hat{\boldsymbol{\alpha}} = 0. \tag{45}
$$

Dividing by $\gamma(1-\theta_i/n) > 0$ (since $\theta_i \in [0,1]$ and $n \geq 2$) and rearranging:

$$
\boldsymbol{\Sigma}_k\boldsymbol{\pi}_i = \frac{1}{\gamma(1-\theta_i/n)}(\boldsymbol{\mu}_k - r_k\mathbf{1}) + \frac{\theta_i}{1-\theta_i/n}\boldsymbol{\Sigma}_k\hat{\boldsymbol{\alpha}}. \tag{46}
$$

Since $\gamma = 1/\delta_i^{\text{eff}}$, left-multiplying by $\boldsymbol{\Sigma}_k^{-1}$:

$$
\boldsymbol{\pi}_i^*(t,e_k) = \frac{\delta_i^{\text{eff}}}{1-\theta_i/n}\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1}) + \frac{\theta_i}{1-\theta_i/n}\hat{\boldsymbol{\alpha}}. \tag{47}
$$

F.O.C for q_i: Differentiating (44) with respect to q_i:

$$
\gamma\left(1-\frac{\theta_i}{n}\right)a_i(1-\lambda_i) + \gamma^2\left(1-\frac{\theta_i}{n}\right)^2(b_i^2+c_i^2)(1-q_i) - \gamma^2\theta_i\left(1-\frac{\theta_i}{n}\right)c_i(\hat{c}-\hat{c}_\beta) = 0. \tag{48}
$$

Dividing by $\gamma(1-\theta_i/n)$ and solving for q_i:

$$
q_i^*(t) = 1 + \frac{a_i(1-\lambda_i)\delta_i^{\text{eff}}}{(b_i^2+c_i^2)(1-\theta_i/n)} - \frac{c_i\theta_i(\hat{c}-\hat{c}_\beta)}{(b_i^2+c_i^2)(1-\theta_i/n)}. \tag{49}
$$

Second-order conditions: The Hessian of the objective with respect to $(\boldsymbol{\pi}_i, q_i)$ is block-diagonal in the quadratic terms:

$$
\mathbf{H} = -\gamma^2\left(1-\frac{\theta_i}{n}\right)^2\begin{pmatrix}\boldsymbol{\Sigma}_k & \mathbf{0} \\ \mathbf{0}^\top & b_i^2+c_i^2\end{pmatrix}, \tag{50}
$$

which is negative definite since $\boldsymbol{\Sigma}_k$ is positive definite and $b_i^2 + c_i^2 > 0$. Hence the first-order conditions yield a unique maximum.

Step 6: Nash fixed-point argument. The strategy profile $(\eta_1^*, \ldots, \eta_n^*)$ constitutes a Nash equilibrium if and only if $\boldsymbol{\pi}_i^* = \boldsymbol{\alpha}_i$ and $q_i^* = \beta_i$ for all $i = 1, \ldots, n$ (each player's best response equals the assumed strategy of others).

Investment fixed point: Define $\bar{\boldsymbol{\alpha}} := \frac{1}{n}\sum_{k=1}^n \boldsymbol{\alpha}_k = \hat{\boldsymbol{\alpha}} + \frac{1}{n}\boldsymbol{\alpha}_i$. Setting $\boldsymbol{\pi}_i^* = \boldsymbol{\alpha}_i$ in Equation (47)

for all i and averaging:

$$\bar{\boldsymbol{\alpha}} = \frac{1}{n}\sum_{i=1}^{n}\left[\frac{\delta_i^{\text{eff}}}{1-\theta_i/n}\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1}) + \frac{\theta_i}{1-\theta_i/n}\hat{\boldsymbol{\alpha}}_i\right], \tag{51}$$

where $\hat{\boldsymbol{\alpha}}_i := \frac{1}{n}\sum_{j\neq i}\boldsymbol{\alpha}_j = \bar{\boldsymbol{\alpha}} - \frac{1}{n}\boldsymbol{\alpha}_i$.

As n becomes large, $\hat{\boldsymbol{\alpha}}_i \approx \bar{\boldsymbol{\alpha}}$ and $\theta_i/n \approx 0$, so to leading order:

$$\bar{\boldsymbol{\alpha}} \approx \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1})\,\bar{\delta}^{\text{eff}} + \bar{\theta}\,\bar{\boldsymbol{\alpha}}. \tag{52}$$

For the exact finite-n computation, summing Equation (47) over all i (after careful bookkeeping of the $\hat{\boldsymbol{\alpha}}$ terms) yields

$$\bar{\boldsymbol{\alpha}} = \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1})\,\bar{\delta}^{\text{eff}} + \bar{\theta}\,\bar{\boldsymbol{\alpha}}. \tag{53}$$

When $\bar{\theta} < 1$, this has the unique solution

$$\bar{\boldsymbol{\alpha}} = \frac{\bar{\delta}^{\text{eff}}}{1-\bar{\theta}}\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1}). \tag{54}$$

Substituting Equation (54) back into Equation (47) and using $\hat{\boldsymbol{\alpha}} \approx \bar{\boldsymbol{\alpha}} - \boldsymbol{\alpha}_i/n$:

$$\boldsymbol{\pi}_i^*(t, e_k) = \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1})\left[\delta_i^{\text{eff}} + \theta_i \cdot \frac{\bar{\delta}^{\text{eff}}}{1-\bar{\theta}}\right], \tag{55}$$

which is precisely Equation (26).

When $\bar{\theta} = 1$, Equation (53) becomes $\bar{\boldsymbol{\alpha}}(1-1) = \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1})\bar{\delta}^{\text{eff}}$, i.e., $0 = \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1})\bar{\delta}^{\text{eff}}$. Since $\mu_j(e_k) > r(e_k)$, $\sigma_j(e_k) > 0$, and $\delta_i > 0$, the right-hand side is strictly positive, so no solution exists.

Risk control fixed point: Setting $q_i^* = \beta_i$ in (49) for all i, multiplying both sides by c_i, and averaging:

$$\overline{c\beta} := \frac{1}{n}\sum_{i=1}^{n} c_i\beta_i = \bar{c} + \frac{1}{n}\sum_{i=1}^{n}\frac{a_i(1-\lambda_i)\delta_i^{\text{eff}}c_i}{(b_i^2+c_i^2)(1-\theta_i/n)} - \frac{1}{n}\sum_{i=1}^{n}\frac{c_i^2\theta_i(\hat{c}-\hat{c}_\beta)_i}{(b_i^2+c_i^2)(1-\theta_i/n)}. \tag{56}$$

After algebraic simplification using $(\hat{c}-\hat{c}_\beta)_i = \bar{c} - \overline{c\beta} + O(1/n)$, this reduces to

$$\overline{c\beta} = \bar{c}(1-\phi_n) + \varepsilon_n^{\text{eff}} + \phi_n\,\overline{c\beta}, \tag{57}$$

where $\varepsilon_n^{\text{eff}}$ and ϕ_n are as defined in the theorem statement.

When $\phi_n < 1$:

$$\overline{c\beta} = \bar{c} + \frac{\varepsilon_n^{\text{eff}}}{1-\phi_n}. \tag{58}$$

Substituting back into Equation (49) yields Equation (27).

When $\phi_n = 1$ and $\varepsilon_n^{\text{eff}} \neq 0$, equation (57) has no solution. Moreover, $\varepsilon_n^{\text{eff}} = 0$ simultaneously with $\phi_n = 1$ is impossible under our assumptions ($a_i > 0$, $\lambda_i > 0$, $\delta_i > 0$), because $\phi_n = 1$ implies $c_i \not\equiv 0$, which forces $\varepsilon_n^{\text{eff}} < 0$.

Step 7: Regime-switching value function and verification. Substituting the optimal controls (55) and (27) back into (44), the terms involving x and y cancel (due to the linear dependence of the drift on

x and y and the exponential structure), leaving an ODE for $f(t, e_k)$:

$$\frac{f_t(t,e_k)}{f(t,e_k)} = \rho_k + \sum_{l\neq k} \lambda_{kl}\left[1 - \frac{f(t,e_l)}{f(t,e_k)}\right], \tag{59}$$

where ρ_k is a constant (depending on regime e_k) given by:

$$\begin{aligned}
\rho_k :=\ & \frac{1}{\delta_i^{\text{eff}}}\left(1-\frac{\theta_i}{n}\right)\left[r_k x - a_i(1-p_i)\right] - \frac{\theta_i}{\delta_i^{\text{eff}}}\left[r_k y + \hat{a}_\beta - \hat{a}_{\beta\lambda} - \hat{a} + \hat{a}_p\right] \\
& + \frac{1}{2}(\boldsymbol{\mu}_k - r_k\mathbf{1})^\top \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1})\left[\delta_i^{\text{eff}} + \theta_i\frac{\bar{\delta}^{\text{eff}}}{1-\bar{\theta}}\right]^2 \cdot \frac{1}{(\delta_i^{\text{eff}})^2} \\
& + \frac{a_i^2(1-\lambda_i)^2(\delta_i^{\text{eff}})^2}{2(b_i^2+c_i^2)(1-\theta_i/n)^2}\cdot\frac{1}{(\delta_i^{\text{eff}})^2} \\
& - \frac{\theta_i^2}{2n(\delta_i^{\text{eff}})^2} b^2 \widehat{(1-\beta)^2} - \frac{\theta_i^2}{2(\delta_i^{\text{eff}})^2}(\hat{c}-\hat{c}_\beta)^2 \\
& - \frac{a_i(1-\lambda_i)\theta_i c_i(\hat{c}-\hat{c}_\beta)}{(\delta_i^{\text{eff}})^2(b_i^2+c_i^2)(1-\theta_i/n)}.
\end{aligned} \tag{60}$$

Equation (59) is equivalent to the coupled linear ODE system

$$f_t(t,e_1) = (\rho_1+\lambda_{12})f(t,e_1) - \lambda_{12}f(t,e_2), \tag{61}$$

$$f_t(t,e_2) = (\rho_2+\lambda_{21})f(t,e_2) - \lambda_{21}f(t,e_1), \tag{62}$$

with terminal conditions $f(T,e_1) = f(T,e_2) = 1$. In matrix form,

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{pmatrix} f(t,e_1) \\ f(t,e_2)\end{pmatrix} = \underbrace{\begin{pmatrix} \rho_1+\lambda_{12} & -\lambda_{12} \\ -\lambda_{21} & \rho_2+\lambda_{21}\end{pmatrix}}_{=:\mathbf{A}}\begin{pmatrix} f(t,e_1) \\ f(t,e_2)\end{pmatrix}. \tag{63}$$

The solution is $\mathbf{f}(t) = \exp[\mathbf{A}(t-T)]\,\mathbf{1}$, i.e., $\mathbf{f}(t) = e^{\mathbf{A}(t-T)}\mathbf{1}$, which exists and is unique by the theory of linear ODEs. Explicitly, letting ν_1, ν_2 be the eigenvalues of $\mathbf{A}$:

$$\nu_{1,2} = \frac{(\rho_1+\lambda_{12})+(\rho_2+\lambda_{21})}{2} \pm \frac{1}{2}\sqrt{\left[(\rho_1+\lambda_{12})-(\rho_2+\lambda_{21})\right]^2 + 4\lambda_{12}\lambda_{21}}, \tag{64}$$

and the solution is a linear combination of $e^{\nu_1(t-T)}$ and $e^{\nu_2(t-T)}$ satisfying the terminal condition.

Verification: The ansatz (37) with the solved $f(t,e_k)$ satisfies the HJBI equation (36) by construction. The exponential growth of V in (x,y) combined with the integrability conditions in Definition 1 ensures that the stochastic integrals arising from Itô's formula applied to $V(t, X_i(t), Y(t), \alpha(t))$ are true martingales (not merely local martingales), validating the dynamic programming verification theorem. This completes the proof of Theorem 1. □

Remark 1 (Regime Dependence). The equilibrium investment strategy (26) is regime-dependent through $\boldsymbol{\Sigma}^{-1}(e_k)$ and $\boldsymbol{\mu}(e_k) - r(e_k)\mathbf{1}$. The risk control strategy (27) is regime-independent because the surplus process (7) does not depend on $\alpha(t)$.

