

Mean Field Games and Global Arms Races: Strategic Dynamics in a Multipolar World

Heng-Fu Zou

May 15, 2025

Abstract

This paper develops a dynamic game-theoretic framework to model global arms races in a multipolar world using mean field game (MFG) theory. We analyze the strategic behavior of a continuum of minor countries influenced by the military decisions of three major powers—the United States, China, and Russia—who engage in a finite-player differential game. Each country chooses its military expenditure over time to minimize a cost function that reflects internal costs and strategic positioning relative to others. We derive both general nonlinear and linear-quadratic-Gaussian (LQG) formulations, solve the coupled HJB–FPK systems, and simulate both time-dependent and stationary equilibria. Our results show how strategic interdependence, peer pressure, and deterrence incentives drive excessive militarization in decentralized equilibrium. We compare decentralized and centralized outcomes and analyze policy interventions such as caps and taxes. The framework offers a rigorous foundation for understanding military competition and evaluating arms control policies under uncertainty.

Key Words: Mean Field Games, Arms Race, Strategic Competition, Militarization, Linear-Quadratic-Gaussian Control, Decentralized Equilibrium

1 Introduction

The strategic accumulation of military power is one of the most enduring and consequential dynamics in international relations. From the dreadnought races of the early twentieth century to the nuclear competition of the Cold War, and now to the emergent tensions over artificial intelligence, cyberwarfare, and missile defense systems, the logic of the arms race continues to shape global security architectures. Today, the rivalry between the United States, China, and Russia has reawakened the dynamics of militarized competition on a global scale. Their strategic postures not only affect their own security, but also reverberate across the international system, prompting reactive behavior among smaller states and reshaping the distribution of military capabilities worldwide.

At the core of these dynamics lies a problem of strategic interdependence under uncertainty. No country determines its military stockpile in isolation. Instead, each faces a forward-looking optimization problem under incomplete information: it must decide how much to invest in military power based on its own perceived threats and objectives, its expectations about the future actions of others, and the costs associated with falling behind strategically. When many countries simultaneously respond to a rising militarization frontier—often defined by a small number of dominant states—the result is a self-reinforcing cycle of escalation, even when no actor desires open conflict. This phenomenon is both a collective action failure and a systemic feature of an anarchic international order.

This paper proposes a dynamic game-theoretic framework to study such arms races using the tools of mean field game (MFG) theory. MFG theory provides a scalable approach for modeling strategic interactions among many agents who are individually negligible but collectively impactful. First developed by Lasry and Lions (2006) and Huang, Caines, and Malhamé (2006), MFG theory merges optimal control, game theory, and stochastic analysis into a unified system that is both analytically tractable and empirically interpretable. In this setting, each agent optimizes over a stochastic trajectory, taking the distribution of the population as given; equilibrium requires that this distribution is consistent with the agents' joint behavior.

We extend this approach to model a multipolar arms race, with explicit heterogeneity across players. Specifically, we consider a system comprising three major powers—the United States, China, and Russia—and a continuum of minor countries. The major powers interact strategically with one another in a finite-player dynamic game, while the minor states respond to the evolving mean field generated by their peers and the aggregate behavior of the dominant powers. Each country chooses its military expenditure path to minimize a cost functional that reflects three key factors: (i) the cost of maintaining a high military stock, (ii) the cost of exerting high military effort, and (iii) the strategic cost of falling behind the global norm.

We begin by developing the general nonlinear mean field game system, consisting of a Hamilton–Jacobi–Bellman (HJB) equation for individual value functions, a Fokker–Planck–Kolmogorov (FPK) equation for the evolution of the population state distribution, and a best response control law linked through a McKean–Vlasov fixed-point condition. We then specialize to a linear-quadratic-Gaussian (LQG) framework, where optimal feedback controls and equilibrium distributions can be derived in closed form. This allows us to simulate the system and provide comparative statics with respect to key parameters, including peer pressure, discount rates, and initial conditions.

Our simulations yield several key findings. First, the military stocks of the three major powers converge dynamically toward a stable equilibrium that reflects their strategic competition and desire for deterrence dominance. Second, the continuum of minor players gradually aligns its military levels with the dominant powers, leading to a rising and persistent global mean field of militarization. Third, increasing peer pressure among minor countries leads to tighter

clustering of military stocks and a rightward shift in the distribution—indicating higher militarization and reduced diversity. Finally, we show that strategic inefficiencies arise from the decentralized nature of the game: countries overinvest in military power relative to a centralized solution that internalizes global externalities.

The policy implications are far-reaching. If left unregulated, decentralized arms races generate path-dependent equilibria marked by over-militarization and systemic vulnerability. However, coordinated action—such as arms control agreements among major powers or penalties on excessive militarization—can alter the feedback mechanisms in the system, shifting the equilibrium toward a more peaceful and efficient outcome. Our model allows for quantitative evaluations of such interventions.

In sum, this paper develops a novel and rigorous framework for understanding the endogenous dynamics of global arms races in a multipolar world. By combining the realism of strategic heterogeneity with the tractability of mean field analysis, we provide both theoretical and empirical insights into the causes, consequences, and control of militarized competition in the twenty-first century.

The remainder of the paper unfolds in nine sections, each progressively building the theoretical and empirical foundation of our model. In Section 2, we formulate the mean field game in two stages. First, we construct the most general nonlinear model for a continuum of agents, capturing complex, dynamic interdependence under stochastic uncertainty. We then specialize this general formulation to the linear-quadratic-Gaussian (LQG) setting, where the linearity of state dynamics, quadratic costs, and Gaussian noise structure permit explicit solutions for value functions, optimal controls, and population distributions. This LQG framework forms the backbone of our analytical and numerical exploration.

Section 3 presents the detailed derivation and solution of both the time-dependent (finite-horizon) and stationary (infinite-horizon) versions of the model. We solve the forward–backward system of Hamilton–Jacobi–Bellman (HJB) and Fokker–Planck–Kolmogorov (FPK) equations and examine the dynamic feedback loop between individual strategies and the evolving population distribution. In the stationary case, we analyze the long-run equilibrium configuration under constant strategic environments and discounting.

In Section 4, we conduct numerical simulations of the time-dependent model, highlighting the evolution of military stocks over time. We show how the strategic arms race unfolds from different initial conditions, and how countries’ military postures converge toward stable equilibrium levels. The simulations illustrate key dynamics, such as escalation, convergence, and diffusion of strategic norms, across heterogeneous agents under endogenous feedback.

Section 5 introduces policy interventions into the mean field game framework. We analyze three types of instruments: (i) penalty functions for excessive militarization, (ii) hard caps on military stock levels, and (iii) fiscal tools such as taxes or subsidies on military expenditure. Each policy alters the optimization structure of the agents, leading to modified HJB–FPK dynamics. We evaluate the comparative effectiveness of these policies in reducing average militarization,

controlling variance, and improving overall welfare.

In Section 6, we contrast the decentralized Nash equilibrium of the mean field game with the centrally planned solution in which a social planner minimizes total global cost. By internalizing the externalities that arise from mutual fear and relative positioning, the planner selects a coordinated path that achieves lower levels of militarization and lower strategic expenditure. This section provides a benchmark for evaluating the efficiency losses due to decentralized decision-making in the arms race.

Section 7 extends the model to reflect the real-world heterogeneity of global power. We develop a hybrid model with three major players—United States, China, and Russia—engaged in a finite-player game, while the rest of the world consists of a continuum of minor countries solving a mean field game in response to the aggregate behavior of the major powers. This structure captures asymmetric influence, hierarchical interactions, and the propagation of strategic norms. We derive the system of coupled HJB–FPK equations and feedback laws for both major and minor players.

Section 8 presents simulation results for the multipolar world model, including both transient and stationary analyses. We simulate the evolution of military stocks among major powers and track how their behavior induces gradual militarization among minor states. We visualize the emergence of a global strategic equilibrium and analyze the distributional implications of increased peer pressure, asymmetries in initial conditions, and changes in discount rates. These simulations demonstrate the contagious logic of militarization and the interdependence of strategic incentives across levels of power.

Finally, Section 9 concludes by synthesizing the main results and offering a policy-oriented interpretation of the model. We highlight the dangers of decentralized arms races in a world of rising multipolar competition and identify key levers for coordination and intervention. We also discuss extensions and future directions for modeling strategic interactions in international security using mean field and hybrid game frameworks.

2 General Mean Field Game Model: Nonlinear Framework

In this section, we construct a general continuous-time mean field game (MFG) model in a one-dimensional state space. We begin by considering a representative agent from a continuum of indistinguishable agents (such as countries), each making decisions over time in response to their individual states and the distribution of the population. The state dynamics are governed by a stochastic differential equation (SDE), and each agent minimizes a cost that depends on both their own trajectory and the aggregate population distribution.

Let $x(t) \in \mathbb{R}$ denote the state variable of a representative agent at time $t \in [0, T]$. In the arms race context, $x(t)$ represents the level of military stock or capability. Let $u(t) \in U \subset \mathbb{R}$ denote the control variable, representing an

action such as military expenditure.

The agent's controlled dynamics follow the stochastic differential equation:

$$dx(t) = f(x(t), u(t), t) dt + \sigma(x(t), u(t), t) dW(t),$$

where:

- $f(x, u, t)$: drift term — deterministic trend in the dynamics.
- $\sigma(x, u, t)$: diffusion term — uncertainty or stochastic volatility.
- $W(t)$: standard Brownian motion.
- $u(t)$: control variable chosen by the agent.

The agent chooses a control process $u(\cdot)$ to minimize a cost functional of the form:

$$J(u(\cdot)) = \mathbb{E} \left[\int_0^T L(x(t), u(t), \mu(t), t) dt + G(x(T), \mu(T)) \right],$$

where:

- $L(x, u, \mu, t)$: running cost depending on the state, control, and population distribution $\mu(t)$.
- $G(x, \mu)$: terminal cost.
- $\mu(t)$: the distribution of states across the population at time t , i.e., $\mu(t) = \text{Law}(x(t))$.

This cost functional introduces mean field dependence through $\mu(t)$, capturing the externality that each country's behavior depends on the average militarization level.

2.1 Hamilton–Jacobi–Bellman (HJB) Equation

The HJB equation characterizes the value function $V(t, x)$, which represents the minimal expected cost for an agent starting at state x and time t . That is:

$$V(t, x) = \inf_{u(\cdot)} \mathbb{E} \left[\int_t^T L(x(s), u(s), \mu(s), s) ds + G(x(T), \mu(T)) \mid x(t) = x \right].$$

Using dynamic programming and Itô's lemma, we derive the Hamilton–Jacobi–Bellman (HJB) equation:

$$\frac{\partial V}{\partial t}(t, x) + \inf_{u \in U} \left\{ f(x, u, t) \frac{\partial V}{\partial x}(t, x) + \frac{1}{2} \sigma^2(x, u, t) \frac{\partial^2 V}{\partial x^2}(t, x) + L(x, u, \mu(t), t) \right\} = 0,$$

with terminal condition:

$$V(T, x) = G(x, \mu(T)).$$

This is a backward PDE solved from $t = T$ to $t = 0$. The term inside the infimum is the Hamiltonian $H(x, u, t, \mu(t), \partial_x V, \partial_{xx} V)$.

2.2 Best Response Control

The best response (BR) or optimal feedback control $u^*(t, x)$ is defined as the argument minimizing the Hamiltonian:

$$u^*(t, x) = \arg \min_{u \in U} \left\{ f(x, u, t) \frac{\partial V}{\partial x}(t, x) + \frac{1}{2} \sigma^2(x, u, t) \frac{\partial^2 V}{\partial x^2}(t, x) + L(x, u, \mu(t), t) \right\}.$$

This gives the feedback law $u^*(t, x; \mu(t))$, i.e., the optimal decision as a function of state and the mean field.

2.3 Fokker–Planck–Kolmogorov (FPK) Equation

Given the best response control $u^*(t, x)$, the evolution of the population distribution $p(t, x)$ over time is governed by the Fokker–Planck–Kolmogorov (FPK) equation, a forward PDE:

$$\frac{\partial p}{\partial t}(t, x) = -\frac{\partial}{\partial x} (f(x, u(t, x), t)p(t, x)) + \frac{1}{2} \frac{\partial^2}{\partial x^2} (\sigma^2(x, u(t, x), t)p(t, x)),$$

with initial condition:

$$p(0, x) = p_0(x),$$

where $p_0(x)$ is the initial distribution of agent states.

2.4 Coupling and Fixed Point: McKean–Vlasov Structure

The key feature of mean field games is the coupling between individual optimization and population dynamics. The optimal control $u(t, x)$ depends on the population distribution $\mu(t)$, while the evolution of $\mu(t)$ depends on u . This mutual dependency creates a fixed-point problem.

This is often interpreted through a McKean–Vlasov SDE:

$$dx(t) = f(x(t), u(t, x(t); \mu(t)), t)dt + \sigma(x(t), u(t, x(t); \mu(t)), t)dW(t), \quad \mu(t) = \text{Law}(x(t)).$$

A mean field equilibrium is a pair (u^*, μ) such that:

- $u^*(t, x)$ solves the HJB equation given $\mu(t)$,
- $\mu(t)$ evolves under the FPK equation with $u^*(t, x)$.

2.5 Specialization: Linear-Quadratic Mean Field Game

We now specialize to the linear-quadratic (LQ) setting, which is analytically tractable and ideal for modeling strategic arms races.

State Dynamics:

$$dx(t) = (ax(t) + bu(t))dt + \sigma dW(t),$$

where a is the natural decay or depreciation of military stock and b is the effectiveness of military expenditure.

Cost Functional:

$$J(u) = \mathbb{E} \left[\int_0^T \left(\frac{q}{2} x(t)^2 + \frac{r}{2} u(t)^2 + \frac{\eta}{2} (x(t) - \bar{x}(t))^2 \right) dt + \frac{q_T}{2} x(T)^2 \right],$$

where $\bar{x}(t) = \int x p(t, x) dx$ is the mean field (average militarization level).

Quadratic Value Function Ansatz:

We guess:

$$V(t, x) = \frac{1}{2} P(t) x^2 + R(t) x + S(t),$$

which leads to the optimal control:

$$u^*(t, x) = -\frac{b}{r} (P(t) x + R(t)).$$

Plugging this into the HJB yields a system of Riccati ODEs for $P(t)$, $R(t)$, $S(t)$, and the FPK equation simplifies to a linear PDE with drift and diffusion defined by this optimal feedback.

The LQ case allows explicit solutions, which we exploit in later sections to perform simulations and policy comparisons.

3 Time-Dependent and Stationary Solutions with McKean–Vlasov Coupling

We now examine the full dynamic solution of the mean field game (MFG) system introduced in Section 2, beginning with the general nonlinear formulation and then specializing to the linear-quadratic-Gaussian (LQG) case. For both the general and specialized settings, we distinguish between the finite-horizon time-dependent case, where the optimization problem spans a fixed interval $[0, T]$, and the infinite-horizon stationary case, where solutions are time-invariant and reflect long-run strategic behavior.