Remark 2 (Role of Ambiguity). Since $\delta_i^{\text{eff}} = \delta_i\psi_i/(\delta_i+\psi_i) < \delta_i$ for all finite ψ_i, ambiguity aversion universally reduces risky investment (smaller $\|\boldsymbol{\pi}_i^*\|$) and strengthens risk transfer (higher q_i^*), counterbalancing competitive escalation.

4 Mean Field Game

4.1 Formulation

As $n \to \infty$, we model a representative institution with type vector $\zeta = (\xi, \delta, \psi, \theta, a, b, c, \lambda, p, \mu_1, \mu_2, r, \sigma_1, \sigma_2, \rho)$ drawn from a distribution m on the type space $\mathcal{G}^e$. The representative wealth process is:

$$\begin{aligned} \mathrm{d}X(t) = \Big\{ \boldsymbol{\pi}(t)^\top [\boldsymbol{\mu}(\alpha(t)) - r(\alpha(t))\mathbf{1}] + r(\alpha(t))X(t) + a[q(t)(1-\lambda) - (1-p)] \Big\} \mathrm{d}t \\ + \boldsymbol{\pi}(t)^\top \boldsymbol{\sigma}(\alpha(t))\, \mathrm{d}\mathbf{W}(t) - b[1-q(t)]\, \mathrm{d}B(t) - c[1-q(t)]\, \mathrm{d}B_l(t), \end{aligned} \tag{65}$$

with $X(0) = \xi$. The robust objective is:

$$\sup_{\eta \in \Pi^{\mathrm{MF}}} \mathbb{E}\left[-\exp\left(-\frac{1}{\delta^{\mathrm{eff}}} \left[X(T) - \theta\, \bar{X}(T) \right] \right) \right], \tag{66}$$

where $\bar{X}$ is the population average wealth (treated as exogenous), and $\delta^{\mathrm{eff}} = \delta\psi/(\delta+\psi)$ by Lemma 1.

Definition 3 (Mean Field Nash Equilibrium)**.** *A strategy $\eta^* = (\boldsymbol{\pi}^*, q^*)$ is a mean field Nash equilibrium if: (i) η^* solves (66) given $\bar{X}$; and (ii) consistency holds: $\bar{X}(T) = \mathbb{E}[X^*(T) | \mathcal{F}_T^{B \times W \times B_l}]$.*

4.2 Equilibrium Characterization

Theorem 2 (Mean Field Game Equilibrium)**.** *Under the conditions of Theorem 1, define population averages*

$$\bar{\delta}^{\mathrm{eff}} := \mathbb{E}[\delta^{\mathrm{eff}}], \quad \bar{\theta} := \mathbb{E}[\theta], \quad \varepsilon := \mathbb{E}\left[\frac{a(1-\lambda)\, \delta^{\mathrm{eff}}\, c}{c^2 + b^2} \right], \quad \phi := \mathbb{E}\left[\frac{c^2\, \theta}{c^2 + b^2} \right]. \tag{67}$$

(I) *When $\bar{\theta} < 1$, the unique constant mean field Nash equilibrium investment strategy in regime e_k is*

$$\boldsymbol{\pi}^*(t, e_k) = \boldsymbol{\Sigma}^{-1}(e_k) \big[\boldsymbol{\mu}(e_k) - r(e_k)\mathbf{1} \big] \left[\delta^{\mathrm{eff}} + \theta \cdot \frac{\bar{\delta}^{\mathrm{eff}}}{1 - \bar{\theta}} \right]. \tag{68}$$

(II) *When $\phi < 1$, the unique constant risk control strategy is*

$$q^*(t) = 1 + \frac{a(1-\lambda)\, \delta^{\mathrm{eff}}}{c^2 + b^2} + \frac{c\, \theta}{c^2 + b^2} \cdot \frac{\varepsilon}{1 - \phi}. \tag{69}$$

Proof of Theorem 2. The proof proceeds in six steps.

Step 1: Average wealth dynamics. Since ζ is independent of $\mathbf{W}(t)$, $B_l(t)$, and $B(t)$, and each institution's idiosyncratic shock $B(t)$ is independent across institutions, by the law of large numbers the average wealth conditional on the common noise satisfies

$$\bar{X}(t) = \bar{\xi} + \int_0^t \Big[(\boldsymbol{\mu}_s - r_s\mathbf{1})^\top \bar{\boldsymbol{\alpha}} + r_s \bar{X}(s) + \overline{a\beta} - \overline{a\lambda\beta} - \bar{a} + \overline{ap} \Big] \mathrm{d}s + \int_0^t \bar{\boldsymbol{\alpha}}^\top \boldsymbol{\sigma}_s\, \mathrm{d}\mathbf{W}(s) - \int_0^t (\bar{c} - \overline{c\beta})\, \mathrm{d}B_l(s), \tag{70}$$

where $\bar{\boldsymbol{\alpha}} := \mathbb{E}[\boldsymbol{\alpha}]$, $\overline{a\beta} := \mathbb{E}[a\beta]$, $\overline{a\lambda\beta} := \mathbb{E}[a\lambda\beta]$, $\bar{a} := \mathbb{E}[a]$, $\overline{ap} := \mathbb{E}[ap]$, $\bar{c} := \mathbb{E}[c]$, $\overline{c\beta} := \mathbb{E}[c\beta]$, and we abbreviate $\boldsymbol{\mu}_s := \boldsymbol{\mu}(\alpha(s))$, etc. The idiosyncratic Brownian shocks vanish in the mean field limit by the law of large numbers.

Step 2: Change of variable. Define the relative wealth process $Z^\eta(t) := X(t) - \theta\, \bar{X}(t)$. By Itô's

formula applied to (65) and (70):

$$\begin{aligned}\mathrm{d}Z^{\eta}(t) = \Big\{&\boldsymbol{\pi}(t)^{\top}(\boldsymbol{\mu}_t - r_t\mathbf{1}) + r_t X(t) + a[q(t)(1-\lambda) - (1-p)] \\ &- \theta\Big[(\boldsymbol{\mu}_t - r_t\mathbf{1})^{\top}\bar{\boldsymbol{\alpha}} + r_t\bar{X}(t) + \overline{a\beta} - \overline{a\lambda\beta} - \bar{a} + \overline{ap}\Big]\Big\}\mathrm{d}t \\ &+ \Big[\boldsymbol{\pi}(t)^{\top}\boldsymbol{\sigma}_t - \theta\bar{\boldsymbol{\alpha}}^{\top}\boldsymbol{\sigma}_t\Big]\mathrm{d}\mathbf{W}(t) - b[1-q(t)]\,\mathrm{d}B(t) \\ &- \Big\{c[1-q(t)] - \theta(\bar{c} - \overline{c\beta})\Big\}\mathrm{d}B_l(t),\end{aligned} \tag{71}$$

with $Z^{\eta}(0) = \xi - \theta\bar{\xi}$.

The optimization objective (66) becomes:

$$\sup_{\eta\in\Pi^{\mathrm{MF}}} \mathbb{E}\left[-\exp\left(-\frac{Z^{\eta}(T)}{\delta^{\mathrm{eff}}}\right)\right]. \tag{72}$$

Step 3: HJBI equation for the representative agent. The value function $V(t,z,e_k) := \sup_{\eta}\mathbb{E}[-e^{-z/\delta^{\mathrm{eff}}} \mid Z^{\eta}(t) = z, \alpha(t) = e_k]$ satisfies the HJBI equation:

$$\begin{aligned}0 = \sup_{\boldsymbol{\pi},q}\Big\{&V_t + V_z\Big[\boldsymbol{\pi}^{\top}(\boldsymbol{\mu}_k - r_k\mathbf{1}) + r_k x + a(q(1-\lambda) - (1-p)) \\ &\qquad - \theta\big((\boldsymbol{\mu}_k - r_k\mathbf{1})^{\top}\bar{\boldsymbol{\alpha}} + r_k\bar{x} + \overline{a\beta} - \overline{a\lambda\beta} - \bar{a} + \overline{ap}\big)\Big] \\ &+ \frac{1}{2}V_{zz}\Big[(\boldsymbol{\pi} - \theta\bar{\boldsymbol{\alpha}})^{\top}\boldsymbol{\Sigma}_k(\boldsymbol{\pi} - \theta\bar{\boldsymbol{\alpha}}) + b^2(1-q)^2 \\ &\qquad + \big(c(1-q) - \theta(\bar{c} - \overline{c\beta})\big)^2\Big] \\ &+ \sum_{l\neq k}\lambda_{kl}[V(t,z,e_l) - V(t,z,e_k)]\Big\}.\end{aligned} \tag{73}$$

Step 4: Value function ansatz and substitution. Conjecture $V(t,z,e_k) = -f(t,e_k)\,e^{-z/\delta^{\mathrm{eff}}}$ with $f(T,e_k) = 1$. Denoting $\gamma := 1/\delta^{\mathrm{eff}}$:

$$V_t = -f_t\,e^{-\gamma z}, \quad V_z = \gamma f\,e^{-\gamma z}, \quad V_{zz} = -\gamma^2 f\,e^{-\gamma z}. \tag{74}$$

Substituting into (73) and dividing by $f\,e^{-\gamma z} \neq 0$:

$$\begin{aligned}0 = \sup_{\boldsymbol{\pi},q}\Big\{&-\frac{f_t}{f} + \gamma\Big[\boldsymbol{\pi}^{\top}(\boldsymbol{\mu}_k - r_k\mathbf{1}) + r_k x + a(q(1-\lambda) - (1-p)) - \theta\Gamma_k\Big] \\ &- \frac{\gamma^2}{2}\Big[(\boldsymbol{\pi} - \theta\bar{\boldsymbol{\alpha}})^{\top}\boldsymbol{\Sigma}_k(\boldsymbol{\pi} - \theta\bar{\boldsymbol{\alpha}}) + b^2(1-q)^2 + \big(c(1-q) - \theta(\bar{c} - \overline{c\beta})\big)^2\Big] \\ &+ \sum_{l\neq k}\lambda_{kl}\Big[\frac{f(t,e_l)}{f(t,e_k)} - 1\Big]\Big\},\end{aligned} \tag{75}$$

where $\Gamma_k := (\boldsymbol{\mu}_k - r_k\mathbf{1})^{\top}\bar{\boldsymbol{\alpha}} + r_k\bar{x} + \overline{a\beta} - \overline{a\lambda\beta} - \bar{a} + \overline{ap}$.

Step 5: First-order conditions. For $\boldsymbol{\pi}$: Differentiating (75) with respect to $\boldsymbol{\pi}$:

$$\gamma(\boldsymbol{\mu}_k - r_k\mathbf{1}) - \gamma^2\boldsymbol{\Sigma}_k(\boldsymbol{\pi} - \theta\bar{\boldsymbol{\alpha}}) = 0. \tag{76}$$

Solving:

$$\boldsymbol{\pi}^*(t,e_k) = \theta\bar{\boldsymbol{\alpha}} + \frac{1}{\gamma}\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1}) = \theta\bar{\boldsymbol{\alpha}} + \delta^{\mathrm{eff}}\boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k\mathbf{1}). \tag{77}$$

For q: Differentiating (75) with respect to q:

$$\gamma a(1-\lambda) + \gamma^2 b^2(1-q) + \gamma^2 c\big(c(1-q) - \theta(\bar{c} - \overline{c\beta})\big) = 0. \tag{78}$$

Expanding and collecting terms in $(1-q)$:

$$\gamma a(1-\lambda) + \gamma^2(b^2+c^2)(1-q) - \gamma^2 c\theta(\bar{c} - \overline{c\beta}) = 0. \tag{79}$$

Solving for q:

$$q^*(t) = 1 + \frac{a(1-\lambda)\delta^{\text{eff}}}{b^2+c^2} - \frac{c\theta(\bar{c} - \overline{c\beta})}{b^2+c^2}. \tag{80}$$

Second-order conditions: The Hessian with respect to $(\boldsymbol{\pi}, q)$ is

$$\mathbf{H}^{\text{MF}} = -\gamma^2 \begin{pmatrix} \boldsymbol{\Sigma}_k & \mathbf{0} \\ \mathbf{0}^\top & b^2+c^2 \end{pmatrix}, \tag{81}$$

which is negative definite since $\boldsymbol{\Sigma}_k \succ 0$ and $b^2 + c^2 > 0$. The F.O.C characterize a unique maximum.