We proceed in four steps:

1. General nonlinear time-dependent MFG system.
2. General nonlinear stationary MFG system.
3. LQG time-dependent solution with Riccati ODEs.
4. LQG stationary solution with algebraic Riccati equations.

3.1 General Nonlinear Time-Dependent MFG System

Consider a representative agent whose scalar state $x(t) \in \mathbb{R}$ evolves according to a controlled stochastic differential equation:

$$dx(t) = f(x(t), u(t), t) dt + \sigma(x(t), u(t), t) dW(t),$$

where $u(t) \in U \subset \mathbb{R}$ is the control, and $W(t)$ is a standard Brownian motion.

The objective is to minimize the cost functional:

$$J(u(\cdot)) = \mathbb{E} \left[\int_0^T L(x(t), u(t), \mu(t), t) dt + G(x(T), \mu(T)) \right],$$

where $\mu(t) = \text{Law}(x(t))$ is the time-varying distribution of the population.

The value function is defined as:

$$V(t, x) = \inf_{u(\cdot)} \mathbb{E} \left[\int_t^T L(x(s), u(s), \mu(s), s) ds + G(x(T), \mu(T)) \mid x(t) = x \right].$$

Applying the dynamic programming principle and Itô's lemma, we obtain the Hamilton–Jacobi–Bellman (HJB) equation:

$$\frac{\partial V}{\partial t}(t, x) + \inf_{u \in U} \left\{ f(x, u, t) \frac{\partial V}{\partial x}(t, x) + \frac{1}{2} \sigma^2(x, u, t) \frac{\partial^2 V}{\partial x^2}(t, x) + L(x, u, \mu(t), t) \right\} = 0,$$

with terminal condition $V(T, x) = G(x, \mu(T))$.

The optimal control $u^*(t, x)$ is obtained by minimizing the integrand of the HJB equation:

$$u^*(t, x) = \arg \min_{u \in U} \left\{ f(x, u, t) \frac{\partial V}{\partial x}(t, x) + \frac{1}{2} \sigma^2(x, u, t) \frac{\partial^2 V}{\partial x^2}(t, x) + L(x, u, \mu(t), t) \right\}.$$

Given this optimal control, the population distribution $p(t, x)$ evolves under the Fokker–Planck–Kolmogorov (FPK) equation:

$$\frac{\partial p}{\partial t}(t, x) = -\frac{\partial}{\partial x} (f(x, u(t, x), t)p(t, x)) + \frac{1}{2} \frac{\partial^2}{\partial x^2} (\sigma^2(x, u(t, x), t)p(t, x)),$$

with initial condition $p(0, x) = p_0(x)$.

The MFG system is closed through the fixed-point consistency condition:
 $\mu(t) = p(t, x) dx$, and $u(t, x) = u(t, x; \mu(t))$.

3.2 General Nonlinear Stationary MFG System

For long-run analysis, we consider an infinite-horizon version with discount factor $\rho > 0$. The optimization objective becomes:

$$J(u(\cdot)) = \mathbb{E} \left[\int_0^\infty e^{-\rho t} L(x(t), u(t), \mu(t)) dt \right].$$

Assuming time-invariant coefficients and a stationary distribution μ , we seek a value function $V(x)$ satisfying the stationary HJB equation:

$$\rho V(x) = \inf_{u \in U} \left\{ f(x, u) V'(x) + \frac{1}{2} \sigma^2(x, u) V''(x) + L(x, u, \mu) \right\}.$$

The optimal stationary control $u(x)$ is then a function of μ , and the invariant distribution $p(x)$ solves the stationary FPK equation:

$$0 = -\frac{d}{dx} (f(x, u(x))p(x)) + \frac{1}{2} \frac{d^2}{dx^2} (\sigma^2(x, u(x))p(x)),$$

subject to normalization $\int p(x) dx = 1$. The equilibrium requires:
 $\mu = p(x) dx$, and $u(x) = u(x; \mu)$.

This formulation establishes the stationary mean field equilibrium as a fixed point.

3.3 LQG Time-Dependent Mean Field Game

We now specialize to the linear-quadratic-Gaussian (LQG) case. The agent's dynamics are:

$$dx(t) = [ax(t) + bu(t)]dt + \sigma dW(t),$$

and the cost functional is:

$$J = \mathbb{E} \left[\int_0^T \left(\frac{q}{2} x(t)^2 + \frac{r}{2} u(t)^2 + \frac{\eta}{2} (x(t) - \bar{x}(t))^2 \right) dt \right].$$

We conjecture a quadratic value function:

$$V(t, x) = \frac{1}{2} P(t)x^2 + R(t)x + S(t),$$

and derive the optimal control:

$$u^*(t, x) = -\frac{b}{r} (P(t)x + R(t)).$$

The coefficients $P(t), R(t), S(t)$ solve:

$$\dot{P}(t) = -2aP(t) + \frac{b^2}{r} P(t)^2 - q - \eta,$$

$$\dot{R}(t) = -\left(a - \frac{b^2}{r} P(t)\right) R(t) + \eta \bar{x}(t),$$

$$\dot{S}(t) = -\frac{b^2}{2r} R(t)^2 + \frac{\sigma^2}{2} P(t) - \frac{\eta}{2} \bar{x}(t)^2.$$

Under this control, the state process follows:

$$dx(t) = \left[\left(a - \frac{b^2}{r} P(t)\right) x(t) - \frac{b^2}{r} R(t) \right] dt + \sigma dW(t),$$

and the FPK equation becomes:

$$\frac{\partial p}{\partial t}(t, x) = -\frac{\partial}{\partial x} \left[\left(a - \frac{b^2}{r} P(t)\right) x - \frac{b^2}{r} R(t) \right] p(t, x) + \frac{\sigma^2}{2} \frac{\partial^2 p}{\partial x^2}.$$

3.4 LQG Stationary Mean Field Game and McKean–Vlasov Equilibrium

In the stationary case, the algebraic Riccati equations become:

$$\rho P + 2aP - \frac{b^2}{r}P^2 + q + \eta = 0,$$

$$\rho R + \left(a - \frac{b^2}{r}P\right)R - \eta\bar{x} = 0.$$

The optimal control is:

$$u^*(x) = -\frac{b}{r}(Px + R),$$

and the FPK equation becomes:

$$0 = -\frac{d}{dx} \left(\left[\left(a - \frac{b^2}{r}P\right)x - \frac{b^2}{r}R \right] p(x) \right) + \frac{\sigma^2}{2} \frac{d^2 p}{dx^2}.$$

Its solution is:

$$p(x) = C \exp \left(\frac{2}{\sigma^2} \left(\frac{1}{2} \left(a - \frac{b^2}{r}P\right)x^2 - \frac{b^2}{r}Rx \right) \right),$$

with $\int p(x) dx = 1$. The fixed point requires:

$$\bar{x} = \int x p^*(x) dx,$$

matching the value used in R.

This stationary solution, characterized by self-consistent expectations and behavior, is the McKean–Vlasov equilibrium of the arms race system.