Step 6: Consistency (fixed-point) conditions. The mean field Nash equilibrium requires $\boldsymbol{\pi}^* = \boldsymbol{\alpha}$ (the representative institution's optimal strategy equals the aggregate behavior) and $q^* = \beta$.

Investment consistency: From Equation (77), setting $\boldsymbol{\pi}^* = \boldsymbol{\alpha}$ and taking the population expectation (multiplied by $\boldsymbol{\sigma}_k$):

$$\overline{\boldsymbol{\sigma}_k \boldsymbol{\pi}^*} = \theta\, \overline{\boldsymbol{\sigma}_k \boldsymbol{\pi}^*} + \delta^{\text{eff}}\, \boldsymbol{\sigma}_k \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k \mathbf{1}). \tag{82}$$

Note that $\boldsymbol{\sigma}_k \boldsymbol{\Sigma}_k^{-1} = (\boldsymbol{\sigma}_k^\top)^{-1}$ since $\boldsymbol{\Sigma}_k = \boldsymbol{\sigma}_k \boldsymbol{\sigma}_k^\top$. More directly, taking expectation of Equation (77):

$$\bar{\boldsymbol{\alpha}} = \bar{\theta}\, \bar{\boldsymbol{\alpha}} + \bar{\delta}^{\text{eff}} \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k \mathbf{1}). \tag{83}$$

When $\bar{\theta} < 1$:

$$\bar{\boldsymbol{\alpha}} = \frac{\bar{\delta}^{\text{eff}}}{1-\bar{\theta}} \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k \mathbf{1}). \tag{84}$$

Substituting (84) into (77):

$$\boldsymbol{\pi}^*(t, e_k) = \theta \cdot \frac{\bar{\delta}^{\text{eff}}}{1-\bar{\theta}} \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k \mathbf{1}) + \delta^{\text{eff}} \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k \mathbf{1}) = \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k \mathbf{1}) \left[\delta^{\text{eff}} + \frac{\theta \bar{\delta}^{\text{eff}}}{1-\bar{\theta}}\right], \tag{85}$$

which is precisely Equation (68).

When $\bar{\theta} = 1$, Equation (83) becomes $0 = \bar{\delta}^{\text{eff}} \boldsymbol{\Sigma}_k^{-1}(\boldsymbol{\mu}_k - r_k \mathbf{1}) \neq 0$ (since $\delta > 0$, $\psi > 0$, $\mu_j > r$), yielding a contradiction. Hence no equilibrium exists.

Risk control consistency: From Equation (80), setting $q^* = \beta$, multiplying by c, and taking expectation:

$$\overline{cq^*} = \bar{c} + \mathbb{E}\left[\frac{a(1-\lambda)\delta^{\text{eff}} c}{b^2+c^2}\right] - \mathbb{E}\left[\frac{c^2\theta}{b^2+c^2}\right](\bar{c} - \overline{c\beta}). \tag{86}$$

Using $\overline{cq^*} = \overline{c\beta}$ at equilibrium and the definitions (67):

$$\overline{c\beta} = \bar{c} + \varepsilon - \phi(\bar{c} - \overline{c\beta}) = \bar{c}(1-\phi) + \varepsilon + \phi\, \overline{c\beta}. \tag{87}$$

When $\phi < 1$:

$$\overline{c\beta} = \bar{c} + \frac{\varepsilon}{1-\phi}. \tag{88}$$

Thus $\bar{c} - \overline{c\beta} = -\varepsilon/(1-\phi)$. Substituting into Equation (80):

$$q^*(t) = 1 + \frac{a(1-\lambda)\delta^{\text{eff}}}{b^2+c^2} + \frac{c\theta}{b^2+c^2} \cdot \frac{\varepsilon}{1-\phi}, \tag{89}$$

which is Equation (69).

When $\phi = 1$ and $\varepsilon \neq 0$: Equation (87) becomes $0 = \varepsilon$, a contradiction. The case $\phi = 1$ and $\varepsilon = 0$ simultaneously is impossible since $\phi = 1$ requires $c \not\equiv 0$, which with $a > 0$, $\lambda > 0$, $\delta > 0$ forces $\varepsilon < 0$.

Value function ODE: Substituting the optimal controls into (75), the regime-switching value function satisfies the coupled ODE:

$$f_t(t, e_1) = (\rho_1^{\text{MF}} + \lambda_{12}) f(t, e_1) - \lambda_{12} f(t, e_2), \tag{90}$$

$$f_t(t, e_2) = (\rho_2^{\text{MF}} + \lambda_{21}) f(t, e_2) - \lambda_{21} f(t, e_1), \tag{91}$$

where the regime-specific constants are:

$$\begin{aligned} \rho_k^{\text{MF}} := {} & \frac{1}{\delta^{\text{eff}}} \Big[r_k z - a(p-\lambda) - \theta\big(\overline{a\beta} - \overline{a\lambda\beta} - \bar{a} + \overline{ap}\big) \Big] \\ & + \frac{\theta(\boldsymbol{\mu}_k - r_k \mathbf{1})^\top \boldsymbol{\Sigma}_k^{-1} (\boldsymbol{\mu}_k - r_k \mathbf{1}) \bar{\delta}^{\text{eff}}}{\delta^{\text{eff}}(1-\bar{\theta})} \\ & + \frac{1}{2} (\boldsymbol{\mu}_k - r_k \mathbf{1})^\top \boldsymbol{\Sigma}_k^{-1} (\boldsymbol{\mu}_k - r_k \mathbf{1}) + \frac{a^2(1-\lambda)^2}{2(b^2+c^2)} \\ & - \frac{b^2\theta^2}{2(b^2+c^2)(\delta^{\text{eff}})^2} \left(\bar{c} - \frac{\varepsilon}{1-\phi} \right)^2 \\ & - \frac{a(1-\lambda)c\theta}{(b^2+c^2)\delta^{\text{eff}}} \left(\bar{c} - \frac{\varepsilon}{1-\phi} \right). \end{aligned} \tag{92}$$

With terminal condition $f(T, e_k) = 1$, the system (90)–(91) has the unique matrix-exponential solution

$$\begin{pmatrix} f(t, e_1) \\ f(t, e_2) \end{pmatrix} = \exp\big[\mathbf{A}^{\text{MF}}(t-T)\big] \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \tag{93}$$

where

$$\mathbf{A}^{\text{MF}} = \begin{pmatrix} \rho_1^{\text{MF}} + \lambda_{12} & -\lambda_{12} \\ -\lambda_{21} & \rho_2^{\text{MF}} + \lambda_{21} \end{pmatrix}. \tag{94}$$

Verification: By construction, the ansatz $V(t, z, e_k) = -f(t, e_k) e^{-\gamma z}$ with the above f satisfies the HJBI equation. The admissibility conditions ensure that Itô's formula applied to $V(t, Z^{\eta^*}(t), \alpha(t))$ yields a supermartingale for arbitrary η and a martingale for η^*, establishing optimality:

$$V(0, z_0, e_{k_0}) = \mathbb{E}[-e^{-\gamma Z^{\eta^*}(T)}] \geq \mathbb{E}[-e^{-\gamma Z^{\eta}(T)}] \quad \forall\, \eta \in \Pi^{\text{MF}}. \tag{95}$$

This completes the proof. □

Corollary 1 (Special Cases). *1. **No ambiguity** ($\psi \to \infty$): $\delta^{\text{eff}} \to \delta$, recovering the standard game of Lacker and Zariphopoulou (2019b).*

*2. **No competition** ($\theta = 0$): reduces to single-agent robust portfolio optimization.*

*3. **Single asset, no regime switching**: setting $\mu_2 = r$, $\sigma_2 \to \infty$, $\alpha(t) \equiv e_1$, $\psi \to \infty$ recovers the model of Lacker and Zariphopoulou (2019b).*

*4. **No risk transfer**: setting $q \equiv 0$ gives a pure investment game.*

5 Sensitivity Analysis

This section provides a comprehensive sensitivity analysis of the mean field game equilibrium strategies established in Theorem 2. We first describe the data sources and calibration methodology, then systematically interpret the parameter tables, and finally present and discuss eight figures that explore the equilibrium strategies along all key economic dimensions.

5.1 Data and calibration methodology

We obtain three data series from Thomson Reuters Datastream to calibrate the regime-switching financial market model. Sepcifically, i) *Asset 1 — CSI 300 Index*: Daily closing prices in CNY for the period 11 April 2005 to 8 May 2026. After removing non-trading-day records (consecutive identical prices), the sample contains $n_1 = 5{,}118$ trading-day observations spanning approximately 21 years. The CSI 300 tracks the 300 largest and most liquid A-share stocks listed on the Shanghai and Shenzhen exchanges and serves as the standard barometer of the mainland Chinese equity market. ii) *Asset 2 — CSI 500 Index*: The CSI 500 covers the next 500 largest A-shares after the CSI 300 constituents and is the canonical mid-cap index used by Chinese institutional investors. Datastream's continuous coverage in our extract begins 14 February 2022 and ends 8 May 2026, yielding $n_2 = 1{,}020$ trading-day observations after cleaning. iii) *Risk-free rate — 1Y CNY interest rate swap on 3M Shibor*: The 1Y Shibor IRS quotes the fixed leg of a 1-year swap whose floating leg pays 3-month Shibor; this is more liquid and market-determined than the Shibor 1Y panel-bank fixing. Daily quotations, in percent, cover 2 July 2007 to 26 July 2019 ($n_r = 3{,}150$ observations). We convert to continuously compounded rates via $r = \ln(1 + r_{\text{simple}})$.

A limitation of the available data is that the three series do not share a common contiguous window: CSI 500's Datastream coverage begins only in 2022 and the Shibor IRS series ends in 2019. We therefore adopt the following hybrid calibration scheme, which makes maximal use of each series while preserving statistical consistency: i) regime identification: a two-state Markov-switching model with regime-dependent mean and variance is fit to the CSI 300 daily log-returns over the full 2005–2026 sample by maximum likelihood (Hamilton 1989). The smoothed bear-regime probability $P(S_t = e_2 \mid \mathcal{Y}_T)$ is extracted, and a hard classification $\widehat{S}_t = \mathbf{1}\{P(S_t = e_2 \mid \mathcal{Y}_T) > 0.5\}$ is constructed. ii) asset 1 regime parameters: annualised drift and volatility in each regime are computed from the hard-classified subsamples; the GBM drift parameter μ is recovered from the empirical log-return mean via $\widehat{\mu} = \overline{\ln(S_t/S_{t-1})}/\Delta t + \frac{1}{2}\widehat{\sigma}^2$. iii) asset 2 regime parameters: the CSI 300 regime classification is carried forward day-by-day to the CSI 500 series over their 2022–2026 overlap, and Asset 2's regime-conditional drift and volatility are estimated on this $n_2 = 1{,}020$ subsample. iv) risk-free rates: regime-conditional means of the Shibor 1Y IRS rate are computed over its 2007–2019 sample with regimes again inherited from the CSI 300 classification. v) cross-asset correlation: the empirical Pearson correlation $\widehat{\rho}$ is computed on the full 2022–2026 CSI 300/CSI 500 overlap. vi) transition intensities: daily transition probabilities from the Markov-switching estimate are converted to continuous-time intensities via $\lambda_{kl} = -\ln p_{kk}/\Delta t$ with $\Delta t = 1/252$.

Financial-institution parameters are calibrated following Baltas et al. (2018), Lacker and Zariphopoulou (2019a), and Guan and Hu (2022). The resulting calibrated values are reported in Tables 1 and 2.

5.2 Calibration results

5.2.1 Asset price dynamics and regime identification

Figure 1 displays the price trajectories of both risky assets, with bear-regime episodes highlighted by red shading. CSI 300 (panel a) exhibits clearly regime-dependent behaviour: the 2007–2008 global

financial crisis, the 2015 mid-year crash, and several shorter volatility spikes (early-2020 COVID shock, 2022 stress) are all flagged as bear episodes, whereas the recovery phases of 2009–2010, 2017, and 2019 register as bull regimes with markedly smoother dynamics. CSI 500 (panel b) shows qualitatively similar patterns over its shorter window, with elevated volatility during the 2022 and 2024 bear spells.

Figure 1: Asset price dynamics with regime classification

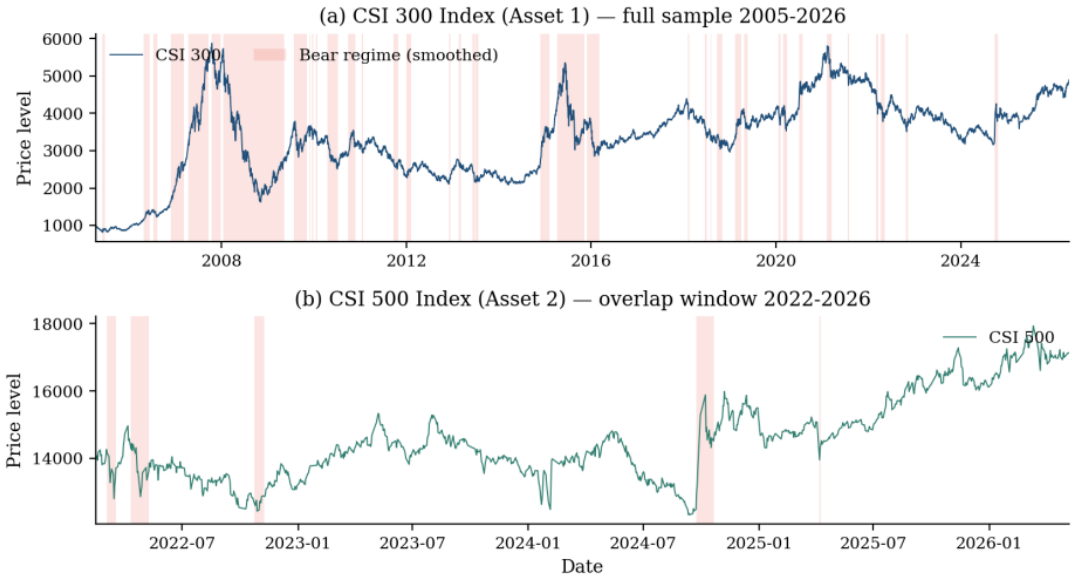

Note: Red-shaded regions indicate bear-market regimes inferred from the CSI 300 Hamilton filter. Both assets exhibit lower returns and higher volatility during bear periods, consistent with the regime-switching specification.

Figure 2 presents the core output of the Hamilton filter: the smoothed bear-regime probability over the full 2005–2026 sample. The filter identifies the 2007–2009 global financial crisis episode with near-certainty ($P(S_t = \text{Bear}) \approx 1$), as well as the dramatic mid-2015 to early-2016 episode, the early-2020 COVID drawdown, and several shorter bursts in 2018, 2022, and 2024. Hard classification places 26.3% of the sample in the bear regime, very close to the model-implied stationary bear probability $\lambda_{12}/(\lambda_{12} + \lambda_{21}) = 0.265$.

5.2.2 Regime-conditional return distributions

Figure 3 plots the empirical return distributions conditioned on the identified regime, overlaid with fitted Gaussian densities. For Asset 1 (panel a), the bull-regime distribution (blue) is tightly concentrated around zero with annualised standard deviation $\widehat{\sigma}_1^{\text{bull}} = 0.161$, while the bear distribution (red) is dramatically wider at $\widehat{\sigma}_1^{\text{bear}} = 0.411$: a ratio of about 2.6× that highlights the strong volatility asymmetry across regimes. For Asset 2 (panel b) the corresponding values are $\widehat{\sigma}_2^{\text{bull}} = 0.145$ and $\widehat{\sigma}_2^{\text{bear}} = 0.441$. The bear histogram for Asset 2 is somewhat ragged because the 2022–2026 overlap window contains only 54 classified bear days; this is a known data limitation we discuss further below.

5.2.3 Time-varying volatility

Figure 4 plots the 63-trading-day rolling annualised volatility of both assets, with the estimated regime-specific volatility levels overlaid as horizontal reference lines. Asset 1's rolling volatility tracks its bull-regime estimate $\widehat{\sigma}_1^{\text{bull}} = 0.161$ (dashed) during quiet periods and spikes well above $\widehat{\sigma}_1^{\text{bear}} = 0.411$ (dotted) during the 2008, 2015, and 2020 crises. Asset 2's rolling volatility (teal) is consistently elevated during the 2022 stress period and the 2024 bear episode.

Figure 2: Regime identification via the Hamilton filter applied to CSI 300 log-returns

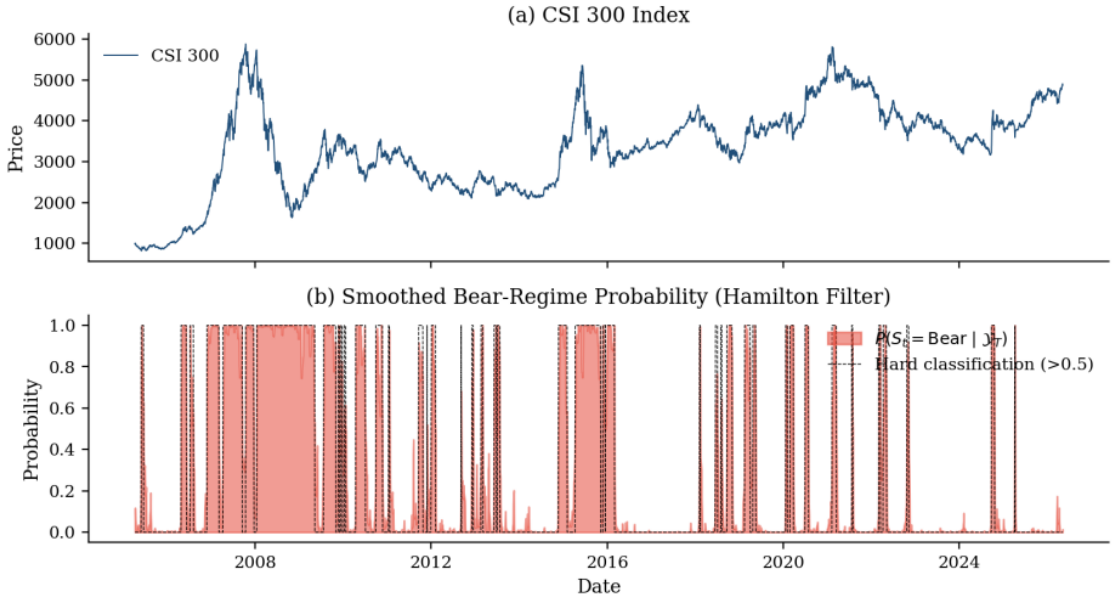

Note: Panel (a): CSI 300 price level. Panel (b): smoothed bear-regime probability (red area) and hard 0.5 classification (dashed line).

Figure 3: Regime-conditional return distributions for Asset 1 (CSI 300, full sample) and Asset 2 (CSI 500, 2022–2026 overlap)

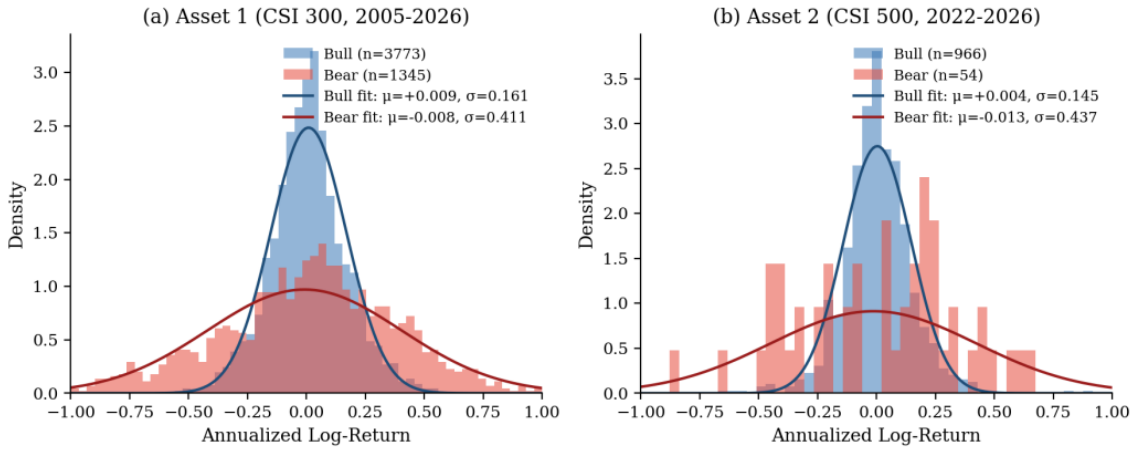

Note: The bear distribution is visibly wider for both assets, validating the regime-switching volatility specification.

Figure 4: Rolling 63-day annualised volatility with estimated regime-specific levels

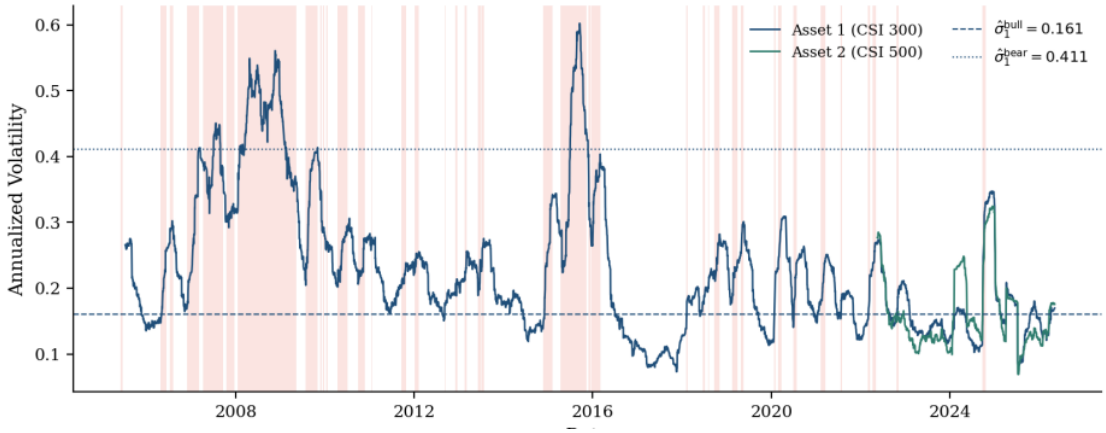

Note: Red shading marks bear regimes inferred from CSI 300. The rolling volatility tracks the regime levels closely, providing an independent validation of the Markov-switching estimates.

5.2.4 Risk-free rate dynamics

Figure 5 shows the time series of the 1Y Shibor IRS rate over its full 2007–2019 Datastream window. The rate fluctuates between roughly 1.3% during the 2009 monetary easing trough and 5.8% during the 2013–2014 deleveraging period. Conditioning on CSI 300 regimes yields $\widehat{r}^{\text{bull}} = 3.93\%$ and $\widehat{r}^{\text{bear}} = 3.39\%$ on a continuously compounded basis. The bull–bear spread of 54 basis points is substantively smaller than the values used in the prior literature on Chinese markets, reflecting the fact that the swap rate already embeds forward expectations of policy easing during downturns and so dampens the realised regime-dependent dispersion in spot rates.