4 Simulations and Interpretation of Strategic Military Equilibrium

In this section, we investigate how the dynamic and stationary solutions to the MFG system translate into military equilibrium behavior across countries. Our focus is the linear-quadratic-Gaussian (LQG) case, for which the system of Riccati equations and the Fokker–Planck–Kolmogorov (FPK) equation derived in Section 3 can be solved either analytically or numerically. We analyze both transient dynamics over a finite horizon and stationary equilibrium behavior in the infinite-horizon setting. We simulate the MFG dynamics under a variety of global conditions and interpret how different forces—such as peer pressure, long-term planning, and international norms—shape the collective behavior of countries in a strategic arms race.

To begin, recall that in the stationary case, the optimal control policy for each country is given by:

$$u^*(x) = -\frac{b}{r}(Px + R),$$

where P and R solve the algebraic Riccati system:

$$\rho P + 2aP - \frac{b^2}{r}P^2 + q + \eta = 0,$$

$$\rho R + \left(a - \frac{b^2}{r}P\right)R - \eta\bar{x} = 0.$$

These coefficients determine how aggressively countries respond to their own military capital level x , and how much they adjust in response to the global av-

erage $\bar{x}$. The steady-state distribution $p^*(x)$ of military capital across countries satisfies the stationary FPK equation:

$$0 = -\frac{d}{dx} \left(\left[\left(a - \frac{b^2}{r} P \right) x - \frac{b^2}{r} R \right] p(x) \right) + \frac{\sigma^2}{2} \frac{d^2}{dx^2} p(x).$$

This is a linear second-order ordinary differential equation with a closed-form solution. Let us denote the drift term by:

$$\mu(x) = \left(a - \frac{b^2}{r} P \right) x - \frac{b^2}{r} R.$$

Then the FPK equation becomes:

$$0 = -\frac{d}{dx} (\mu(x)p(x)) + \frac{\sigma^2}{2} \frac{d^2 p}{dx^2}.$$

This type of equation has a known solution via integrating factor method.

First, we write the stationary density as:

$$p^*(x) = C \exp \left(\frac{2}{\sigma^2} \int^x \mu(z) dz \right),$$

where C is the normalization constant ensuring $\int p^*(x) dx = 1$. In our case, since $\mu(x)$ is affine, the integral becomes:

$$\int^x \mu(z) dz = \left(\frac{1}{2} \left(a - \frac{b^2}{r} P \right) x^2 - \frac{b^2}{r} R x \right),$$

which yields:

$$p^*(x) = C \exp \left(\frac{2}{\sigma^2} \left[\frac{1}{2} \left(a - \frac{b^2}{r} P \right) x^2 - \frac{b^2}{r} R x \right] \right).$$

This expression describes a shifted and scaled Gaussian-like distribution of military stock. The shape of this distribution depends critically on the coefficients P and R , which in turn depend on the parameters ρ (discount rate), η (peer pressure), and $\bar{x}$ (global military norm).

To interpret this distribution, we compute the mean and variance:

$$\bar{x} = \int x p(x) dx, \quad \text{Var}(x) = \int (x - \bar{x})^2 p(x) dx.$$

These moments describe the long-run military equilibrium across countries. We find that the mean $\bar{x}$ generally exceeds the social planner's optimum due to the decentralized externalities embedded in the peer pressure term $\eta(x - \bar{x})^2$. Each country wants to avoid falling behind, which pushes the entire distribution rightward.

To better understand how equilibrium changes under global conditions, we simulate three comparative statics:

1. Higher peer pressure ($\eta \uparrow$): The coefficient P increases, which steepens the cost associated with falling behind. This leads countries to invest more heavily in military capital, shifting the distribution $p^*(x)$ rightward. The peak of the distribution flattens and variance increases, indicating more dispersion across countries.

2. Lower discount rate ($\rho \downarrow$): Countries value future security more, and thus tolerate higher up-front costs to maintain long-term military advantage. This increases P and marginally increases R , which again shifts the distribution rightward but less drastically than in the high-peer-pressure case.

3. Changes in perceived global norm $\bar{x}$: Since the optimal control and the Riccati coefficients depend directly on $\bar{x}$, raising the expected average militarization level pushes R downward (more negative), which intensifies military spending in the control formula. The entire distribution $p^*(x)$ moves right. Conversely, lowering $\bar{x}$ to simulate a more pacified global environment significantly compresses the distribution toward lower values.

Time-dependent simulations (finite horizon) show similar trends. Beginning from an initial distribution (e.g., concentrated at low x), the FPK equation gradually spreads the mass over time. As agents follow the best response dynamics derived from the HJB equation, the population's distribution converges toward the stationary equilibrium. These transient dynamics are critical for understanding how rapidly an arms race escalates or how quickly demilitarization takes hold under a new norm.

In all cases, the joint behavior of countries converges to a distribution that reflects their incentives, expectations, and exposure to uncertainty. Importantly, because the MFG model is based on symmetric and decentralized decision-making, the equilibrium can lead to over-investment in military power due to mutual imitation and the tragedy of strategic expectations. These insights highlight the necessity for coordinated policy efforts to shift the mean field (through treaties or norms) or directly alter strategic incentives (via taxes, penalties, or caps).

5 Policy Interventions and Their Effects on Strategic Military Equilibrium

Given the potential for self-reinforcing over-militarization highlighted in Section 4, it becomes important to study whether and how effective policy instruments can steer the global equilibrium toward more peaceful outcomes. In this section, we examine the formal incorporation of policy interventions into the mean field game (MFG) framework. We explore three classes of interventions: (i) penalty functions that discourage excessive militarization, (ii) military caps that impose hard constraints on the state variable, and (iii) subsidies or taxes that alter the cost of control (i.e., military spending). We analyze each intervention both mathematically and strategically by modifying the HJB-FPK system accordingly and examining the resulting changes to the equilibrium distribution.

Let us denote the baseline individual cost functional as:

$$J(u(\cdot)) = \mathbb{E} \left[\int_0^T \left(\frac{q}{2} x(t)^2 + \frac{r}{2} u(t)^2 + \frac{\eta}{2} (x(t) - \bar{x}(t))^2 \right) dt + \frac{q_T}{2} x(T)^2 \right],$$

with dynamics:

$$dx(t) = (ax(t) + bu(t))dt + \sigma dW(t).$$

To model penalties for excessive militarization, we modify the running cost by adding a convex penalty function $\phi(x(t))$ that rises rapidly beyond a threshold $x_{\max}$. A simple choice is a quadratic penalty:

$$\phi(x) = \frac{\lambda}{2} (x - x_{\max})^2 \cdot \mathbf{1}_{x > x_{\max}},$$

where $\lambda > 0$ controls the severity. The new cost functional becomes:

$$J_{\text{pen}}(u) = \mathbb{E} \left[\int_0^T \left(\frac{q}{2} x^2 + \frac{r}{2} u^2 + \frac{\eta}{2} (x - \bar{x})^2 + \phi(x) \right) dt + \frac{q_T}{2} x(T)^2 \right].$$

In the HJB equation, this adds $\phi(x)$ directly to the Hamiltonian. Since the penalty term is convex and only activated for $x > x_{\max}$, the optimal control $u^*(x)$ becomes more conservative in this region. The effect on the distribution $p(t, x)$ is a leftward truncation and compression: countries are deterred from moving beyond the penalized threshold.

Alternatively, one may impose a hard cap on the state variable, requiring $x(t) \leq x_{\text{cap}}$. Mathematically, this creates a reflecting barrier or boundary condition in the FPK equation. In the stationary case, we replace the domain $\mathbb{R} + \text{with}[0, x_{\text{cap}}]$, and the stationary FPK equation becomes:

$$0 = -\frac{d}{dx} (\mu(x)p(x)) + \frac{\sigma^2}{2} \frac{d^2}{dx^2} p(x),$$

with boundary conditions:

$$p(0) = 0, \quad \left(\mu(x)p(x) - \frac{\sigma^2}{2} \frac{d}{dx} p^*(x) \right) \Big|_{x=x_{\text{cap}}} = 0.$$

This forces all countries' military levels to remain within a legal bound. The result is a truncated equilibrium distribution with sharp mass concentration near the cap, especially if peer pressure η remains strong. Importantly, hard caps may generate undesirable clustering near x_{cap} , signaling strategic risk if countries attempt to "touch the ceiling" without exceeding it.

A third intervention modifies the cost of military expenditure. Suppose a supranational organization subsidizes disarmament or taxes military buildup. Let the new cost of control become:

$$\frac{r(1+\tau)}{2} u^2,$$

where $\tau > 0$ is a tax and $\tau < 0$ is a subsidy. Then, the optimal control derived from minimizing the HJB Hamiltonian becomes:

$$u^*(x) = -\frac{b}{r(1+\tau)} (Px + R),$$

and the drift in the FPK equation adjusts accordingly:

$$\mu(x) = \left(a - \frac{b^2}{r(1+\tau)} P \right) x - \frac{b^2}{r(1+\tau)} R.$$

We now observe that taxation increases the effective cost of control and thus reduces the responsiveness of countries to their current state, flattening the control response. The overall distribution $p^*(x)$ shifts leftward and becomes more peaked at low values. Conversely, subsidies reduce the effective cost of control, possibly inducing more militarization unless targeted at low states x only.

These interventions can be compared by examining how they alter the Riccati system. In the case of a tax $\tau > 0$, the Riccati equation becomes:

$$\rho P + 2aP - \frac{b^2}{r(1+\tau)} P^2 + q + \eta = 0,$$

which clearly shows that increasing τ (i.e., stronger taxation) reduces the coefficient on the negative quadratic term. This leads to a lower equilibrium value of P , meaning a reduced marginal penalty on high x . Paradoxically, too much taxation might reduce the slope of the control function and cause countries to respond less forcefully to increases in their state, thus slowing disarmament unless peer pressure or penalties are increased correspondingly.

In all cases, the interventions modify the equilibrium not by directly controlling the distribution, but by reshaping incentives and indirectly steering collective behavior. The strength of the MFG framework lies in its ability to represent such strategic adjustments endogenously. Policymakers can, through appropriate parameter adjustments (e.g., through treaties, incentive schemes, or regulations), shift the Nash equilibrium of the game toward more globally optimal outcomes.

6 Comparison Between Decentralized Mean Field Equilibria and the Central Planner Solution

In this section, we formally analyze the divergence between decentralized equilibrium outcomes in a mean field game (MFG) and the socially optimal allocation of resources under a centralized planner. This comparison is particularly illuminating in the context of arms races, where individual incentives drive countries to maintain high levels of military capital, even when global security and welfare might be improved through restraint and coordination. We begin by characterizing the planner's problem and deriving the optimality conditions. We then contrast the planner's solution to the decentralized MFG outcome and quantify the structural inefficiencies that emerge.

We maintain the linear-quadratic-Gaussian (LQG) setting introduced earlier. In the decentralized case, each country minimizes the individual cost functional:

$$J^{\text{MFG}}(u) = \mathbb{E} \left[\int_0^T \left(\frac{q}{2} x(t)^2 + \frac{r}{2} u(t)^2 + \frac{\eta}{2} (x(t) - \bar{x}(t))^2 \right) dt + \frac{q_T}{2} x(T)^2 \right],$$

where the state evolves according to:

$$dx(t) = (ax(t) + bu(t))dt + \sigma dW(t).$$

In contrast, the central planner minimizes the aggregate expected cost of all countries, internalizing the impact each agent has on the overall distribution. Since all agents are symmetric and indistinguishable, the planner's problem becomes one of minimizing the average cost functional:

$$J^{\text{Planner}}(u) = \mathbb{E} \left[\int_0^T \left(\frac{q}{2} x(t)^2 + \frac{r}{2} u(t)^2 + \eta (x(t) - \bar{x}(t))^2 \right) dt + \frac{q_T}{2} x(T)^2 \right],$$

with the same dynamics. Note that in the planner's problem, the externality term $\eta(x - \bar{x})^2$ appears once, whereas in the MFG setting, it is perceived as part of each agent's private cost. Consequently, the planner recognizes that the average deviation term aggregates to the population variance:

$$\mathbb{E}[(x - \bar{x})^2] = \text{Var}(x(t)),$$

and therefore, the planner aims to minimize not only the mean military capital level but also its dispersion across countries. This insight will be reflected in the derivation of optimal policies.

To proceed analytically, we again assume the control policy is linear in state:

$$u(t) = -K(t)x(t) + h(t),$$

and derive the corresponding optimal control laws using the dynamic programming principle. The planner's value function retains the quadratic form:

$$V^{\text{Planner}}(t, x) = \frac{1}{2} \tilde{P}(t)x^2 + \tilde{R}(t)x + \tilde{S}(t),$$

and the planner's HJB equation becomes:

$$\frac{\partial V}{\partial t} + \min_u \left\{ (ax + bu)(\tilde{P}x + \tilde{R}) + \frac{1}{2} \sigma^2 \tilde{P} + \frac{q}{2} x^2 + \frac{r}{2} u^2 + \eta (x - \bar{x})^2 \right\} = 0.$$

Solving for the optimal control in the same way as in the MFG yields:

$$u^{\text{Planner}}(t, x) = -\frac{b}{r}(\tilde{P}(t)x + \tilde{R}(t)).$$

Substituting into the HJB equation and matching coefficients results in the Riccati system for the planner:

$$\begin{aligned}
\dot{\tilde{P}}(t) + 2a\tilde{P}(t) - \frac{b^2}{r}\tilde{P}^2(t) + q + \eta &= 0, \\
\dot{\tilde{R}}(t) + \left(a - \frac{b^2}{r}\tilde{P}(t)\right)\tilde{R}(t) - 2\eta\bar{x}(t) &= 0, \\
\dot{\tilde{S}}(t) + \frac{b^2}{2r}\tilde{R}^2(t) - \frac{\sigma^2}{2}\tilde{P}(t) + \eta\bar{x}^2(t) &= 0.
\end{aligned}$$

The key difference lies in the second equation: the coefficient in front of $\bar{x}$ is now 2η rather than η , as it was in the MFG Riccati system. This discrepancy reflects the planner's internalization of the externality, recognizing that each country's deviation contributes to a collective inefficiency. As a result, the planner's coefficient $\tilde{P}(t)$ is typically smaller than the decentralized coefficient $P(t)$, leading to more moderate control policies.

In the stationary case (infinite horizon, time-invariant), the planner solves the algebraic system:

$$\begin{aligned}
\rho\tilde{P} + 2a\tilde{P} - \frac{b^2}{r}\tilde{P}^2 + q + \eta &= 0, \\
\rho\tilde{R} + \left(a - \frac{b^2}{r}\tilde{P}\right)\tilde{R} - 2\eta\bar{x} &= 0.
\end{aligned}$$

The contrast becomes sharpest when comparing the optimal control laws:

- MFG control: $u^{\text{MFG}}(x) = -\frac{b}{r}(Px + R)$
- Planner control: $u^{\text{Planner}}(x) = -\frac{b}{r}(\tilde{P}x + \tilde{R})$

Since $\tilde{P} < P$ and $|\tilde{R}| < |R|$, the planner prescribes weaker militarization for all countries. This leads to a lower drift rate in the population-level dynamics, and a more concentrated and left-shifted equilibrium distribution $\tilde{p}(x)$.