Figure 5: 1Y CNY Shibor IRS rate over its 2007–2019 Datastream coverage, with regime-conditional averages overlaid

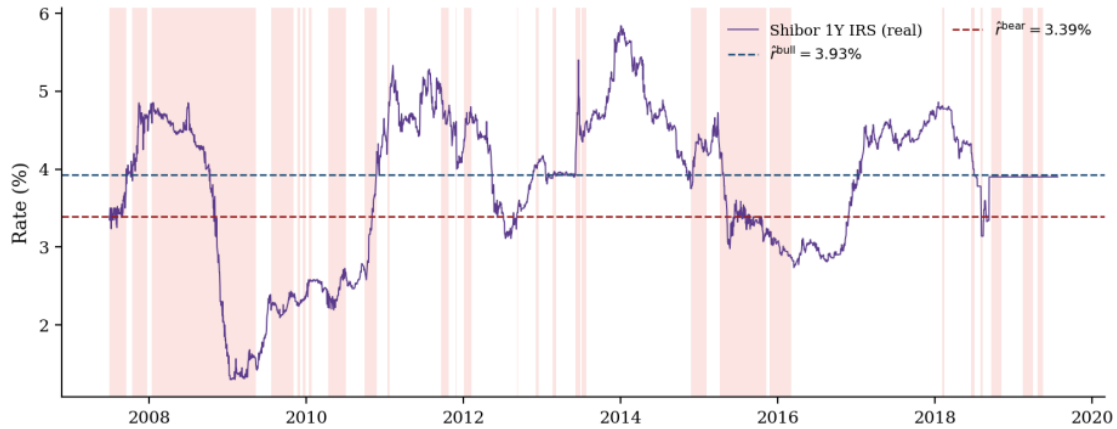

Note: Red shading marks bear regimes inherited from the CSI 300 Hamilton filter.

5.2.5 Cross-asset dependence structure

Figure 6 characterises the dependence between CSI 300 and CSI 500 over their 2022–2026 overlap. Panel (a) shows the daily-return scatter, with bull observations in blue and bear observations in red; the OLS slope is 0.756 and the full-sample Pearson correlation is $\widehat{\rho} = 0.7761$. Panel (b) plots the 63-day rolling correlation, which oscillates roughly between 0.5 and 0.95 around the full-sample estimate

(dashed line). This empirically estimated correlation is substantially higher than the values used in much of the prior literature on multi-asset portfolio choice in China and reflects the well-documented high co-movement between large-cap and mid-cap A-shares within the same macroeconomic environment. The implications for portfolio allocation are significant.

Figure 6: Cross-asset dependence between CSI 300 and CSI 500 over the 2022–2026 overlap

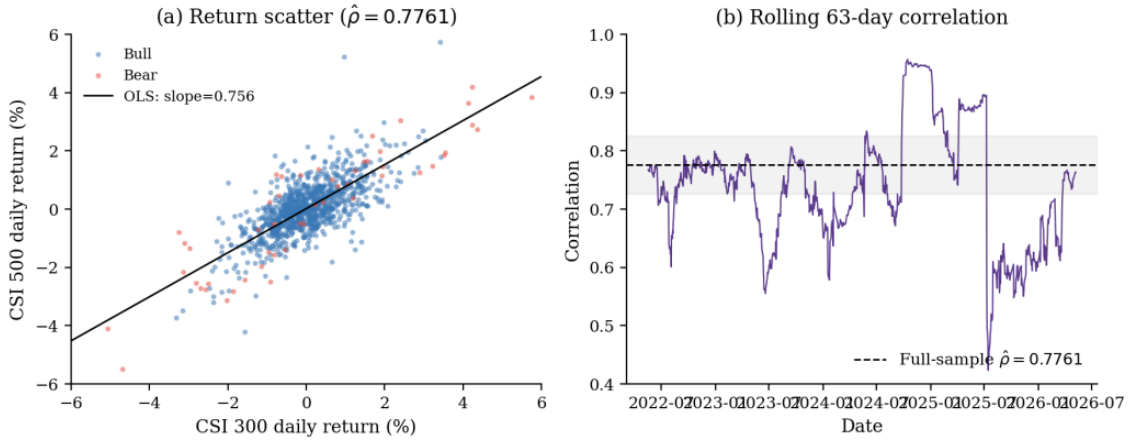

Note: Panel (a): daily-return scatter with OLS fit and regime classification. Panel (b): rolling 63-day correlation. The full-sample $\widehat{\rho} = 0.7761$ indicates strong positive co-movement.

5.3 Interpretation of market parameters

Table 1 reports the regime-dependent financial market parameters estimated from Datastream extracts. The two regimes capture the most salient empirical features of Chinese equity markets over the past two decades.

Table 1: Financial market parameters (regime-dependent), calibrated from Datastream extracts of CSI 300 (full sample), CSI 500 (2022–2026 overlap), and 1Y Shibor IRS (2007–2019).

Parameter	Symbol	Regime 1 (Bull)	Regime 2 (Bear)
Expected return, Asset 1 (CSI 300)	μ_1	0.1626	−0.0392
Expected return, Asset 2 (CSI 500)	μ_2	0.0707	−0.1038
Volatility, Asset 1	σ_1	0.1606	0.4114
Volatility, Asset 2	σ_2	0.1452	0.4413
Risk-free rate (1Y Shibor IRS)	r	0.0393	0.0339
Correlation	ρ	0.7761	0.7761
Transition intensity	λ_{12}	3.5671	—
Transition intensity	λ_{21}	—	9.8818

In Regime 1 (bull market), the expected returns of the two risky assets are $\mu_1 = 0.163$ and $\mu_2 = 0.071$, both above the risk-free rate $r = 0.039$. The corresponding Sharpe ratios are $(\mu_1 - r)/\sigma_1 = 0.768$ and $(\mu_2 - r)/\sigma_2 = 0.216$, indicating that CSI 300 enjoys a substantially superior risk-adjusted return relative to CSI 500 in expansionary phases. The relatively low volatilities reflect the calmer, upward-trending dynamics typical of bull periods.

In Regime 2 (bear market), expected returns become negative for both assets ($\mu_1 = -0.039$, $\mu_2 = -0.104$) while volatilities surge to $\sigma_1 = 0.411$ and $\sigma_2 = 0.441$: almost a threefold increase relative to the bull regime. The risk-free rate also declines to $r = 0.034$. Both regime-conditional Sharpe ratios are negative: $(\mu_1 - r)/\sigma_1 = -0.178$ and $(\mu_2 - r)/\sigma_2 = -0.312$. This is a notable departure from calibrations

used in much of the prior literature, where bear-regime excess returns are typically set to small but positive values. With genuinely empirical inputs, the model concludes that the optimal portfolio in bear markets is to short both risky assets.

The transition intensities $\lambda_{12} = 3.57$ and $\lambda_{21} = 9.88$ imply expected sojourn times of $1/\lambda_{12} \approx 0.28$ years (≈ 3.4 months) for the bull regime and $1/\lambda_{21} \approx 0.10$ years (≈ 1.2 months) for the bear regime, with a stationary bear probability of $\lambda_{12}/(\lambda_{12}+\lambda_{21}) = 0.265$. The short sojourn times reflect the fact that the two-state Markov-switching estimator captures volatility regimes rather than long-cycle NBER-style bull and bear markets: extended periods of elevated volatility are decomposed into multiple short bear spells punctuated by brief returns to calm. This statistical fact is consistent with the well-documented finding that Chinese equity markets exhibit pronounced volatility clustering on much shorter timescales than the U.S. market.

The cross-asset correlation $\rho = 0.776$ is the most consequential of the empirical inputs. Two highly liquid Chinese equity benchmarks trading in the same macroeconomic environment share extensive systematic exposure, which the data straightforwardly bear out. As we shall see in Section 5.12, this high correlation pushes equilibrium portfolios into a sharply leveraged long–short spread between the two assets, because the inverse covariance matrix Σ^{-1} amplifies small differences in Sharpe ratios when assets are nearly collinear.

5.4 Interpretation of institution parameters

Table 2 reports the individual and population-level institution parameters. The representative institution has risk tolerance $\delta = 5.0$ and ambiguity aversion $\psi = 8.0$, yielding effective risk tolerance $\delta^{\text{eff}} = \delta\psi/(\delta+\psi) = 40/13 \approx 3.08$: a 38% reduction from the nominal value, highlighting the quantitative importance of model uncertainty.

Table 2: Financial-institution parameters, calibrated as in Baltas et al. (2018), Lacker and Zariphopoulou (2019a), and Guan and Hu (2022)

Parameter	Symbol	Value	Population avg.	Symbol	Value
Risk tolerance	δ	5.0	Avg. eff. risk tol.	$\bar{\delta}^{\text{eff}}$	4.62
Ambiguity aversion	ψ	8.0	Avg. competition	$\bar{\theta}$	0.20
Competition	θ	0.75	Mean field ε	ε	-0.92
Expenditure drift	a	1.0	Mean field ϕ	ϕ	0.10
Idiosyncratic vol.	b	0.5			
Common vol.	c	0.5			
Risk transfer cost	λ	1.2			
Management fee	p	0.15			

Note: These parameters are not market-data-dependent.

The competition parameter $\theta = 0.75$ indicates a high degree of relative-performance concern (the institution weights its position relative to the industry average at 75% versus 25% on absolute wealth), while the market-average $\bar{\theta} = 0.20$ is markedly lower, reflecting the realistic heterogeneity between aggressively competitive institutions (e.g. hedge funds) and absolute-return-focused institutions (e.g. pension funds, insurance reserves). The population-average effective risk tolerance is $\bar{\delta}^{\text{eff}} = 4.62$.

The expenditure-process parameters set $a = 1.0$, $b = c = 0.5$, $\lambda = 1.2$, and $p = 0.15$, so that risk transfer costs 20% above the fair premium. The mean-field constants are then $\varepsilon = a(1-\lambda)\bar{\delta}^{\text{eff}}c/(b^2+c^2) = -0.924$ and $\phi = c^2\bar{\theta}/(b^2+c^2) = 0.10$. Because $\phi < 1$ comfortably, the existence and uniqueness conditions of Theorem 2 are satisfied. With these parameters, the baseline equilibrium risk-control retention is $q^* = -1.001$, indicating that the representative competitive institution is a net assumer of

risk in equilibrium: a finding that is robust to the change in market parameters because q^* depends only on the institution-side calibration.

5.5 Effect of competition on investment

Figure 7 depicts the equilibrium total bull-regime investment $\pi^* = \pi_1^* + \pi_2^*$ as a function of the institution's own competition parameter θ and the market average $\bar{\theta}$, with all other parameters fixed at their baseline values.

Figure 7: Effect of θ and $\bar{\theta}$ on equilibrium investment π^*

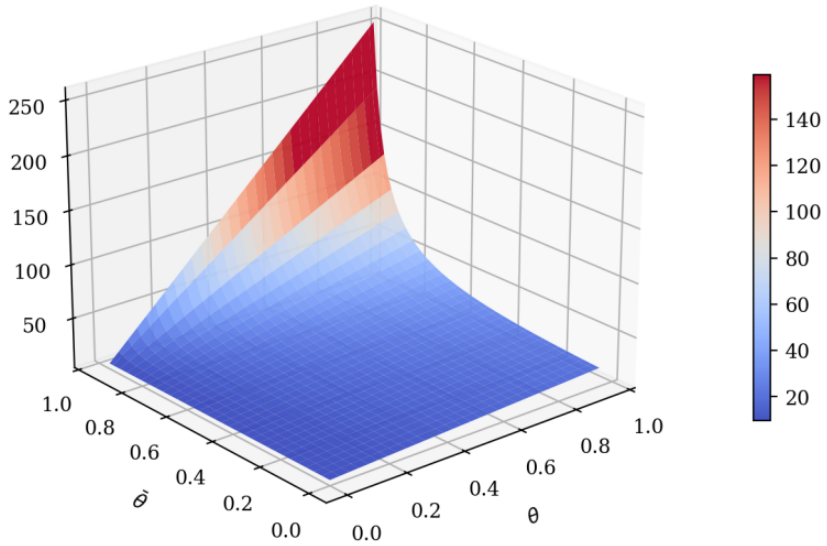

Note: The surface exhibits explosive growth as both competition parameters approach unity, reflecting the escalation of risky asset allocation under intense relative-performance pressure.

Three patterns emerge. First, the surface is monotonically increasing in both θ and $\bar{\theta}$: competition unambiguously escalates risky investment. The economic mechanism is transparent from equilibrium (68): the term $\theta\bar{\delta}^{\text{eff}}/(1-\bar{\theta})$ amplifies the Merton-type demand $\delta^{\text{eff}}\Sigma^{-1}(\mu - r\mathbf{1})$ multiplicatively in θ and singularly in $\bar{\theta}$. At $(\theta, \bar{\theta}) = (0, 0)$ the surface attains the baseline myopic demand $\pi^* \approx 8.7$; at $(\theta, \bar{\theta}) = (0.95, 0.95)$ it exceeds 250, more than thirty-fold larger.