Quantitatively, we find that:

- The mean military stock $\bar{x}^{\text{Planner}} < \bar{x}^{\text{MFG}}$,
- The variance under the planner is smaller: $\text{Var}^{\text{Planner}}(x) < \text{Var}^{\text{MFG}}(x)$,
- The total expected cost is lower under the planner: $J^{\text{Planner}} < J^{\text{MFG}}$.

This confirms the classic result that decentralized strategies with externalities (such as peer effects) lead to over-provision of strategic goods like military capital. In the MFG setting, each country ignores how its actions raise the incentive for others to militarize, generating an arms race with inflated costs and inefficiencies.

The policy implication is conceptually straightforward: in the absence of coordination, decentralized equilibrium leads to a distorted global outcome characterized by excessive militarization and synchronized strategic risk. However, the geopolitical reality is far more complex. Centralized planning at the global level is politically impossible. There is no supranational authority with the legitimacy or coercive capacity to dictate defense policy across sovereign states. Even multilateral institutions like the United Nations or NATO face deep-seated constraints rooted in power asymmetries, divergent national interests, and historical mistrust.

Consequently, the task of global governance becomes one of approximating the planner's logic through feasible second-best mechanisms. These include bilateral and multilateral treaties, non-proliferation regimes, arms control agree-

ments, and economic incentives or sanctions. While such institutions cannot directly enforce optimal behavior, they can shape expectations, internalize some externalities, and influence the feedback structures that drive military decision-making.

But even this is optimistic. In reality, geopolitical strategy is shaped not only by costs and optimization, but by ideology, misperception, nationalism, and domestic politics. Coordination fails not just because of strategic divergence, but also because of systemic mistrust and opportunism. Arms control efforts are regularly undermined by cheating, ambiguity, or technological shifts that render old agreements obsolete.

Thus, while our model identifies clear distortions and possible remedies, the real world remains far from the ideal of rational coordination. What remains possible, however, is the construction of flexible institutional frameworks that nudge the system toward more stable equilibria—not by solving the game, but by reshaping it at the margins. Effective arms control, in this view, is less about achieving global optimality and more about minimizing catastrophic feedback loops within a fundamentally anarchic order.

7 Multipolar Arms Race: Three Major Powers and a Continuum of Minor Countries

The homogeneous mean field framework developed in earlier sections captures the decentralized arms race dynamics among symmetric countries. However, the real-world international system is far from symmetric. A small number of major powers exert disproportionate influence on global security architecture, while the majority of countries behave as followers, reacting to the strategic environment set by those powers. In this section, we build a two-tier mean field game model in which three major countries—United States, China, and Russia—act strategically and directly influence the equilibrium, while a continuum of minor countries respond to the joint behavior of the major powers and the distribution of peers.

7.1 General Nonlinear MFG System with Major–Minor Asymmetry

Let us denote the major powers by $i \in \{1, 2, 3\}$, representing the United States, China, and Russia respectively. Let $x_i(t) \in \mathbb{R}_+$ be the military stock of major power i at time t , and $u_i(t) \in U_i \subset \mathbb{R}$ be its military expenditure (control variable). The dynamics of each major power follow a controlled stochastic differential equation (SDE):

$$dx_i(t) = f_i(x_i(t), u_i(t), t) dt + \sigma_i(x_i(t), u_i(t), t) dW_i(t),$$

where $W_i(t)$ is a Brownian motion representing geopolitical uncertainty, and the function f_i governs the deterministic drift of the military stock.

Each major power aims to minimize a cost functional that includes: (i) internal costs of military stock and expenditure, (ii) strategic deviation from

other major powers, and (iii) global deterrence measured by deviation from the mean field of minor countries:

$$J_i(u_1(\cdot), u_2(\cdot), u_3(\cdot), \mu(\cdot)) = \mathbb{E} \left[\int_0^T \left(L_i(x_i(t), u_i(t), \bar{x}_M(t), \mu(t), t) \right) dt \right. \\ \left. + G_i(x_i(T), \bar{x}_M(T), \mu(T)) \right],$$

where:

- $\bar{x}_M(t) = \frac{1}{3} \sum x_i(t)$ with $i \in \{1, 2, 3\}$ is the average military stock of three major powers,

- $\mu(t)$ is the distribution of the minor countries' military stocks at time t ,
- L_i and G_i are running and terminal cost functions.

Each minor country is indexed over a continuum $m \in [0, 1] \setminus \{1, 2, 3\}$, and its state $x(t) \in \mathbb{R}_+$ evolves according to:

$$dx(t) = f_m(x(t), u(t), t) dt + \sigma_m(x(t), u(t), t) dW(t),$$

with cost functional:

$$J_m(u(\cdot)) = \mathbb{E} \left[\int_0^T L_m(x(t), u(t), \bar{x}_M(t), \mu(t), t) dt + G_m(x(T), \bar{x}_M(T), \mu(T)) \right].$$

Here, minor players respond only to the joint behavior of major powers (through $\bar{x}_M(t)$) and the population distribution of other minor players (through $\mu(t)$). Their decisions do not directly affect the behavior of the major powers.

7.2 HJB Equation (Minor Players)

Let $V_m(t, x)$ denote the value function for a representative minor country. Then the HJB equation is:

$$\frac{\partial V_m}{\partial t}(t, x) + \inf_{u \in U} \left\{ f_m(x, u, t) \frac{\partial V_m}{\partial x} + \frac{1}{2} \sigma_m^2(x, u, t) \frac{\partial^2 V_m}{\partial x^2} + L_m(x, u, \bar{x}_M(t), \mu(t), t) \right\} = 0,$$

with terminal condition $V_m(T, x) = G_m(x, \bar{x}_M(T), \mu(T))$.

Best Response (Minor Players):

The optimal control is defined pointwise via:

$$u_m^*(t, x) = \arg \min_{u \in U} \left\{ f_m(x, u, t) \frac{\partial V_m}{\partial x} + \frac{1}{2} \sigma_m^2(x, u, t) \frac{\partial^2 V_m}{\partial x^2} + L_m(x, u, \bar{x}_M(t), \mu(t), t) \right\}.$$

FPK Equation (Minor Players):

Given the optimal control $u_m^*(t, x)$, the distribution $p(t, x)$ of the minor countries' states evolves according to:

$$\frac{\partial p}{\partial t}(t, x) = -\frac{\partial}{\partial x} (f_m(x, u_m^*(t, x), t) p(t, x)) + \frac{1}{2} \frac{\partial^2}{\partial x^2} (\sigma_m^2(x, u_m^*(t, x), t) p(t, x)).$$

This is the Fokker–Planck–Kolmogorov (FPK) equation, describing how the probability density of minor players' states evolves forward in time.

McKean–Vlasov Coupling:

The system is closed through the consistency condition:

$$\mu(t) = p(t, x) dx, \quad \text{and} \quad u^m(t, x) = u^m(t, x; \bar{x}_M(t), \mu(t)).$$

The full equilibrium requires solving:

- A finite-player game among the major powers (with their own HJB systems), and
- A mean field fixed point for the minor players' HJB-FPK system, coupled through $\bar{x}_M(t)$.

7.3 Linear-Quadratic-Gaussian (LQG) MFG System

We now specialize the model for explicit derivations. Each major player $i \in \{1, 2, 3\}$ controls:

$$dx_i(t) = a_i x_i(t) dt + b_i u_i(t) dt + \sigma_i dW_i(t),$$

and minimizes:

$$J_i = \mathbb{E} \left[\int_0^T \left(\frac{q_i}{2} x_i^2 + \frac{r_i}{2} u_i^2 + \frac{\eta_i}{2} (x_i - \bar{x}_M)^2 + \frac{\theta_i}{2} (x_i - \bar{x}_\mu)^2 \right) dt \right].$$

Each minor player controls:

$$dx(t) = ax(t) dt + bu(t) dt + \sigma dW(t),$$

and minimizes:

$$J_m = \mathbb{E} \left[\int_0^T \left(\frac{q}{2} x^2 + \frac{r}{2} u^2 + \frac{\eta}{2} (x - \bar{x}_M)^2 \right) dt \right].$$

HJB for Minor Player (LQG):

Guess the value function:

$$V_m(t, x) = \frac{1}{2} P_m(t) x^2 + R_m(t) x + S_m(t).$$

Compute the derivatives:

$$\frac{\partial V}{\partial x} = P_m x + R_m, \quad \frac{\partial^2 V}{\partial x^2} = P_m.$$

Optimal control:

$$u_m^*(t, x) = -\frac{b}{r} (P_m x + R_m).$$

Plug into the HJB equation, and collect terms to get the Riccati system:

$$\dot{P}_m = -2aP_m + \frac{b^2}{r} P_m^2 - q - \eta,$$

$$\dot{R}_m = -\left(a - \frac{b^2}{r} P_m\right) R_m + \eta \bar{x}_M,$$

$$\dot{S}_m = -\frac{b^2}{2r} R_m^2 + \frac{\sigma^2}{2} P_m - \frac{\eta}{2} \bar{x}_M^2.$$

FPK for Minor Players (LQG):

Drift under optimal control is:

$$\mu(x, t) = \left(a - \frac{b^2}{r} P_m(t)\right) x - \frac{b^2}{r} R_m(t).$$

Then the FPK equation becomes:

$$\frac{\partial p}{\partial t} = -\frac{\partial}{\partial x} (\mu(x, t) p) + \frac{\sigma^2}{2} \frac{\partial^2 p}{\partial x^2}.$$

The mean field $\bar{x}_\mu(t) = \int x p(t, x) dx$ feeds back into the major players' cost.

Major Powers' Best Responses:

Each major power solves a 3-player LQ differential game. Their optimal controls take the form:

$$u_i^*(t, x_i) = -\frac{b_i}{r_i} (P_i(t) x_i + R_i(t)),$$

where $P_i(t)$, $R_i(t)$ satisfy Riccati ODEs depending on $\bar{x}_M(t)$, $\bar{x}_\mu(t)$, and the strategies of other major players.

8 Simulations and Interpretation of Strategic Military Equilibrium

This section presents simulations of the heterogeneous mean field game (MFG) system featuring three major powers and a continuum of minor states. We numerically solve the coupled Riccati and Fokker–Planck–Kolmogorov (FPK) equations derived in Section 7 and analyze how the strategic interactions shape the global distribution of military power.

We examine:

1. Time-evolution of optimal military stocks for major powers.
2. Transition of the minor player distribution over time.
3. Stationary equilibrium distributions of all agents.
4. Comparative statics: how changes in key parameters affect global militarization.

We begin by discretizing the system over a finite time horizon $[0, T]$, with $T = 50$, and then we examine the stationary case as $T \rightarrow \infty$.

8.1 Numerical Implementation: Dynamic System of Major and Minor Players

We restate the key dynamic equations in the LQG formulation:

8.1.1 Minor Player Dynamics and Controls

Each minor player follows:

$$dx(t) = \left(a - \frac{b^2}{r} P_m(t)\right) x(t) dt - \frac{b^2}{r} R_m(t) dt + \sigma dW(t).$$

The density $p(t, x)$ evolves according to:

$$\frac{\partial p}{\partial t} = -\frac{\partial}{\partial x} \left(\left[\left(a - \frac{b^2}{r} P_m(t)\right) x - \frac{b^2}{r} R_m(t) \right] p \right) + \frac{\sigma^2}{2} \frac{\partial^2 p}{\partial x^2}.$$

We solve this numerically on a bounded domain $x \in [0, x_{\max}]$, with reflecting boundaries or fast decay conditions. The Riccati equations for (P_m, R_m, S_m) are:

$$\begin{aligned} \dot{P}_m &= -2aP_m + \frac{b^2}{r} P_m^2 - q - \eta, \\ \dot{R}_m &= -\left(a - \frac{b^2}{r} P_m\right) R_m + \eta \bar{x}_M(t), \\ \dot{S}_m &= -\frac{b^2}{2r} R_m^2 + \frac{\sigma^2}{2} P_m - \frac{\eta}{2} \bar{x}_M^2(t). \end{aligned}$$

The mean field $\bar{x}_\mu(t) = \int x p(t, x) dx$ is computed at each time step to influence the major players.

8.1.2 Major Powers' Controls and Dynamics

Each major power $i \in \{1, 2, 3\}$ has:

$$\begin{aligned} dx_i(t) &= a_i x_i(t) dt + b_i u_i(t) dt + \sigma_i dW_i(t), \\ u_i^*(t, x_i) &= -\frac{b_i}{r_i} (P_i(t) x_i + R_i(t)). \end{aligned}$$

The Riccati equations for each major player involve the average behaviors:

$$\begin{aligned}\dot{P}_i &= -2a_i P_i + \frac{b_i^2}{r_i} P_i^2 - q_i - \eta_i - \theta_i, \\ \dot{R}_i &= -\left(a_i - \frac{b_i^2}{r_i} P_i\right) R_i + \eta_i(\bar{x}_M(t) - x_i(t)) + \theta_i(\bar{x}_\mu(t) - x_i(t)), \\ \dot{S}_i &= -\frac{b_i^2}{2r_i} R_i^2 + \frac{\sigma_i^2}{2} P_i - \frac{\eta_i}{2}(\bar{x}_M(t) - x_i(t))^2 - \frac{\theta_i}{2}(\bar{x}_\mu(t) - x_i(t))^2.\end{aligned}$$

This system is solved jointly across the three major players, with feedback coupling via:

- $\bar{x}_M(t) = \frac{1}{3} \sum x_i(t)$, $i \in \{1, 2, 3\}$.
- $\bar{x}_\mu(t)$ from the evolving FPK.

8.2 Simulation Results: Transient Dynamics and Convergence

In a representative simulation calibrated with:

- $a_i = 0.05, b_i = 0.1, q_i = 1, r_i = 1, \eta_i = 0.3, \theta_i = 0.4, \sigma_i = 0.1$,
- $a = 0.03, b = 0.08, q = 1, r = 1, \eta = 0.5, \sigma = 0.1$,
- Initial condition for minor players $p(0, x) = \text{Gaussian}(x; 0.5, 0.1^2)$,
- Initial major stocks: $x_1(0) = 1.2, x_2(0) = 0.9, x_3(0) = 0.8$,

we observe:

- Rapid initial growth in all three major powers' military stocks, with China and Russia increasing faster due to their lower initial values and the presence of competitive pressure.
- Minor players gradually increase their military stocks in response to $\bar{x}_M(t)$, with the distribution $p(t, x)$ shifting rightward and flattening.
- The system stabilizes around $t \approx 25$, with $x_1 \approx 1.6$, $x_2 \approx 1.55$, $x_3 \approx 1.5$, and $\bar{x}_\mu \approx 1.3$.