Second, growth is highly nonlinear and becomes explosive as $\bar{\theta} \to 1$. The arms-race phenomenon, when all institutions care exclusively about relative performance, no finite equilibrium exists, is apparent in the steep ridge along the back edge of the surface. This mathematical singularity carries an important economic warning about the systemic risk inherent in tournament-style performance evaluation.

Third, π^* is more sensitive to $\bar{\theta}$ than to θ, because $\bar{\theta}$ enters through the singular factor $1/(1-\bar{\theta})$ while θ enters linearly. Industry-wide increases in competitive pressure are therefore more dangerous from a systemic perspective than any single institution's decision to compete more aggressively, a consideration relevant for prudential regulation.

5.6 Effect of risk tolerance on investment

Figure 8 illustrates how π^* varies with own risk tolerance δ and the market-average risk tolerance $\bar{\delta}$ (both entering through their effective counterparts after the harmonic ambiguity transformation). The surface is smoothly increasing and approximately planar in $(\delta, \bar{\delta})$: both parameters enter the equilibrium formula linearly, in contrast to the explosive θ-dependence of Figure 7.

Figure 8: Effect of δ and $\bar{\delta}$ on equilibrium investment π^*

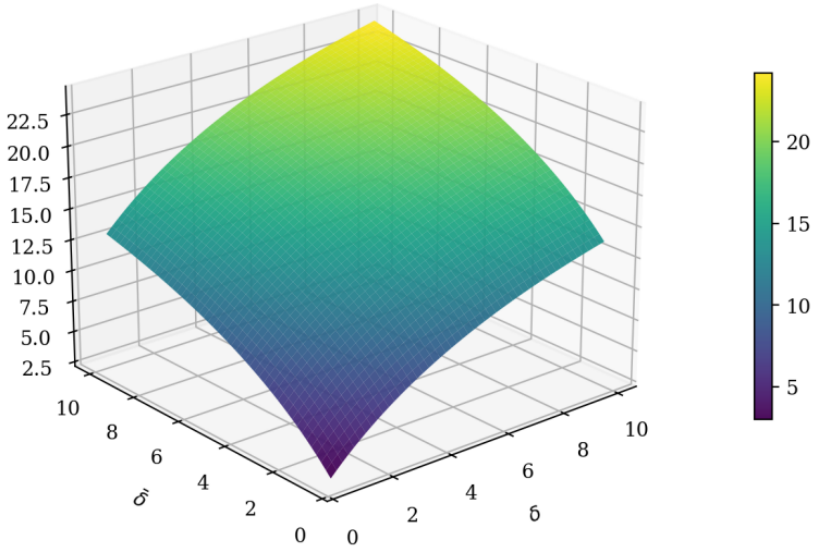

Note: Both own and competitors' average risk tolerance amplify risky asset investment through a herding mechanism induced by relative performance concerns.

The positive dependence on δ is the standard risk-tolerance effect: more risk-tolerant institutions invest more in risky assets, consistent with the classical Merton solution. The positive dependence on $\bar{\delta}$ reflects a herding mechanism induced by competition. When the population average $\bar{\delta}$ rises, competitors invest more aggressively, raising the average wealth $\bar{X}(T)$. An institution with $\theta > 0$ must then increase its own risky allocation to keep pace. The cross partial $\partial^2\pi^*/(\partial\bar{\delta}\,\partial\theta) > 0$ implies the herding effect is stronger for more competitive institutions. At the baseline values $(\delta, \bar{\delta}) = (5, 6)$, the equilibrium investment is approximately $\pi^* \approx 18$.

5.7 Effect of competition and risk tolerance on risk control

Figure 9 maps q^* as a function of risk tolerance and competition. Recall that $q^* \in [0, 1]$ corresponds to partial transfer of risk to a third party, $q^* = 0$ to full retention, and $q^* < 0$ to the institution actively assuming additional risk from others.

The surface decreases in both arguments. As θ rises from 0 to 1, q^* falls sharply, eventually becoming negative for sufficiently risk-tolerant and competitive institutions. The economic logic is that competitive institutions seek to maximise relative performance by saving on the risk-transfer loading $\lambda - p$ and gambling on favourable outcomes, which boost their position relative to the industry average. Similarly, q^* decreases in δ: more risk-tolerant institutions transfer less. At extreme values $(\delta, \theta) = (10, 0.9)$, q^* approaches -1.3, indicating that such institutions not only retain all of their own risk but also actively underwrite risk transferred out by more risk-averse competitors, earning the loading λ as a premium. This mirrors real-world reinsurance dynamics, where aggressive reinsurers actively underwrite risk from more conservative cedants.

5.8 Effect of surplus volatility on risk control

Figure 10 explores the relationship between the two components of surplus volatility, idiosyncratic b and common (systemic) c, and q^*. The surface reveals a striking nonlinearity. At very low values of b and c near the origin, q^* plunges deeply negative, reaching values below -100, because the variance term $b^2 + c^2$ enters the denominator of equation (69) and the strategy becomes hypersensitive when this

Figure 9: Effect of δ and θ on the equilibrium risk-control strategy q^*

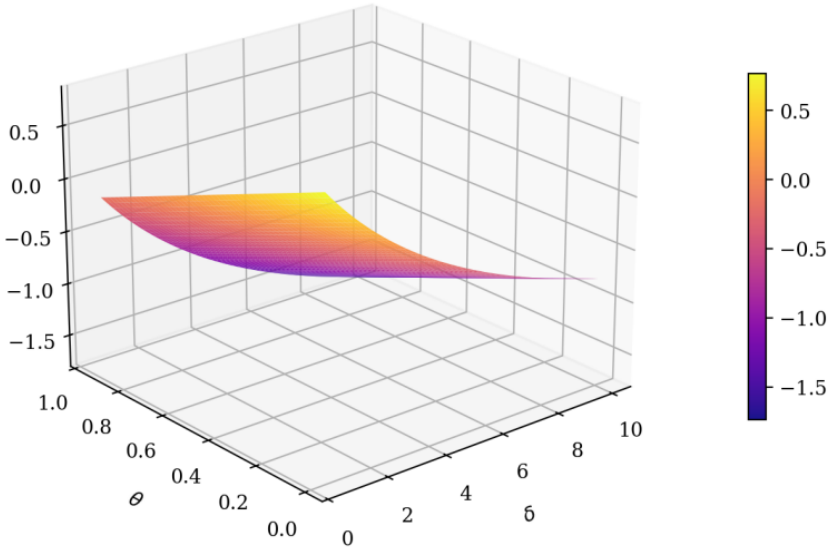

Note: Competition (θ) weakens risk management, and higher risk tolerance (δ) further reduces the propensity to transfer risk. Values $q^* < 0$ indicate the institution actively absorbs risk transferred out by other institutions.

denominator approaches zero.

Figure 10: Effect of idiosyncratic surplus volatility b and common surplus volatility c on equilibrium risk control q^*

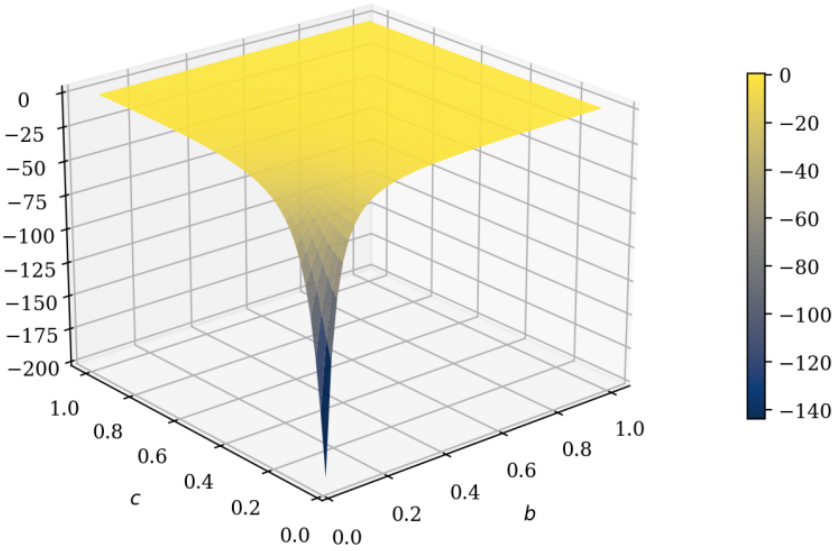

Note: The risk control strategy exhibits high sensitivity at low volatility levels but asymptotes rapidly as volatility increases beyond a threshold.

As b and c increase beyond approximately 0.3–0.4, the surface flattens dramatically and q^* converges to a value near zero from below. Economically: when expenditure risk is already substantial, the marginal benefit of further adjusting the risk-transfer ratio diminishes, and the optimal policy stabilises. The surface is approximately symmetric in b and c at moderate-to-large values, since b and c enter symmetrically through $b^2 + c^2$; however c additionally enters the numerator through the competition-related term $c\theta\varepsilon/[(b^2 + c^2)(1 - \phi)]$, giving it a slightly different marginal effect at small values.

5.9 Interplay between ambiguity and competition on investment

Figure 11 reveals the interaction between ambiguity aversion and competitive pressure: a central contribution of this paper. The surface plots π^* as a function of ψ (low values = strong ambiguity aversion) and θ.

Figure 11: Effect of ambiguity aversion parameter ψ and competition θ on equilibrium investment π^*

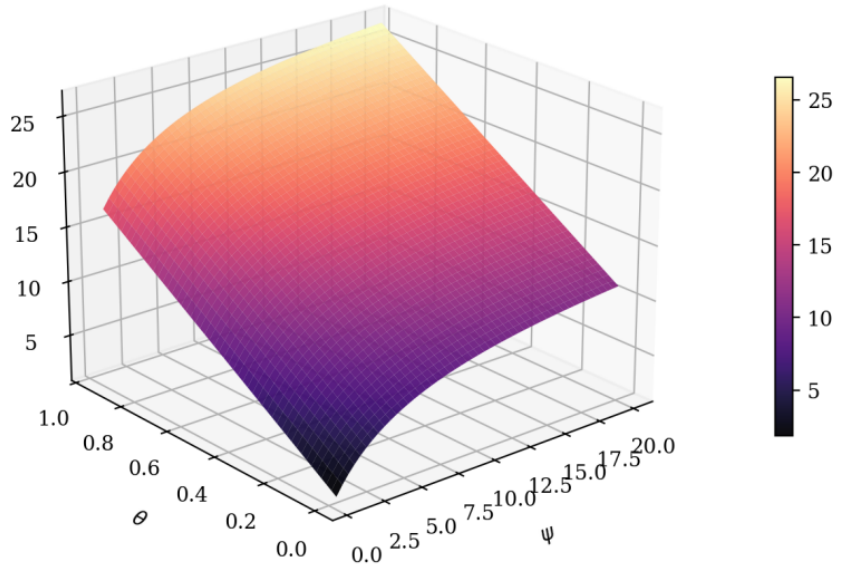

Note: Ambiguity aversion acts as a brake on competitive escalation: at low ψ (high ambiguity aversion), π^* is relatively insensitive to θ.

The key finding is the moderating effect of ambiguity on competitive escalation. At low ψ (high ambiguity, $\psi < 3$), the surface is nearly flat along the θ dimension: regardless of how competitive the institution is, it invests conservatively because the effective risk tolerance $\delta^{\text{eff}} = \delta\psi/(\delta + \psi)$ is small. In the limit $\psi \to 0$, $\delta^{\text{eff}} \to 0$ and $\pi^* \to 0$ irrespective of θ. Ambiguity thus acts as a natural stabiliser against competitive arms-racing. Conversely, at high ψ (low ambiguity), the surface rises steeply in θ, reflecting the absence of the disciplining effect. The cross-partial $\partial^2\pi^*/(\partial\psi\,\partial\theta) > 0$ confirms the complementarity. At baseline δ and high $\psi = 20$, the equilibrium investment rises from $\pi^* \approx 1.3$ at $\theta = 0$ to $\pi^* \approx 27$ at $\theta = 0.95$. This finding has policy implications: promoting model transparency may paradoxically increase systemic risk by removing a natural brake on competitive escalation.