8.3 Stationary Distribution and Long-Term Strategic Balance

Solving the stationary Riccati system (as $\dot{P} = \dot{R} = 0$):

$$\begin{aligned}\rho P_m + 2aP_m - \frac{b^2}{r} P_m^2 + q + \eta &= 0, \\ \rho R_m + \left(a - \frac{b^2}{r} P_m\right) R_m - \eta \bar{x}_M &= 0,\end{aligned}$$

we find a stationary FPK distribution for minor players of the form:

$$p^*(x) = C \exp\left(\frac{2}{\sigma^2} \left(\frac{1}{2} \left(a - \frac{b^2}{r} P_m\right) x^2 - \frac{b^2}{r} R_m x\right)\right),$$

where C normalizes the distribution over $\mathbb{R}_+$. This distribution is unimodal and right-skewed when $\bar{x}_M$ is high.

In our simulation, the stationary distribution has:

- Mean $\bar{x}_\mu \approx 1.25$,
- Moderate variance, centered near the average major power level.

This reflects a rational convergence of minor players toward the strategic frontier established by the major powers.

8.4 Comparative Statics and Strategic Interpretation

We explore how key parameters affect the equilibrium:

- Increasing η_i (major power rivalry): raises dispersion in major stocks and shifts minor player distributions rightward.
- Increasing θ_i (hegemonic concern): flattens the control policy, leading to more aggressive militarization to dominate minor states.
- Increasing η (peer pressure among minors): reduces variance in the minor distribution but may raise the mean militarization.
- Lowering discount rate ρ : causes all countries to care more about long-term positioning, amplifying the arms race.

The model shows how strategic deterrence, relative positioning, and interdependence collectively drive long-run military stock accumulation—even in the absence of direct conflict.

9 Conclusion and Policy Outlook

This paper developed a rigorous dynamic framework to analyze global arms races using mean field game (MFG) theory. We modeled the strategic behavior of a continuum of minor countries, each optimizing its military investment relative to evolving global norms, while three major powers—the United States, China, and Russia—interacted in a finite-player strategic setting. The result is a coupled, decentralized system in which each country’s optimal decision depends not only on internal costs and constraints but also on the expectations and behavior of others.

We formalized this interdependence using Hamilton–Jacobi–Bellman (HJB) and Fokker–Planck–Kolmogorov (FPK) equations linked through a McKean–Vlasov fixed-point structure. The linear-quadratic-Gaussian (LQG) formulation enabled analytical derivations and exact feedback controls, which we used to simulate both time-dependent transition dynamics and long-run stationary equilibria.

Our simulations reveal that in the absence of regulation, rational strategic behavior produces excessive militarization and coordination on inflated levels of military stock. Minor powers gradually converge toward the global norm established by dominant states, and peer pressure compresses this convergence into synchronized patterns of buildup. These effects are magnified in high peer-pressure environments, where the incentive to conform dominates internal preferences for moderation. In this sense, decentralized equilibrium is not merely inefficient—it is path-dependent and self-reinforcing, with the potential to entrench militarized standoffs even without open conflict.

We also compared this decentralized outcome with the planner’s optimal trajectory and showed that centralized coordination would lead to less costly and more stable equilibria. However, this idealized comparison exposes a deeper tension: the policy implication is conceptually clear—absent coordination, the equilibrium is distorted—but in real-world geopolitics, centralized global plan-

ning is politically impossible. There is no supranational authority with the legitimacy, neutrality, or enforcement power to dictate defense policy across sovereign states. Even when cooperation is theoretically beneficial, the international system lacks the institutional scaffolding and trust necessary to enforce compliance.

As a result, global governance must proceed not through ideal coordination, but through second-best approximations. Treaties, deterrence frameworks, arms control regimes, and economic incentives can all approximate aspects of the planner’s logic—if imperfectly—by embedding the right incentives and shaping the strategic environment. These instruments help align national objectives with system-wide stability by modifying the feedback loop rather than solving it. But we must also acknowledge the limits of this approach. Realpolitik is not governed solely by optimization and cost minimization—it is shaped by fear, nationalism, historical grievances, and domestic political pressures. Many of the mechanisms we model as continuous and rational unfold in reality through crises, bluffs, and cascading miscalculations.

Therefore, while our framework offers clear analytical insights into how military equilibria form and persist, it also serves as a caution: decentralized rationality in an anarchic system can lead to collectively irrational outcomes. In this environment, even partial institutional success—through verifiable caps, transparency mechanisms, or normative soft power—can dramatically reduce global risk. Effective policy, then, is not about achieving optimal control, but about minimizing catastrophic feedback loops in a world where coordination is fragile, and the consequences of failure are severe.

Future research can extend this framework in several directions: incorporating asymmetric information, modeling alliance networks, analyzing stochastic regime shifts, and calibrating empirical data from historical arms races. Ultimately, the theory of mean field games offers not only a mathematical toolset but a new lens for understanding how rational agents trapped in irrational systems might still find strategies for peace.

References

- Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., & Moll, B. (2022). Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *Review of Economic Studies*, 89(1), 45–86.
- Cardaliaguet, P. (2012). Notes on Mean Field Games. Retrieved from <https://www.ceremade.dauphine.fr/~cardalia/MFG20130420.pdf>
- Cardaliaguet, P., Delarue, F., Lasry, J.-M., & Lions, P.-L. (2019). *The Master Equation and the Convergence Problem in Mean Field Games*. Princeton University Press.
- Carmona, R., & Delarue, F. (2018). *Probabilistic Theory of Mean Field Games with Applications I: Mean Field FBSDEs, Control, and Games*. Springer.
- Carmona, R., & Delarue, F. (2018). *Probabilistic Theory of Mean Field Games with Applications II: Mean Field Games with Common Noise and Master Equations*. Springer.
- Fearon, J. D. (1995). Rationalist Explanations for War. *International Organization*, 49(3), 379–414.
- Gong, L., & Zou, H.-F. (2003). Military spending and stochastic growth. *Journal of Economic Dynamics and Control*, 28(1), 153–170.
- Huang, M., Caines, P. E., & Malhamé, R. P. (2007). Large-population cost-coupled LQG problems with nonuniform agents: Individual-mass behavior and decentralized ε -Nash equilibria. *IEEE Transactions on Automatic Control*, 52(9), 1560–1571.
- Huang, M., Malhamé, R. P., & Caines, P. E. (2006). 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.
- Jervis, R. (1976). *Perception and Misperception in International Politics*. Princeton University Press.
- Kydd, A. H. (2005). *Trust and Mistrust in International Relations*. Princeton University Press.
- Lasry, J.-M., & Lions, P.-L. (2006). Jeux à champ moyen. I – Le cas stationnaire. *Comptes Rendus Mathématique*, 343(9), 619–625.
- Lasry, J.-M., & Lions, P.-L. (2006). Jeux à champ moyen. II – Horizon fini et contrôle optimal. *Comptes Rendus Mathématique*, 343(10), 679–684.
- Lasry, J.-M., & Lions, P.-L. (2007). Mean Field Games. *Japanese Journal of Mathematics*, 2(1), 229–260.
- Powell, R. (2006). War as a Commitment Problem. *International Organization*, 60(1), 169–203.
- Schelling, T. C. (1960). *The Strategy of Conflict*. Harvard University Press.
- Zou, H.-F. (1995). A dynamic model of capital and arms accumulation. *Journal of Economic Dynamics and Control*, 19(1–2), 371–393.
- Zou, H.-F. (1997). Dynamic analysis in the Viner model of mercantilism. *Journal of International Money and Finance*, 16(4), 637–651.