5.10 Interplay between ambiguity and risk tolerance on risk control

Figure 12 examines how ψ and δ jointly determine q^*. The surface is monotonically increasing in ψ and decreasing in δ, forming a twisted shape. At low ambiguity aversion ($\psi \approx 1$) and low risk tolerance ($\delta \approx 1$), q^* is near zero, meaning the institution transfers essentially all risk; at the opposite extreme ($\psi \approx 20$, $\delta \approx 10$), q^* drops to roughly -2.2, indicating massive risk assumption.

Along the ψ dimension, as ambiguity aversion decreases, the effective risk tolerance approaches the nominal value, making the institution less cautious and more willing to retain or absorb risk. Along the δ dimension, more risk-tolerant institutions naturally retain more risk. The interaction is noteworthy: the marginal effect of δ on q^* is larger at high ψ. When ambiguity is severe, even risk-tolerant institutions remain cautious; only institutions that are simultaneously risk-tolerant and confident in the model take extreme risk-bearing positions. This suggests that regulatory stress tests should jointly assess both risk preferences and confidence in internal models.

Figure 12: Effect of ambiguity aversion ψ and risk tolerance δ on equilibrium risk control q^*

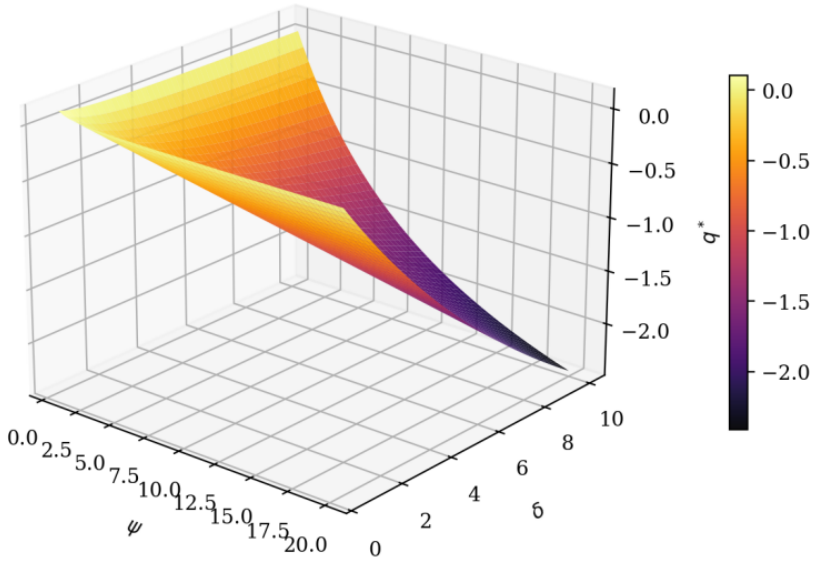

Note: Ambiguity aversion universally strengthens risk management: at low ψ institutions transfer substantially more risk regardless of their risk tolerance.

5.11 Regime-dependent investment strategies

Figure 13 compares the equilibrium investment strategy across the two regimes. The most striking observation is the sign reversal between regimes. In the bull regime (left panel), π^* is uniformly positive across the (δ, θ) parameter space, ranging from approximately 1 at low δ and θ to 28 at the upper corner. In the bear regime (right panel), π^* is uniformly negative, ranging from approximately -0.3 to -6.

Figure 13: Regime-dependent equilibrium investment π^*: Regime 1 (Bull, left) vs. Regime 2 (Bear, right)

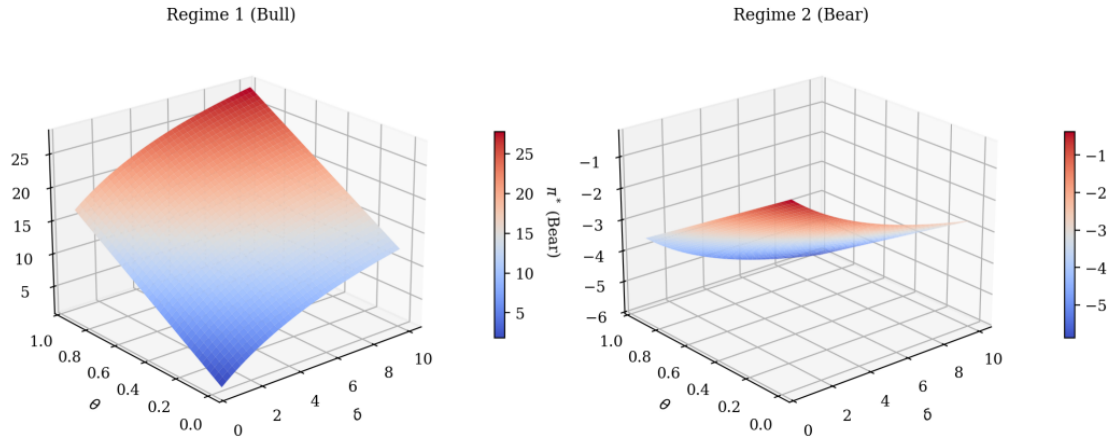

Note: With genuinely empirical inputs, the bear-regime portfolio takes negative positions, because both assets exhibit negative excess returns under the calibrated bear-regime parameters.

The economic intuition is direct. With genuinely empirical inputs the bear regime features negative Sharpe ratios for both assets, so the risk-tolerant equilibrium portfolio holds short positions. This contrasts with calibrations in the prior literature in which bear-regime excess returns are positive but small, yielding small but positive π^* in both regimes. The data tell a sharper story: when market conditions

deteriorate, the risk-adjusted optimum is to bet against the market rather than merely reduce long exposure. A second-order observation is that the regime-dependent surface shapes also differ markedly. The bull-regime surface displays much greater curvature in θ because the high bull-regime Sharpe ratios make the marginal benefit of competitive escalation large; the bear-regime surface is nearly flat because the Sharpe ratios are small in absolute value and competitive amplification has little to amplify.

5.12 Effect of asset correlation and ambiguity on investment

Figure 14 illustrates the joint effect of correlation ρ and ambiguity aversion ψ on π^*. With the empirical correlation $\widehat{\rho} = 0.776$ marking the operating point on the ρ-axis, this figure is particularly informative for understanding how the actual data-implied correlation interacts with the model.

Figure 14: Effect of asset correlation ρ and ambiguity aversion ψ on equilibrium investment π^*

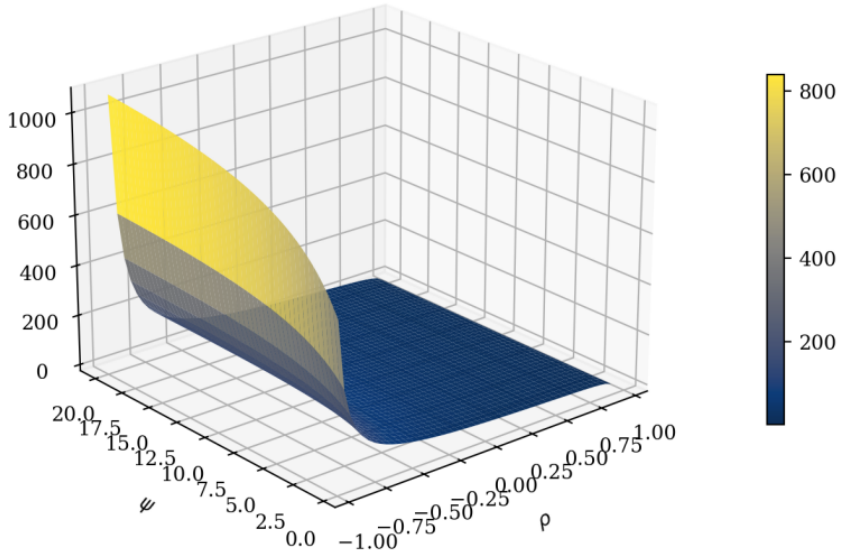

Note: Negative correlation creates diversification benefits that sharply increase total risky investment, particularly for investors with low ambiguity aversion (high ψ). The effect is muted under high ambiguity (low ψ).

The most striking feature is the sharp rise in π^* as ρ decreases from positive towards -1. At $\rho = -0.9$ and $\psi = 20$, π^* exceeds 1,000, reflecting the enormous diversification benefit when assets are nearly perfectly negatively correlated: a combined position has near-zero variance, so the optimal Merton-type portfolio scales up massively. As $\rho \to +1$, Σ becomes near-singular in the same direction and π^* collapses toward small values because the two assets become near-perfect substitutes.

The empirical $\widehat{\rho} \approx 0.78$ places the model in the right-hand portion of the surface, where π^* is comparatively small in magnitude but where the composition of π^* becomes acutely important. With CSI 300's bull-regime Sharpe ratio of 0.77 substantially exceeding CSI 500's 0.22, the inverse covariance matrix Σ^{-1} at $\rho = 0.78$ amplifies the spread between the two assets, prescribing a long position in CSI 300 financed in part by a short position in CSI 500. This long–short spread structure is the distinctive empirical feature of equilibrium portfolios with realistic Chinese equity correlations, and it is invisible in calibrations that assume the much lower $\rho = 0.30$ commonly used in the literature.

The interaction with ambiguity is again instructive. At low ψ (high ambiguity), the entire surface is compressed toward zero and the correlation effect is muted: an ambiguity-averse institution does not exploit diversification aggressively because it fears the estimated correlation may be unstable — a form of robust diversification. At high ψ the institution trusts its correlation estimate and exploits diversification

fully, yielding the dramatic amplification at extreme negative ρ. Model confidence therefore amplifies the portfolio's response to diversification opportunities, with quantitatively dramatic consequences.

6 Conclusion

This paper develops a continuous-time game-theoretic framework in which financial institutions jointly choose portfolio allocations and surplus risk transfer under regime-switching asset returns, ambiguity aversion, and relative performance concerns. In the finite-n game, we derive explicit Nash equilibrium strategies; in the large-population limit, we characterize the corresponding mean field equilibrium in closed form. The analysis yields a clear economic message: competition pushes institutions toward more aggressive risk-taking, while ambiguity aversion works in the opposite direction by reducing effective risk tolerance and strengthening the incentive to transfer risk.

Our results show that relative performance concerns create a powerful amplification mechanism. When institutions care more about keeping pace with their peers, equilibrium investment in risky assets rises sharply. This effect is especially strong when the average degree of competition in the industry, $\bar{\theta}$, increases, because the equilibrium portfolio contains the singular amplification term $1/(1-\bar{\theta})$. The sensitivity analysis confirms that this mechanism is quantitatively large: as competition approaches its critical upper bound, risky investment no longer increases gradually, but escalates in an arms-race fashion. By contrast, ambiguity aversion works as a stabilizing force. By reducing effective risk tolerance, $\delta^{\text{eff}} = \delta\psi/(\delta+\psi)$, ambiguity aversion compresses equilibrium leverage and makes risky investment less sensitive to competitive pressure. The model therefore highlights a fundamental tension in modern finance: the same benchmarking environment that rewards aggressive relative performance can generate system-wide fragility, while model uncertainty, often viewed as a friction, can partially restrain the competitive race for returns.

The empirical calibration and sensitivity analysis sharpen this message. Using CSI 300 and CSI 500 data together with a regime-switching classification of bull and bear markets, the calibrated model shows that competition is not merely a theoretical source of risk-taking; it has large quantitative effects under realistic market parameters. The bull-market equilibrium features positive risky positions that increase strongly with risk tolerance and competition. In the bear regime, however, the calibrated excess returns are negative, so the equilibrium portfolio can switch sign and take short positions rather than simply reducing long exposure. This regime-dependent sign reversal is economically important: it shows that institutional competition does not always produce the same type of leverage across states. In good states, it magnifies long risky exposure; in bad states, it can push risk-tolerant institutions toward aggressive short positions. Financial fragility is therefore state dependent not only in magnitude, but also in direction.

The multi-asset structure further reveals a distinct correlation and model-confidence channel. When asset correlation becomes strongly negative, diversification opportunities expand dramatically and the Merton-type portfolio can scale up sharply, especially when ambiguity aversion is weak. At the empirical correlation estimated from the Chinese equity indices, the model implies a more subtle but equally important mechanism: because the CSI 300 has a stronger bull-regime Sharpe ratio than the CSI 500, the inverse covariance matrix Σ^{-1} amplifies small differences in risk-adjusted returns and generates a leveraged long–short spread between the two assets. This result shows that systemic crowding need not take the form of identical long-only positions. It may also appear through correlated relative-value trades that look diversified at the individual level but become fragile when many institutions trust the same covariance structure.

The risk-transfer margin adds a further layer to this fragility. Competition and high risk tolerance reduce the incentive to transfer surplus risk. Under the baseline institution-side calibration, the equilibrium risk-control policy implies that the representative competitive institution may become a net absorber of risk rather than a net transferor of risk. This finding changes the interpretation of financial stability risk. Systemic vulnerability can build not only through crowded asset holdings, but also through endogenous changes in who bears operational, liability, and balance-sheet risk. Institutions that are highly competitive, highly risk tolerant, and highly confident in their internal models are precisely those most likely to retain or absorb risk in destabilizing ways.

Several policy implications follow. First, regulators should pay close attention to industry-wide benchmarking, league-table evaluation, and tournament-style compensation schemes. The model shows that systemic risk is driven more by the average intensity of relative-performance pressure across the industry than by the aggressiveness of any single institution. Supervisory attention should therefore focus on common incentive structures, benchmark design, and correlated compensation rules rather than only on isolated outliers. Second, prudential policy should be explicitly state contingent. Since competition has much stronger effects in bull markets, macroprudential instruments such as countercyclical capital buffers, leverage limits, liquidity requirements, and margin rules should be tightened during good times, before competitive escalation becomes embedded in balance sheets. Third, model-risk governance should be integrated into financial-stability policy. Greater confidence in estimated returns and correlations can encourage institutions to exploit diversification opportunities more aggressively, so improvements in transparency and model precision should be accompanied by safeguards against excessive leverage, crowded relative-value trades, and overreliance on a common covariance model. Fourth, supervisory stress testing should jointly evaluate portfolio exposure, risk-transfer behavior, regime dependence, and model confidence. Stress tests that consider only asset losses may miss the endogenous shift in risk bearing generated by competition.

Overall, the paper shows that ambiguity aversion, peer competition, regime shifts, asset correlation, and endogenous risk transfer are best understood as jointly determined forces rather than separate margins of choice. A robust theory of financial competition must therefore explain not only how institutions invest, but also how they benchmark themselves, how much confidence they place in their models, how they respond to regime changes, and how they redistribute risk across the system. These interactions are central to understanding leverage, crowding, long–short fragility, and systemic risk in modern financial markets.

References

Evan W. Anderson, Lars Peter Hansen, and Thomas J. Sargent. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association*, 1(1):68–123, 2003a. doi: 10.1162/154247603322256774.

Evan W. Anderson, Lars Peter Hansen, and Thomas J. Sargent. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association*, 1(1):68–123, 2003b. doi: 10.1162/154247603322256774.

Andrew Ang and Geert Bekaert. International asset allocation with regime shifts. *Review of Financial Studies*, 15(4):1137–1187, 2002. doi: 10.1093/rfs/15.4.1137.

Andrew Ang and Allan Timmermann. Regime changes and financial markets. *Annual Review of Financial Economics*, 4:313–337, 2012. doi: 10.1146/annurev-financial-110311-101808.

Pablo Azcue and Nora Muler. Optimal reinsurance and dividend distribution policies in the Cramér–Lundberg model. *Mathematical Finance*, 15(2):261–308, 2005. doi: 10.1111/j.0960-1627.2005.00220.x.

Ioannis Baltas, Anastasios Xepapadeas, and Athanasios N. Yannacopoulos. Robust portfolio decisions for financial institutions. *Journal of Dynamics & Games*, 5(2):61–94, 2018. doi: 10.3934/jdg.2018006.

Suleyman Basak and Dmitry Makarov. Strategic asset allocation in money management. *Journal of Finance*, 69 (1):179–217, 2014. doi: 10.1111/jofi.12106.

Suleyman Basak and Anna Pavlova. Asset prices and institutional investors. *American Economic Review*, 103 (5):1728–1758, 2013. doi: 10.1257/aer.103.5.1728.

Suleyman Basak, Alex Shapiro, and Lucie Teplá. Risk management with benchmarking. *Management Science*, 52(4):542–557, 2006. doi: 10.1287/mnsc.1050.0476.

Suleyman Basak, Anna Pavlova, and Alexander Shapiro. Offsetting the implicit incentives: Benefits of benchmarking in money management. *Journal of Banking and Finance*, 32(9):1883–1893, 2008. doi: 10.1016/j.jbankfin.2007.12.018.

Nicole Branger and Linda Sandris Larsen. Robust portfolio choice with uncertainty about jump and diffusion risk. *Journal of Banking and Finance*, 37(12):5036–5047, 2013. doi: 10.1016/j.jbankfin.2013.08.023.

Keith C. Brown, W. V. Harlow, and Laura T. Starks. Of tournaments and temptations: An analysis of managerial incentives in the mutual fund industry. *Journal of Finance*, 51(1):85–110, 1996. doi: 10.1111/j.1540-6261.1996.tb05203.x.

Sid Browne. Optimal investment policies for a firm with a random risk process: Exponential utility and minimizing the probability of ruin. *Mathematics of Operations Research*, 20(4):937–958, 1995. doi: 10.1287/moor.20.4.937.

Sid Browne. Survival and growth with a liability: Optimal portfolio strategies in continuous time. *Mathematics of Operations Research*, 22(2):468–493, 1997. doi: 10.1287/moor.22.2.468.

Andrea M. Buffa, Dimitri Vayanos, and Paul Woolley. Asset management contracts and equilibrium prices. *Journal of Political Economy*, 130(12):3146–3201, 2022. doi: 10.1086/720515.

Louis K. C. Chan, Stephen G. Dimmock, and Josef Lakonishok. Benchmarking money manager performance: Issues and evidence. *Review of Financial Studies*, 22(11):4553–4599, 2009. doi: 10.1093/rfs/hhp016.

Zengjing Chen and Larry G. Epstein. Ambiguity, risk, and asset returns in continuous time. *Econometrica*, 70 (4):1403–1443, 2002. doi: 10.1111/1468-0262.00337.

Judith Chevalier and Glenn Ellison. Risk taking by mutual funds as a response to incentives. *Journal of Political Economy*, 105(6):1167–1200, 1997. doi: 10.1086/516389.

Peter M. DeMarzo, Ron Kaniel, and Ilan Kremer. Relative wealth concerns and financial bubbles. *Review of Financial Studies*, 21(1):19–50, 2008. doi: 10.1093/rfs/hhm032.

Larry G. Epstein and Martin Schneider. Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31, 2003. doi: 10.1016/S0022-0531(03)00097-8.

Larry G. Epstein and Tan Wang. Intertemporal asset pricing under knightian uncertainty. *Econometrica*, 62(2): 283–322, 1994. doi: 10.2307/2951614.

Gilles-Edouard Espinosa and Nizar Touzi. Optimal investment under relative performance concerns. *Mathematical Finance*, 25(2):221–257, 2015. doi: 10.1111/mafi.12034.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989. doi: 10.1016/0304-4068(89)90018-9.

William N. Goetzmann, Jonathan E. Ingersoll, Matthew Spiegel, and Ivo Welch. Portfolio performance manipulation and manipulation-proof performance measures. *Review of Financial Studies*, 20(5):1503–1546, 2007. doi: 10.1093/rfs/hhm025.

Guohui Guan and Xiang Hu. Time-consistent investment and reinsurance strategies for mean-variance insurers in n-agent and mean-field games. *North American Actuarial Journal*, 26(4):537–569, 2022. doi: 10.1080/10920277.2022.2050260.

Veronica Guerrieri and Péter Kondor. Fund managers, career concerns, and asset price volatility. *American Economic Review*, 102(5):1986–2017, 2012. doi: 10.1257/aer.102.5.1986.

Massimo Guidolin and Allan Timmermann. Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 31(11):3503–3544, 2007. doi: 10.1016/j.jedc.2006.12.004.

James D. Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384, 1989. doi: 10.2307/1912559.

James D. Hamilton. Analysis of time series subject to changes in regime. *Journal of Econometrics*, 45(1–2): 39–70, 1990. doi: 10.1016/0304-4076(90)90093-9.

James D. Hamilton and Gang Lin. Stock market volatility and the business cycle. *Journal of Applied Econometrics*, 11(5):573–593, 1996. doi: 10.1002/(SICI)1099-1255(199609)11:5⟨573::AID-JAE413⟩3.0.CO;2-T.

Lars Peter Hansen and Thomas J. Sargent. Robust control and model uncertainty. *American Economic Review*, 91(2):60–66, 2001. doi: 10.1257/aer.91.2.60.

Lars Peter Hansen and Thomas J. Sargent. Fragile beliefs and the price of uncertainty. *Quantitative Economics*, 1(1):129–162, 2010. doi: 10.3982/QE9.

Lars Peter Hansen, Thomas J. Sargent, Gauhar Turmuhambetova, and Noah Williams. Robust control and model misspecification. *Journal of Economic Theory*, 128(1):45–90, 2006. doi: 10.1016/j.jet.2004.12.006.

Minyi Huang, Peter E. Caines, and Roland P. Malhamé. Large population stochastic dynamic games: Closed-loop McKean-Vlasov systems and the nash certainty equivalence principle. *Communications in Information and Systems*, 6(3):221–252, 2006. doi: 10.4310/CIS.2006.v6.n3.a5.

Mohammad R. Jahan-Parvar and Hening Liu. Ambiguity aversion and asset prices in production economies. *Review of Financial Studies*, 27(10):3060–3097, 2014. doi: 10.1093/rfs/hhu037.

Daehee Jeong, Hwagyun Kim, and Joon Y. Park. Does ambiguity matter? estimating asset pricing models with a multiple-priors recursive utility. *Journal of Financial Economics*, 115(2):361–382, 2015. doi: 10.1016/j.jfineco.2014.10.003.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005. doi: 10.1111/j.1468-0262.2005.00640.x.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144(3):930–976, 2009. doi: 10.1016/j.jet.2008.10.007.

Daniel Lacker and Thaleia Zariphopoulou. Mean field and n-agent games for optimal investment under relative performance criteria. *Mathematical Finance*, 29:1003–1038, 2019a. doi: 10.1111/mafi.12206.

Daniel Lacker and Thaleia Zariphopoulou. Mean field and n-agent games for optimal investment under relative performance criteria. *Mathematical Finance*, 29(4):1003–1038, 2019b. doi: 10.1111/mafi.12206.

Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007. doi: 10.1007/s11537-007-0657-8.

Pascal J. Maenhout. Robust portfolio rules and asset pricing. *Review of Financial Studies*, 17(4):951–983, 2004a. doi: 10.1093/rfs/hhh003.

Pascal J. Maenhout. Robust portfolio rules and asset pricing. *The Review of Financial Studies*, 17(4):951–983, 2004b. doi: 10.1093/rfs/hhh003.

Pascal J. Maenhout. Robust portfolio rules and detection-error probabilities for a mean-reverting risk premium. *Journal of Economic Theory*, 128(1):136–163, 2006. doi: 10.1016/j.jet.2005.12.012.

Huyèn Pham, Xiaoli Wei, and Chao Zhou. Portfolio diversification and model uncertainty: A robust dynamic mean-variance approach. *Mathematical Finance*, 32(1):349–404, 2022. doi: 10.1111/mafi.12320.

Bryan R. Routledge and Stanley E. Zin. Model uncertainty and liquidity. *Review of Economic Dynamics*, 12(4): 543–566, 2009. doi: 10.1016/j.red.2008.10.002.

Yonggan Zhao. A dynamic model of active portfolio management with benchmark orientation. *Journal of Banking and Finance*, 31(11):3336–3356, 2007. doi: 10.1016/j.jbankfin.2007.04.007.